Permissioned vs Permissionless Blockchain Platforms: Tradeoffs in Trust and Performance*

Yannis Bakos† Hanna Halaburda‡

March 15, 2022
— Working Paper – Comments Welcome —

Abstract

The prevalence of blockchain technology is increasingly evidenced in blockchain-based platforms, which can employ either permissioned blockchains (typical in supply-chain applications such as IBM Food Trust) or permissionless blockchains (common in Decentralized Finance platforms such as Compound). It is generally agreed that permissioned blockchains can improve on the operational cost and performance of permissionless blockchains, but it is usually assumed that this improvement comes at the cost of transaction security, especially in low-trust environments. We develop a model of transaction safety in permissioned and permissionless blockchains to study this tradeoff and find that in several settings there may be no tradeoff at all. With a minimal level of trust in the blockchain operators and the supporting institutions, well-designed permissioned blockchains can offer both higher operational efficiency and higher transaction security. While this minimal trust in the "system" is essential to the functioning of permissioned blockchains, it is also inherent in most business relationships, making permissioned blockchains well suited for enterprise applications of the technology. We explore the implications of our analysis for the design of permissioned blockchains, such as the reputation or bonding implications for their validators. This analysis is directly relevant to blockchain-based platforms in selecting an appropriate technology.

*We thank Agostino Capponi, David Cerezo, Guillaume Haeringer, Evgeny Lyandres and Maher Said for helpful comments and suggestions.

†Stern School of Business, New York University; email: bakos@stern.nyu.edu

‡Stern School of Business, New York University; email: hhalaburda@gmail.com

1 Introduction

A well-designed blockchain is essentially a shared decentralized ledger that is trusted by all its users. The integrity of this ledger is based on a combination of cryptography (such as hash functions and digital signatures) and mechanisms for distributed consensus (such as the mining of blocks in the Bitcoin blockchain), which enables all participants to agree on a unique immutable version of the "ground truth."

The prevalence of blockchain technology is increasingly evidenced in blockchain-based platforms that use blockchains to deliver value to their users. Blockchain participants, often referred to as the nodes of the blockchain network, are typically categorized as users and validators, with the latter responsible for maintaining the integrity and consensus of the blockchain; the two categories are not mutually exclusive and nodes often participate in both roles. In the context of blockchain-based platforms, blockchain users would be comprised by the participants in the "sides" of the platform, while the validators would safeguard the operation of the blockchain.

Blockchains can be categorized as permissioned or permissionless, based on whether participation in a certain role requires going through an qualification process that is typically controlled by an entity with operational responsibility for the blockchain, or is open to all comers as long as they satisfy the requirements of the applicable protocol. A blockchain-based platform can employ either permissioned blockchains (as is typical in supply-chain applications such as IBM Food Trust) or permissionless blockchains (which is common in Decentralized Finance platforms such as Compound). While a permissioned blockchain may require qualification for validators only, or for both roles of user and validator, *in this work we focus on whether a blockchain is permissioned in terms of its validators.*

Another important aspect of blockchain settings is "trust." We characterize the setting of a blockchain as "trusted" when enforcement mechanisms external to the blockchain can be used to induce participants to follow the blockchain protocol; for instance, this could be with contractual obligations that can be enforced by courts or arbitrators, or with monetary penalties implemented by escrow agents or other institutions, or with reputation mechanisms that would punish deviations with negative reputation. In the absence of such mechanisms outside the blockchain, we characterize its setting as "trustless." This characterization need not be binary as the level of incentives and penalties provided by a particular setting can vary in a continuous fashion. Trust is not associated with the participants, but with the

setting; participants are simply assumed to be rational and pursue their own self interest.

The choice of permissioned or permissionless validators is an important design decision for blockchain-based platforms. It is generally agreed that permissioned validation offers improved operational cost and performance compared to permissionless blockchains, but it is usually assumed that this improvement comes at the cost of transaction security, especially in low-trust environments. We develop a model of transaction safety in permissioned blockchains based on the qualification requirements for their validators that we use to study this tradeoff.

We find that in several settings there may be no tradeoff at all. That is because with a minimal level of trust in the blockchain operators and the supporting institutions, well-designed permissioned blockchains can offer both higher operational efficiency and higher transaction security. While this minimal trust in the "system" is essential to the functioning of permissioned blockchains, it is also inherent in most business relationships, making permissioned blockchains well suited for platform applications of the technology. Our analysis is directly relevant to blockchain-based platforms in their selection of an appropriate blockchain technology; for instance, we can characterize the type of environments where a platform may still prefer to employ permissionless blockchains, despite the likely penalty in operational cost and performance. Furthermore, if a platform employs a permissioned blockchain, our results have direct implications for the selection of validators based on their reputation or ability to post a bond.

1.1 Permissioned vs. permissionless blockchains

Permissionless blockchains achieve consensus via a decentralized protocol applied across a theoretically unlimited set of participants or nodes. Permissionless protocols do not require the nodes to reveal their identities beyond a pseudonymous identifier. Moreover, participants can freely acquire new identifiers, dispose of old ones and control multiple identifiers at any point of time. Permissionless blockchains thus cannot assume any level of trust in their setting, as their participants remain anonymous and beyond reach of enforcement mechanisms external to the blockchain, such as courts or other institutions. *Permissioned* blockchains, in contrast, require that the validator nodes that can update the blockchain must be approved before assuming that role, and some permissioned blockchains place approval requirements on user nodes as well. The approval process usually requires nodes to

provide full transparency regarding their identities, and as a result consensus mechanisms for permissioned blockchains typically assume at least some level of trust in the underlying institutional setting.

Arguably the major innovation in permissionless blockchains was devising a mechanism to provide trust in the consensus ledger without requiring any trust in the institutional setting, which means that all incentives to follow the consensus mechanism must be provided within the protocol. Bitcoin's consensus mechanism famously was the first successful such protocol, and it requires validators (known as "miners") to show "proof of work" (PoW) as they compete to validate new blocks, combined with a block reward for the validating miner. The innovation in permissioned blockchains is more subtle; given a certain level of trust in the institutional setting, instead of depending on a single party, such as the owner of a distributed database, operation of permissioned blockchains can be delegated to a community of permissioned validators, such as the Trust Anchors in the IBM Food Trust blockchain, who mostly are large established food industry players,[1] or the validators in the XRP Ledger blockchain run by Ripple, who include independent trusted organizations such as MIT.[2]

It is generally agreed that the open participation in permissionless blockchains that requires them to be able to operate in a completely trustless environment, being robust even under the assumption that each participant would undermine the integrity of the system if that results in short term benefit, comes at the cost of efficiency and performance limitations. Proof-of-work blockchains such as Bitcoin require costly mining in terms of the computing hardware resources and the energy consumed, and these costs become even larger when including externalities such as environmental and climate impact. Bitcoin's mechanism has performance limitations as well, such as limits in the feasible size and frequency of new blocks and thus in transaction capacity, which result from the need to accommodate a large and unpredictable number of diverse miners.

There have been many attempts to address the limitations of permissionless blockchains with improved technology, e.g., via more frequent blocks (as in Ethereum) or side-transactions (as in the Lightning Network) or increased block size (as in Bitcoin Gold), and there has been a lot of effort to develop alternatives to proof-of-work, such as proof-of-stake, that avoid the environmental impact and other externalities. While these attempts have had different degrees of success, mechanisms able to reach consensus in permissionless blockchains and at

[1] see https://www.ibm.com/blockchain/solutions/food-trust/food-industry-technology
[2] see https://ripple.com/insights/xrp-ledger-decentralizes-expansion-55-validator-nodes/

the same time preserve transaction security have remained costly.

By contrast, in permissioned blockchains where the number of validators is typically limited,[3] and these validators can be induced to behave according to the protocol based on enforcement mechanisms outside the blockchain, such as legal contracts or reputation, consensus can be established with high performance and efficiency. It is typically assumed, however, that this operational efficiency and higher performance of permissioned blockchains comes at the cost of reduced transaction safety, in the sense that transactions are more vulnerable to being compromised by the operators and validators of the permissioned blockchains.[4]

1.2 Contribution and overview of our results

In this paper we study when permissioned blockchains are more desirable than permissionless, taking into account both operational efficiency and transaction safety. This is an increasingly important question as blockchains are more widely adopted in business settings. We develop a framework to compare transaction safety in permissionless and permissioned blockchains so that we can study this assumed tradeoff. We find that there may be no tradeoff at all as well-designed permissioned blockchains can offer both higher operational efficiency *and* higher transaction safety as long as there is a minimal level of trust in the underlying institutional setting.

Our setting illustrates that permissionless consensus mechanisms are inherently costly as validators cannot be punished outside the blockchain due to their anonymity; instead validators must incur upfront costs to participate in the validation of each block (e.g., via resources consumed for the proof-of-work or the opportunity cost of the assets staked in proof-of-stake) and are rewarded upon successfully validating a block. Punishments are implemented by withholding the block rewards if the validators deviate from the protocol;

[3] E.g., Ripple's XRP Ledger has 50-200 validators *vs.* hundreds of thousands and possibly a million miners for Bitcoin in early 2021–see https://markets.businessinsider.com/news/currencies/bitcoin-miners-bln-earnings-how-they-make-money-transactions-2021-2

[4] This tradeoff is referenced in comparisons of permissioned and permissionless blockchains, e.g.,
https://www.blockchain-council.org/blockchain/permissioned-and-permissionless-blockchains-a-comprehensive-guide/
https://searchcio.techtarget.com/tip/Permissioned-vs-permissionless-blockchains-Key-differences
https://medium.com/@akadiyala/nuances-between-permissionless-and-permissioned-blockchains-f5b566f5d483
https://101blockchains.com/permissioned-vs-permissionless-blockchains/
https://freemanlaw.com/permission-and-permissionless-blockchains/

free entry of validators and lack of punishments means that these rewards at equilibrium will be dissipated by the upfront costs. As a result, the validation costs in permissionless blockchains are incurred *during the normal operation of the consensus mechanism*. We show that in our setting the level of the block reward, and hence the operating cost of the permissionless consensus mechanism, determines the level of transaction security. Since block rewards are typically denominated in cryptocurrency, this implies that given the parameters of its protocol, the exchange rate for the cryptocurrency in a permissionless blockchain determines its transaction security.

Permissioned blockchains have no comparable costs during the normal operation of the consensus mechanism, since the known identity of the validators allows punishments for deviations from the protocol to be imposed ex-post and *outside the blockchain*, such as via a fine or negative reputation. Thus permissioned blockchains do not incur the costs and limitations of permissionless consensus mechanisms, but require a certain level of trust in the underlying institutional setting. We show that if that condition can be satisfied, permissioned blockchains can also offer higher transaction safety than permissionless ones. In most business relationships the participants trust the institutional setting within which they operate, require accountability for the actions of their counterparties, and have a stream of interactions that allows them to build a certain reputation.[5] Thus we believe that our results will apply in most business applications of blockchain, implying that permissioned blockchains dominate permissionless in these settings.

Our analysis also addresses another important question: Given that permissioned blockchains require a certain level of trust based on the underlying institutional setting, why is blockchain necessary at all, instead of relying on a (possibly distributed) database for the ledger, operated by a single party? We show that relying on multiple validators for maintenance of the ledger allows permissioned blockchains to achieve higher security even at lower levels of trust that any single participant can be induced not to deviate from the protocol. In the popular blockchain parlance, a distributed ledger allows for distributed trust in an environment where none of the potential validators has enough trust to run the ledger singlehandedly.

In our analysis we quantify trust in a validator as the expected penalty that the valida-

[5]For instance accountability of the parties involved in a transaction is important in supply chains and thus permissioned blockchains like the IBM Food Trust Blockchain, where the transparency in identities enables accountability, are appropriate for such settings.

tor will incur by "misbehaving," or taking actions not allowed by the blockchain protocol, such as validating inappropriate transactions, e.g., allowing double spending, or executing actions prescribed in smart contracts without satisfying all the specified preconditions. As stated earlier, these penalties depend on enforcement mechanisms external to the blockchain, as permissioned validators are identifiable outside the blockchain, and are thus subject to penalties imposed by the institutional setting. Trust is therefore a continuous variable in our analysis, and the higher a validator's expected penalty for misbehavior, the higher is the trust that validator will not deviate from the consensus protocol.

The expectation component of the punishment captures the probability that the punishment will be indeed enforced if misbehavior is detected, i.e., whether the validator will be held accountable. The accountability can be enforced through courts or arbitrators (in case of monetary penalty outside the blockchain), through publicizing the misbehavior (in case of reputation loss), by being expelled from the blockchain and thus foregoing future benefits, or by a combination of the above. The key difference compared to permissionless blockchains is that such penalties *are off the equilibrium path*, i.e., do not take place in the normal operation of the protocol, resulting in the lower operating costs of permissioned blockchains.

Transaction safety can be quantified as the largest safe transaction for a particular blockchain. The creation of Bitcoin has shown that it is possible to achieve substantial transaction safety in a permissionless system, i.e., it is possible to create a shared trusted ledger without any trust in the participating validators or users, and our analysis allows us to characterize in which settings it will be optimal to opt for such permissionless systems. Furthermore, for settings where permissioned systems will perform better, we suggest design principles that will increase transaction safety.

The transaction safety provided by a permissioned blockchain depends on the penalties that individual potential validators may expect if they were to deviate from the protocol. We show that a larger number of validators and higher expected penalties for these validators result in a higher transaction safety. Interestingly, transaction safety in a permissioned blockchain is not determined by the validators with the highest expected penalty but by the validators with the *lowest* expected penalty, because these become the weakest link in the validation protocol. The designer of a permissioned blockchain needs to trade-off a higher number of validators (which increases transaction safety) and the trust in the weakest link of these validators. Depending on the distribution of the levels of expected penalty among the potential validators, a well-designed permissioned blockchain may offer higher or lower

levels of transaction safety compared to a permissionless blockchain.

The rest of the paper is organized as follows: Section 2 reviews the relevant literature. Section 3 develops a model for the safety of transactions in a permissionless blockchain. Section 4 develops a model for the safety of transactions in a permissioned blockchain. Section 5 compares transaction safety in permissionless and permissioned blockchains. Finally, Section 6 offers some conclusions and areas for future research.

2 Related Literature

There is an emerging literature on the economic incentives in permissioned blockchains. Cao, Cong, and Yang (2018) study how permissioned blockchains can support auditing. Narang, Byali, Dayama, Pandit, and Narahari (2019) propose design principles for permissioned blockchains to facilitate reviews and unbiased information exchange in certain business-to-business settings. Pun, Swaminathan and Hou (2018) find that blockchains can be used to prevent counterfeiting.

Gans and Gandal (2019) look at the performance of permissionless and permissioned blockchains. They find that under certain conditions (namely a fixed block reward), achieving security on permissioned blockchains is less expensive than the same level of security on permissionless blockchain. The advantage, however, disappears when the block reward is endogenous. The main difference between Gans and Gandal (2019) and our model is that Gans and Gandal assume that even for permissioned blockchains the identity of the nodes is not known outside of the blockchain environment. Hence, the incentives for security need to fully come from the design of the blockchain protocol, which requires the nodes to bear the cost up-front and be compensated later, just as in a permissionless blockchain. That leads to both systems being costly to maintain. In contrast, our results show that such cost in a permissioned system is negligible.

Auer, Monnet and Shin (2021) look at the use of a permissioned blockchain for keeping track of credit worthiness in a credit economy. Their analysis focuses on the coordination problem for an exogenously determined set of validators who need to conduct costly validation and reach consensus by sending costly messages. Similarly to our setting, they assume that validators are identified outside of the blockchain network, and thus can be punished by exclusion from future participation if misbehavior is detected. We look at the more general problem of the relationship between validator incentives and transaction security, the

selection of validators that achieves highest security for a given cost, and compare the cost and security to permissionless blockchains.

Amoussou-Guenou, Biais, Potop-Butucaru, and Tucci-Piergiovanni (2019) study the incentive compatibility of consensus mechanisms in permissioned blockchains. They focus on the coordination game in a setting where sending messages is costly, and the validators are only rewarded when sufficiently many other validators send the messages needed to achieve consensus.

Several papers study the use of blockchains within supply chains. Babich and Hilary (2020) discuss key strengths and weaknesses of blockchains in a supply chain context, and propose future research directions. Blaettchen, Jagmohan, Ratakonda, and Franceschini (2020) conduct a simulation analysis to quantify the costs and benefits of blockchain technology applied to food supply chains. Chod, Trichakis, Tsoukalas, Aspegren, and Weber (2019) study blockchains as a technology that facilitates inventory signaling because of the ability to enable transparency.

Cui, Gaur, and Liu (2020a) examine how blockchain technology can affect competing firms purchasing from a supplier with limited capacity. In their setting the blockchain provides a mechanism to share information that reduces the incentive to over-order or under-order when supplier capacity is respectively small or large. Cui, Hu, and Liu (2020b) study the implications of blockchain-enabled traceability across serial and parallel supply chains. Iyengar, Saleh, Sethuraman and Wang (2020) analyze a setting with consumers that can access information stored on the blockchain and examine associated welfare implications.

There is an emerging literature exploring economic characteristics of Bitcoin's blockchain. For instance Cong, He and Li (2019) study the concentration of computational power in mining pools; Lehar and Parlour (2020) and Malik, Aseri, Singh and Srinivasan (2021) study the ability of mining pools to exert pricing power in Bitcoin transaction fees; Huberman, Leshno and Moallemi (2021), Basu et al (2019), Easley, O'Hara and Basu (2019) and Hinzen, John and Saleh (2019) study the relationship between transaction congestion and transaction fees.

There is also a literature related to our study of transaction safety.Budish (2018) and Chiu and Koeppl (2017) study the vulnerabilities of large transactions and the ability to protect them from double spending attacks; Pagnotta (2021) studies the relationship between Bitcoin price and the computational power dedicated to mining, which affects transaction safety; Prat and Walter (2019) offer a theoretical equilibrium model relating the exchange

rate of Bitcoin to the computational power of the Bitcoin network, and empirically calibrate it. Ebrahimi, Routledge and Zetlin-Jones (2019) study the role of the economic incentives to miners (i.e., block rewards) in determining the resilience of blockchain ledgers in averting double-spend attacks; Halaburda et al (forthcoming) offers a review of these vulnerabilities.

Our work is related to the above literature in that we offer an economic analysis of blockchains with a focus on the concern about the integrity of the blockchain and the safety of transactions against attacks. Our contribution is the development of a framework for comparative analysis of permissioned and permissionless blockchains in terms of the transaction safety, and the finding that under conditions likely to be satisfied in a business relationship, permissioned blockchains can offer both higher efficiency and higher transaction safety.

3 Transaction Safety in Permissioned Blockchains

The key distinction between permissioned and permissionless blockchains is that in the former the identities of the permissioned nodes can be established outside of the blockchain, and thus these nodes can accumulate reputation and be subject to contractual enforcement outside the blockchain ecosystem. Our focus is the safety of blockchain transactions, and thus we analyze blockchains with permissioned validating nodes without distinguishing whether the non-validating participants are permissioned or not.

3.1 A Model of Permissioned Blockchains

Strategic behavior in permissionless blockchains is well understood, and the literature has converged to a class of models similar to the one we develop in the next section as the benchmark for our comparative analysis. There are few formal approaches to modeling permissioned blockchains, however, let alone ones that allow for a comparison with permissionless systems, which is the goal of our analysis in this section.

We consider a permissioned blockchain with n validating nodes that we index with $i = \{1, \ldots, n\}$; these nodes are permissioned in the sense that they can be identified outside of the blockchain. We assume that all validating nodes have equal weight as validators, which is typical for permissioned blockchains presently in use, and we denote by $\sigma(n)$ the number of nodes that need to validate a transaction in order for that transaction to be included in the consensus blockchain.

Transaction validation in permissioned blockchains is typically an efficient process verifying that a transaction satisfies the requisite conditions and then posting it on the blockchain; any related operating costs are likely to be minimal and usually reimbursed by the blockchain owner or covered by a per transaction fee. Accordingly, we assume without loss of generality, a zero marginal cost to the validator to process additional transactions.

The validators, however, can also engage in attacks, such as a double spending, where the transfer of already spent funds is validated while the original spending of these funds is expunged from the blockchain, thereby depriving the original recipient of access to these funds. Multiple validators need to coordinate in order to conduct a successful attack, unless $n = 1$, i.e., the system is fully centralized. While an attack almost certainly would be detected by its target, it may be hard for the affected party to provide sufficient evidence to prove occurrence of the attack. In other words, it is possible that a successful attack may occur without "public" detection.

Since the validating nodes have known and established identities, they can be required to enter into contracts outside the blockchain, agreeing, for instance, that participation in a publicly detected attack would incur a certain penalty or forfeit a bond posted outside the blockchain. Validators may also have established reputations outside the blockchain, which could be negatively affected if they are discovered to participate in an attack. On the other hand, the fact that validators have known identities allows them to enter into side agreements in the form of either formal or relational contracts that could facilitate the coordination required for multiple validators to collude and carry an attack. While nodes in permissionless blockchains may also contract based on their blockchain identities and mechanisms like smart contracts, the ability to identify the validators in permissioned blockchains outside of the blockchain allows punishment via negative reputation feedback as well as institutional enforcement mechanisms such as the courts. Permissioned blockchains thus allow a wider set of means to incentivize desired behaviors from their validators.

We model this ability to impose punishments outside the blockchain by assuming that each node i is subject to punishment P_i if it is publicly detected that the node participated in an attack. This punishment could come from a combination of fines, forfeiting future gains from participation, or reputational damage and is specific to each node as individual nodes will have different limits to their potential liability, depending on their reputation and financial resources. We model users' trust in the system with the parameter τ: if an attack is detected, we assume that nodes that participate in the attack will be punished (at their

corresponding P_i's) with probability $\tau \in [0, 1]$. This parameter τ can be thought as reflecting the trust in the system and its institutions, in the sense that it reflects the belief that agreed punishments will actually be carried out when they are called for.

We analyze next the safety of transactions in permissioned blockchains given that validating nodes have known identities and thus can contract to assist an attack, but can also be punished if discovered.

3.2 Attacks on Permissioned Blockchains

An attack in a blockchain can be either a double spending, which we mentioned earlier, or preventing the recording of a valid transaction, such as blocking a confirmation of delivery for goods that were actually received. Let V be the value the attacker obtains from a successful attack.[6] In order to carry a successful attack, the attacker will need $\sigma(n)$ validating nodes to collude in order to include an invalid transaction such as a double spend, and $n-\sigma(n)+1$ validating nodes to collude in order to block a valid transaction. Let $N(\sigma(n))$ be the smallest number of nodes needed to collude in order to compromise the integrity of the blockchain by either type of attack. Then $N(\sigma(n)) = \min\{\sigma(n),\, n - \sigma(n) + 1\}$, and for simplicity, we will denote it by N.

We assume that an attack will be publicly detected with probability f and that the total benefit to the attacker V can be contractually distributed among all the nodes that participate in the attack. For the attack to be successful a minimum of $\sigma(n)$ or $n-\sigma(n)+1$ nodes must participate, depending on the whether the attack is a double spend or blocking a valid transaction. If validation of a transaction requires a simple majority of the validating nodes, and assuming an odd n, the required number of participating nodes is the simple majority for either type of attack: $N = \sigma(n) = \frac{n+1}{2}$.

The expected cost for each node i that participates in the attack is $f\tau P_i$ as if the attack is discovered the punishments P_i will be imposed with probability τ. Thus, when participation in an attack can be contracted, it is worth to conduct one when $V > \sum_{i\in B} f\tau P_i = f\tau \sum_{i\in B} P_i$, where B is the set of N nodes participating in the attack. Therefore, the transactions of such high values are *vulnerable* to attack. Users should avoid such transactions. The threshold designating *vulnerable* transactions in the system defines *resiliency* the system. In order to characterize the resiliency of the permissioned blockchain, we note that

[6]For simplicity in the exposition we assume that the attack will directly benefit only one of the nodes.

the leader of the attack is better off by selecting co-attackers with a low cost, i.e., B can be assumed to be the set of N nodes with the lowest P_i's. This leads us to the following result.

Lemma 1 *In a permissioned blockchain with contractible participation in an attack, transactions of value greater than V_p are* vulnerable *to attack, where $V_p = f\tau \sum_{i\in B} P_i$ and*

$$B = \arg \min_{S_N \subset \{1,\ldots,n\} \text{ s.t. } |S_N|=N} \sum_{i\in S_N} P_i.$$

That is, resiliency *of a permissioned blockchain with contractible attacks is V_p.*

One implication of Lemma 1 is that the safety threshold depends on the "lowest cost" coalition that the attacker can assemble to execute an attack, which will consist of the N validators that face the least individual penalties if they are discovered. A related implication is that if a sufficient number of validators can bear sufficiently high penalties, then V_p can be as high as needed to guarantee the safety of any level of transaction values.

3.3 Permissioned Blockchain Design and Transaction Safety

For a given trust level τ, the number n of validators in a permissioned blockchain, their corresponding liabilities P_i if they are discovered to participate in an attack, and the consensus mechanism for the blockchain, determine the value threshold V_p above which a transaction is not safe from attack. Increasing the size of n can either increase or decrease that threshold, as illustrated in the following example.

Example. Consider an environment with 15 potential validating nodes where five of them, $i = 1, \ldots, 5$, can be penalized to some amount $P_i = X$, and the remaining ten to a smaller amount $\frac{X}{3}$. Assume that the permissioned blockchain uses simple majority rule as its consensus mechanism, i.e., $\sigma(n) = \frac{n+1}{2}$. Moreover, assume that $\tau = 1$. Then in a permissioned blockchain with nodes $i = 1, \ldots, 5$ as validators, and thus $n = 5$, three nodes must collude to carry an attack and the expected cost of such an attack would result in $V_p(n = 5) = f\,\tau\,3X$. On the other hand, if all 15 nodes participate as validators, a majority of eight is needed to carry out an attack. The lowest cost coalition would consist of eight nodes with potential liability $\frac{X}{3}$, and the expected cost of such an attack would result in $V_p(n = 15) = f\,\tau\,\frac{8}{3}X < V_p(n = 5)$.

For a given n and consensus mechanism, the blockchain has the highest safety if it employs as validators the n nodes with the highest P_i's (e.g., with the highest reputation to lose, or the deepest pockets.) Whether increasing the number of validator nodes n itself will increase transaction safety or not depends on the distribution of available P_i's among the potential validators. In the example above, if the remaining ten nodes faced liability of $\frac{X}{2}$ if discovered to participate in an attack instead of $\frac{X}{3}$, then the blockchain with all 15 nodes would be safer than the blockchain with the five nodes $i = 1, \ldots, 5$.

3.4 Non-contractible Attacks

If participation in an attack cannot be contracted, e.g., because such contracts are not enforceable, then attacks can be easily prevented if they require two or more validators to participate in order to succeed. This follows from our assumption that the benefits from an attack accrue to a single participant because the rest of the validators cannot be compensated for the risk of being discovered and penalized for participating in an attack.

Lemma 2 *If participation in an attack is not contractible, $P_i > 0$ for all i, and $N \geq 2$, there will be no attack for transactions of any value V.*

Proof. Participating in an attack incurs expected cost of $f\tau P_i$ for each validator in the attacking coalition. If the attack is non-contractible, the node that directly benefits cannot commit to compensate any other node for incurring the costs of the attack. Thus, no validator other than the one directly benefiting from the attack would join. When $N \geq 2$, the benefiting party cannot single-handedly execute a successful attack. ■

4 Transaction Safety in Permissionless PoW Blockchains

In this section we model the safety of transactions on a permissionless blockchain; our results allow us to compare our setting with the literature, and provide a benchmark against which to evaluate the safety of permissioned blockchains.

Permissionless blockchains like Bitcoin and Ethereum do not restrict the number or identity of the participating nodes; thus we do not know the real number of nodes as each can control an unlimited number of identities, and these nodes normally cannot be identified

outside of the blockchain.[7] These blockchains rely on an associated cryptocurrency (e.g., bitcoins or ether) to incentivize their validators (miners), who receive a block reward for each block they validate and remains on the consensus blockchain.

In this section, we focus our analysis on consensus mechanisms that depend on Proof of Work (PoW), which has been the oldest and the most common consensus mechanism for permissionless blockchains. We discuss Poof of Stake mechanism in the next section. Miners in PoW blockchains receive block rewards in proportion to the computing power they control.

4.1 A Permissionless Blockchain Model

We consider a permissionless blockchain with miners that employ aggregate computing power C in operating the consensus mechanism. We denote with μ the cost (in fiat currency) of operating a unit of this computing power.

We use R to denote the block reward in the corresponding cryptocurrency, and x to denote the associated exchange rate to fiat currency.[8] Thus a miner that controls c units of computational power will receive an expected reward of $\frac{c}{C}Rx$. The miner faces cost $c\mu$.[9] In permissionless blockchains there is free entry into mining, and thus the investment in C rises to the point that[10]

$$\frac{c}{C}Rx = c\mu \iff \frac{1}{C}Rx = \mu.$$

Thus the total cost of mining each block is $C\mu = Rx$.[11]

Permissionless blockchains typically impose waiting periods before block rewards can be spent, to incentivize cooperative behavior by miners. We denote with w the number of blocks that must be added to the blockchain before a miner can spend a block reward earned for mining a specific block. This is a typical restriction in real world blockchains, known as the *block maturation time* or *confirmation delay*, and it is meant to safeguard against spending the block rewards for blocks in forks of the blockchain that do not become part of

[7] Blockchain participants, however, can demonstrate that they control a node on the blockchain by proving that they are in possession of that node's private cryptographic key.

[8] We do not consider transaction fees but the block reward R can be thought as inclusive of expected transaction fees in the block.

[9] The cost $c\mu$ does not depend on whether the transactions in a block are valid; in other words, there is no additional cost, at the validation stage, for validating a block that is part of an attack.

[10] Note that if x increases, C will also increase (cf. Prat and Walter (2019) and Pagnotta (forthcoming)).

[11] Similar derivations can be found in Budish (2018), Gans and Gandal (2019), Leshno and Strack (2020), and Halaburda et al (forthcoming).

the consensus blockchain.[12]

4.2 Attacks on Permissionless Blockchains

In the above setting an attack would typically consist of a double spend; specifically, the attacker would initiate a transaction with a payment that is intended to elicit a certain action by the recipient(s), such as the provision of a good or service or a follow-up payment transaction that would benefit the attacker. We denote the value of this benefit to the attacker with V. Once the good, service or follow-up transaction have been irrevocably provided, the attacker will replace the block containing the original transaction with an alternative block that contains instead a transaction transferring the payment to a recipient controlled by the attacker. The attack will be successful if this substitution succeeds in becoming part of the consensus blockchain, thus cancelling the original transaction in the attack, and also the attacker manages to convert both the cryptocurrency used in the double spend attack and the proceeds obtained as a result of the attack to value outside the attacked blockchain.

It is also possible to have an attack where a valid transaction is prevented from being recorded on the blockchain for a number of blocks, for instance a confirmation of delivery may be excluded from blocks generated under control of the attacker, and that can benefit the attacker in ways similar to a double spend by making it possible to delay, avoid or reverse payment for goods actually received. Such an attack would require the same resources as a double spend attack, and thus we do need to consider this as a separate case.[13]

While an attack will almost certainly be detected by its target, it may take more time for the attack to be recognized by the blockchain ecosystem, or it may not be recognized altogether. By design there is no authority responsible to coordinate a response; the orphaned chain with the original transaction in a double spend may become inaccessible after a certain point; validators may follow the protocol and keep building on the longest blockchain ignoring the attack.[14] We account for this uncertainty by using f to denote the probability that an

[12]For instance, in the Bitcoin protocol $w = 100$.

[13]An attack preventing a transaction from being recorded for d blocks would require the attacker to rewrite α blocks of the blockchain, leaving out the transaction in question, and thus would require the same resources as a double spend after α blocks.

[14]One of the few documented double spending attacks in Bitcoin involved a mining pool against online gambling sites on the Bitcoin blockchain, and these attacks were accidentally discovered by a researcher long after the misappropriated funds and corresponding block rewards were spent. Details can be found at https://www.reddit.com/r/dashpay/comments/51vofr/bitcoin_casino_gets_double_spent/.

attack will be recognized by the blockchain ecosystem, and thus have consequences. For our comparisons of transaction safety, we assume the same probability of detection in both permissioned and permissionless blockchains. In reality the detection probability is likely lower for permissionless systems, for instance because the latency of the network allows for more "excuses" by the attacker. Furthermore, nodes in a permissioned system can be punished if the attack is detected at any point in the future, while in a permissionless system only detection within w blocks has consequences for the attacker. Lowering the detection probability in permissionless blockchains or lengthening the time of consequential punishment in permissioned ones, would strengthen our results in Section 6.

There is conjecture and some empirical evidence that if a successful attack is detected, the fiat price of the blockchain's cryptocurrency will decrease, which will impose a punishment on holders of this cryptocurrency, including the attackers.[15] Another possibility is that if an attack is detected, that section of the blockchain could become orphaned with the main blockchain reconfirming transactions not related to the attack. That would result in loses for the attacker only, rather than the whole ecosystem.[16]

The worst possible outcome for the attacker in our setting is that the attack is detected before the attacker's block rewards can be spent and the attacker's cryptocurrency holdings lose all their value. In that case the attacker will have incurred the cost to mine the blocks necessary for the attack, but will not be able to monetize outside the blockchain up to w block rewards.

In comparing transaction safety between permissioned and permissionless blockchains, users cannot estimate the exact cost of attack in a permissionless blockchain, as it depends on the mining power controlled by the attacker, and because of the permissionless nature of the system it is not known how much mining power any entity controls. We consider safety of transaction from the point of view of the user who neither knows the distribution of the computational power nor whether the transaction counterparty is a miner, who could mount an attack.[17] We consider the maximal cost of a successful attack that is detected, as this is

[15]See, for instance, "Bitcoin falls 11% after report suggests a critical flaw in the cryptocurrency called 'double spend' may have occurred," *Business Insider*, January 21, 2021, accessed at https://markets.businessinsider.com/currencies/news/bitcoin-price-double-spend-flaw-critical-report-suggests-2021-1-1029990921, for the response to the report of a small double spend that turned out not to be malicious as the parties involved only intended to replace the original transaction with one offering a transaction fee.

[16]This would be similar to Ethereum's attempt to revert The DAO attack.

[17]An attacker can rewrite only transactions that he himself originated. Though an attacker can prevent transactions originated by others from being recorded on the blockchain.

the worst case for the attacker, and then compare it with the cost of attack in a permissioned system. Less extreme assumptions would strengthen our results in Section 6.

In a permissionless blockchain, we call *vulnerable* any transactions of value larger than this maximal cost of attack. Such transactions are vulnerable to attack, as it is always profitable to attack them. Even if the counterparty is not the miner, the counterparty to the transaction would find it beneficial to collude with a miner or a set of miners and attack. Therefore, users should avoid *vulnerable* transactions. The threshold designating *vulnerable* transactions in a system, i.e., the maximal cost of successful attack, defines the *resiliency* of the system.

Permissionless systems do not have "safe" transactions, i.e., transactions that are smaller than the least expensive attack and therefore no miner would have an incentive to attack. That is because an attack may be successful, even with a very small fraction of computational poser, i.e., at a very low cost. It is true that the ex-ante probability of such success is low, but not zero. Another reason is that an attacker can benefit from attacking multiple transactions in the same block. In an attack, the attacker changes the whole block. Thus, replacing multiple transactions costs the same as replacing one – provided they are all transactions the attacker can control. Note that when analyzing *vulnerable* transactions, it is enough to look at a single one, as one provides sufficient incentive to attack.

The success probability of an attack that attempts to replace d blocks on the blockchain depends on d and on c/C, the computing power c controlled by the attacker (e.g., by owning corresponding hardware, or by contracting with other miners) as a fraction of the total computing power C. The higher the computational power the attacker, the more likely is the attack to succeed — and the higher its cost. The highest possible cost is when the attacker controls all of the computational power, at which point the attack replacing exactly d blocks is certain to succeed.[18] A miner controlling the full computational power would incur cost $C\mu = xR$ per period to mine — whether honestly or for an attack. In honest mining, the miner receives reward xR. If an attack is detected, however, the attacker loses wRx, i.e., the block reward for w periods.

To avoid having a payment reversed because of a successful attack, it is often advised to the recipients of high value transactions to wait for the transaction to be d blocks deep in the blockchain before irrevocably providing the value purchased to the transaction originator.

[18] In fact, already with more than 50% of mining power the attacker can be sure to succeed for sure, but the time at which the replacement blockchain will become longer than the original one will vary.

This protects the recipient because the cost of an attack increases with d as an attacker would need to replace d blocks on the blockchain. It is expected that if d is large enough so that the cost of the attack is higher than the value of double spending, then double spending can be prevented. This however, does not work if d would need to exceed w.

Given the detection probability f, the total cost of an attack requiring the replacement of d blocks in the blockchain is $fdRx$ if $d \leq w$ and $fwRx$ if $d > w$. Choosing $d > w$ does not further increase the safety of a transaction, as the attacker would be able to cash out block rewards older than w blocks.[19] Thus, transactions with value above $fRxw$ are *vulnerable*, i.e., cannot be protected from an attacker that can access the necessary computing power, leading to the following result:

Lemma 3 *In a permissionless blockchain, transactions of value to an attacker above V_{pl} are* vulnerable*, where $V_{pl} = fwRx$. That is,* resiliency *of a permissionless blockchain is V_{pl}.*

The party receiving value would typically be concerned that the party transferring the value may orchestrate or carry out an attack that reverses the transaction in question by double spending the value transferred; Lemma 3 thus implies that in permissionless blockchains users should be vary of transactions with value above $V_{pl} = fwRx$. Notice that unlike for permissioned blockchains, in the context of permissionless blockchains, the resiliency does not depend on whether attacks are contractible or not. Even if assume that attacks are not contractible, the attacker may singlehandedly control sufficient mining power. Such singlehanded attacks are never successful in permissioned blockchains (unless they are completely centralized).

Corollary 1 *The resiliency of a permissionless blockchain, V_{pl}, increases as the fiat value of the cryptocurrency x increases, and decreases as the block reward R decreases.*

This Corollary is consistent with the literature and the well known heuristics that PoW blockchains are more secure if the price of the native cryptocurrency is high, as well as the concern that the safety of PoW blockchains may decrease as the mining reward decreases.[20]

[19] Since the value of d for each transaction is decided by the transacting parties, we would not expect the recipient to choose a $d \leq w$ if it would put the transaction at risk of double spending. Note, however, that larger values of d imply a longer delay in the transaction becoming actionable, which may be costly to both parties. For instance in the case of Bitcoin, $w = 100$, which implies a delay up to 1000 minutes or more than 16 hours.

[20] See Budish (2018), Huberman, Leshno and Moallemi (2021), Pagnotta (2021), Ciaian et al (2021), Halaburda et al (forthcoming).

5 Permissionless Proof-of-Stake Blockchains

While proof-of-work is still the dominant consensus mechanism for permissionless blockchains, there has been significant recent progress on the main alternative, which is mechanisms based on proof-of-stake (PoS).[21] The largest appeal of PoS systems is that they avoid the costly externalities in PoW systems, where a substantial part of their operating cost is energy consumption that has negative environmental impacts not priced in the cost faced by the miners; for that reason PoS systems are considered to provide an environmentally friendly alternative. In PoS systems, the validators are selected to add a block, and correspondingly rewarded for this action, based on the number of coins they are staking in order to participate in the validation process.

Proof-of-stake was first implemented in the Peercoin blockchain in 2012.[22] While early versions of PoS systems were plagued by the discovery of several vulnerabilities, there was substantial work on improving the PoS consensus mechanisms, and currently we see large-scale blockchains like Solana and Ethereum 2.0 using the new generation of PoS solutions.

In the PoS setting validators (corresponding to the miners in the PoW case) put up individual stakes of s_i coins, with a resulting total number of coins staked $S = \sum_i s_i$. Not all the coins in the system are staked, for there would be no coins to make transactions. Similar to the PoW case, R denotes the block reward, i.e., the number of coins awarded to the creator of each block. We use ρ to denote the rate of return per period outside of the blockchain for an asset of similar risk profile as the coins staked by the prospective validators. In each period, the expected benefit for validator i is $\frac{s_i}{S}Rx$, the opportunity cost of i's stake is $s_i x\rho$; and thus i will stake s_i when $\frac{s_i}{S}Rx \leq s_i x\rho$. Because of free entry, there will be new staking capital entering the staking pool (i.e., S will be increasing) until $\frac{s_i}{S}Rx = s_i x\rho$. Consequently, the total cost of staking is $Sx\rho$ per period, and in equilibrium, $Sx\rho = xR$. Thus, the cost of running a permissionless PoS system is xR per block.

PoS systems require locking the staking tokens in a staking wallet in order to be considered for block creation. There is a period after being selected and creating a block before the staker can take out the staked funds; we use α to denote the number of blocks this period lasts.[23] Typically there is also a period that the tokens need to stay in the staking wallet

[21] For insightful analysis of PoS, see Saleh (forthcoming) and Rosu and Saleh (forthcoming).

[22] See Halaburda and Sarvary (2015).

[23] Ethereum 2.0 currently provides no way to withdraw the stake from the staking wallet, but that ability is expected to be implemented in future updates. There can also be a minimum capital requirement for

before they take part in block creation; we use β to denote the number of blocks this period lasts.

Cost of attack. We focus our analysis on a majority attack in a setting where staker-validators follow the longest-chain rule. As in the PoW case, we assume that x goes to 0 when the attack gets detected (with probability f), because participants lose confidence in the system. Some PoS designs, including Ethereum 2.0, "slash" (i.e., burn) the stake of a validator that is caught misbehaving. This provision guarantees that a validator considering an attack will lose its stake even in the absence of the externality that a successful attack may have on the whole system (where the price of the token goes down to 0), however it does not change the payoffs to the prospective attacker. If the attack is discovered, the attacker's payoff is 0 – just for a different reason. Consequently, in a PoS permissionless blockchain, transactions of value to an attacker above V_{pl}^{PoS} are not safe, where $V_{pl}^{PoS} = f(\alpha + \beta)Rx$.

Given the similarity between the condition $V_{pl} = fwRx$ in the case of PoW systems and $V_{pl}^{PoS} = f(\alpha + \beta)Rx$ in the case of PoS systems, in the rest of our analysis we will focus on PoW systems and the $V_{pl} = fwRx$ condition, but our results also apply to PoS subject to substituting $\alpha + \beta$ for w. Moreover, note that both in PoW and PoS the cost of running the system is Rx per block.

6 Comparing Permissioned and Permissionless Blockchains

6.1 Comparison of Transaction Safety

Two key differences between permissioned and permissionless blockchains in our model is that (a) in permissionless blockchains the participating nodes cannot be identified outside the blockchain, and (b) permissioned blockchains depend on the expectation that nodes that are discovered to compromise the integrity of the blockchain will be punished. We parametrized trust in the latter as the probability τ that such punishments will be actually administered. We can use the above to compare transaction safety in the two types of blockchains, which leads to Proposition 1.

stakes; for instance Ethereum requires a minimum stake of 32 ETH — that can be sidestepped, however, by joining a staking pool.

Proposition 1 *A permissioned blockchain has a higher level of maximum value for transaction safety than a permissionless blockchain (i.e., $V_p > V_{pl}$) if $\tau \sum_{i \in B} P_i > wRx$.*

Proof. Follows from the definitions of V_p and V_{pl}, and the fact that $V_p > V_{pl}$ when $f\tau \sum_{i \in B} P_i > fwRx$, which means that $V_p > V_{pl}$ when $\tau \sum_{i \in B} P_i > wRx$. ■

To illustrate the proposition, presently the Bitcoin protocol specifies $w = 100$ and $R = 6.25$. Let's suppose that the price of Bitcoin is $x = \$50{,}000$.[24] Then $wRx = \$31{,}250{,}000$. If the institutional framework for the permissioned blockchain is fully trusted to administer large penalties for validators that compromise its integrity, i.e., $\tau = 1$, a permissioned blockchain with large technology companies as validators will provide higher levels of transaction safety than bitcoin, as these validators will have joined reputation at stake worth significantly more than \$31.25 M, and the financial resources to credibly be assessed much higher amounts in posting a bond or paying a fine. Given the precedent of fines of hundreds of millions or billions of dollars imposed on large tech companies for data protection and privacy violations, even small values of τ will likely still result in permissioned blockchains being able to offer safety for larger transactions than permissionless ones. This is stated in Corollary 2:

Corollary 2 *For $\tau > 0$ and high enough P_i's, a permissioned blockchain is more resilient than permissionless.*

On the other hand, if trust in the institutional setting is small enough, i.e., if users expect that nodes that participate in an attack and compromise the integrity of the blockchain will be unlikely to be punished even if discovered, permissionless blockchains could offer more transaction safety. That is because if τ is small enough there may not be enough potential nodes with high enough P_i available, as shown in the example from Section 3.3. In the extreme case when there is no trust in the penalty system, or $\tau = 0$, the permissioned system provides no transaction safety at all, as for $\tau = 0$, $V_p = 0$ for any P_i's. In that case, transactions in the permissionless blockchain are safer. This is consistent with the frequent observation that permissionless blockchains do not require trust in the participants or the institutional environment, and is stated in Corollary 3:

[24] As we write this, the exchange rate for Bitcoin is about 1BTC = $\$56,900$, but that rate is very volatile.

Corollary 3 *If a permissioned blockchain operates in a trustless environment, i.e., when $\tau = 0$, then $V_p = 0$. In that case a permissionless blockchain will provide more transaction safety as long as $x > 0$, because then $V_{pl} > 0$.*

Lemma 3 has some interesting implications for permissionless blockchains. The Lemma shows that if participation in attacks is non-contractible, such attacks become impossible when more than one validator is needed for a successful attack. One of the intended benefits of permissionless blockchains was that the lack of verifiable identities would make node actions non-contractible, and specifically would prevent contracting for participation in attacks and thus increase the safety of transactions. That was particularly important for early cryptocurrencies such as Bitcoin, which had very low fiat currency values at the time of their inception and thus low values of wRx. The prevalence of smart contracts, however, and markets for computing power have enabled what seems the worse combination, where the validating nodes can contract among them based on their pseudonymous identities, and can use this contracting to orchestrate an attack, but they are not identifiable outside the blockchain, so cannot enter contracts that would punish them if they participate in attacks.

6.2 Coexistence of Permissioned and Permissionless Blockchains

In settings with heterogeneous participants, functional permissioned and permissionless blockchains can coexist. We define a blockchain, permissioned or permissionless, as functional only if it offers a strictly positive level of transaction safety, i.e., positive V.

Consider a setting with two types of users and two types of validators. Users may have high trust in the system, $\tau_H > 0$, or low trust, $\tau_L < \tau_H$. We will assume for the simplicity of exposition the extreme case where $\tau_L = 0$. The validators may have existing reputation $P_i = P > 0$ or have no reputation $P_i = 0$. We assume that there is a large number of validators available of each type. We also assume that participation in an attack is contractible.

Lemma 4 *A functional permissioned blockchain will only employ validators with $P_i > 0$.*

Proof. This result follows directly from Lemma 1. If $\tau > 0$, adding any node with $P_i = 0$ unambiguously decreases safety of the permissioned blockchain, i.e., decreases V_p. And thus, it is suboptimal to employ as validators nodes with $P_i = 0$. If $\tau = 0$, the permissioned

blockchain is not functional. ■

In our example, the permissioned blockchain will only employ as validators nodes $P_i = P$.

Lemma 5 *Permissioned blockchain with n high-reputation validators and $m(n)$ consensus rule is considered safe up to $V_{p,H} = \tau_H NP$ by high-trust users. And is considered not safe at all by low-trust users, $V_{p,L} = 0$.*

The safety of the permissionless blockchain does not depend on τ. And so, permissionless blockchain is considered safe by all users up to $V_{pl} = wRx$. When there is at least n potential nodes in the environment that $\tau_H NP > wRx$, given $\sigma(n)$, then functional permissionless and permissioned blockchains will coexist. The permissioned blockchain will be run by at least n high-reputation nodes and attract all high-trust users. The low-reputation nodes can only participate in running the permissionless blockchain, but high-reputation nodes can also engage in running it. The permissionless blockchain, however, will only attract low-trust users.

When, however, there is too few high-reputation nodes in the environment or their P is not high enough to satisfy $\tau_H NP > wRx$ for a given τ_H, then only permissionless blockchain will be functional.

6.3 Cost of Transaction Safety

We now abstract from operational costs and compare the cost of providing incentives to the validating nodes not to participate in potential attacks and thus assure the safety of blockchain transactions. Proposition 2 states that when comparing "functional" blockchains where transactions below a certain strictly positive value V_p for permissioned and V_{pl} for permissionless blockchains are safe, the incentive cost for permissioned blockchains is lower.

Proposition 2 *For any $V_p > 0$ and $V_{pl} > 0$ the cost of deterring attacks on transactions with value below V_i, $i = p, pl$, is lower for a permissioned blockchain than for a permissionless one.*

Proof. Consider first an equilibrium in a functional permissioned blockchain with $V_p > 0$. The safety of transactions with values less than V_p is guaranteed by the threat that any

validator nodes i that participate in an attack will suffer penalties P_i's if discovered, which only happens off-the-equilibrium path. Thus the cost at equilibrium of deterring attacks on transactions below V_p is zero.

Consider now an equilibrium in a functional permissionless blockchain with $V_{pl} > 0$. The safety of transactions with values less than V_{pl} is guaranteed by the threat that any validator nodes (i.e., miners) that participate in an attack will forfeit the block reward $Rx = \frac{V_{pl}}{w} > 0$ for each block of the attack if the attack is discovered. Transaction safety in permissionless blockchains thus requires the per-block cost of the corresponding proof-of-work.

Thus providing incentives to deter attacks is more costly in permissionless blockchains with consensus based on PoW than in permissioned blockchains. ■

At equilibrium, the permissioned blockchain only bears the operating cost of validating transactions, which is typically very small and we normalized as zero in our model. The cost of punishment serves as deterrent and prevents participation in attacks but it is never actually incurred at equilibrium because for $\tau > 0$ this deterrent is credible. Thus the equilibrium cost of validator incentives to keep the permissioned blockchain safe for transaction values up to $V_p = f\tau \sum_{i \in B} P_i$ is zero when $V_p > 0$. On the other hand if $\tau = 0$, then $V_p = 0$ and the permissioned blockchain cannot be functional.

A permissionless blockchain is safe up to V_{pl} in equilibrium, but the cost of deterring attacks is born up front as a reward of Rx per block. Thus the equilibrium cost of validator incentives to keep the permissionless blockchain safe for transaction values up to $V_{pl} = f\,w\,R\,x > 0$, is Rx per block.

7 Conclusions and Future Research

In permissionless blockchains validators enjoy anonymity and free entry and exit; thus they must choose to follow the consensus mechanism based on the incentives provided within the blockchain by its protocol. In permissioned blockchains, validators can be induced to follow the protocol based on mechanisms external to the blockchain, such as legal contracts, monetary penalties and reputation. Permissionless blockchains thus rely on their protocol to ensure compliance of the validators, and their participants "trust the code." Permissioned blockchains can ensure validator compliance based on the enforcement mechanisms provided by the institutional setting, and their participants "trust the institutions."

Both the popular and academic blockchain literature agree that permissioned blockchains offer lower cost of operation compared to permissionless blockchains because they don't have to depend on a costly consensus mechanism such as Proof of Work, and higher operational efficiency as they can optimize the validator network for the desired latency and transaction capacity, both of which are challenges in a permissionless blockchain that does not control the number and identity of its validators. It is often argued that these operational benefits come at the cost of transaction safety as transactions in permissioned blockchains cannot achieve the safety offered by permissionless blockchains, especially in settings without trust.

We showed that under certain conditions this tradeoff does not need to be there at all. Unless facing in a completely trustless environment, where permissioned blockchains cannot operate as they depend on at least some minimal level of trust to identify and discipline their participants, well-designed permissioned blockchains can be both more efficient to operate and offer higher transaction safety than permissionless blockchains.

We showed that the specific parameters of the blockchain setting determine whether permissioned or permissionless blockchains will perform better in terms of providing transaction safety. Known identities in permissioned blockchains allow contracting among validators both for attacks and for penalties if the integrity of the blockchain is compromised. Permissionless blockchains dispense with identities that can be established outside the blockchain, and this prevents contracting to not participate in attacks and follow the protocol rules. On the other hand smart contracts and efficient markets for computing power that have developed to facilitate efficient mining, can also be used to acquire computing power or coordinate attacks. This makes high value transactions on permissionless blockchains vulnerable to attack. Which type of blockchain performs best in providing a desired level of transaction safety depends on the parameters of the setting, specifically the degree of trust in the institutional setting of the blockchain, the liability limits and reputation of the validators in the permissioned blockchain, and the level and maturation of block rewards in the permissionless blockchain.

A key parameter in the case of permissioned blockchains is the trust in the system, specifically the contractual promises and reputation of the validators. Permissioned blockchains cannot function in the absence of such trust, in which case permissionless blockchains are the only option. Famously this was the reason for the inception of Bitcoin. In the presence of even modest trust, however, permissioned blockchains can provide higher transaction safety than permissionless by recruiting validators with high enough reputation and financial credi-

bility. Most business relationships involve at least a modest level of trust in the institutional setting and the involved counterparties, and in these cases permissioned blockchains can outperform permissionless ones in transaction safety. This is consistent with the observation that permissioned blockchains dominate business applications of blockchain technology.

Employing validators with high credibility, i.e., validators that would suffer substantial monetary and/or reputational losses if they were to deviate from the consensus protocol, increases transaction safety in permissioned blockchains, and thus both blockchain operators and users would prefer validators with deep pockets and established reputations. Technology companies with strong reputations from other markets thus have an advantage compared to less established entrants in the operation of permissioned blockchains. For instance even though the Hyperledger blockchain platform is open source, IBM Blockchain enjoys the advantage of both its own reputation and its access to a set of reputable validators; it has thus dominated the operation of Hyperledger-based blockchain platforms in corporate settings where transactions are mission-critical and require high degrees of safety.

Our results in this paper suggest several questions for future research. An important question to address is how to design a permissioned blockchain focusing on its governance mechanism in order to maximize its benefits in terms of efficiency, high performance and transaction safety. For instance, should Walmart or IBM be running the FoodTrust blockchain, and who would be the ideal validators? The first step is already indicated in our exploration—the validators in the permissioned blockchain should face large reputational risk or have "deep pockets" so that they can credibly be faced with large fines, loss of reputation and possibly offer restitution if their actions compromise the integrity of the blockchain.

References

Yackolley Amoussou-Guenou, Bruno Biais, Maria Potop-Butucaru, Sara Tucci-Piergiovanni (2019), Rationals vs Byzantines in Consensus-based Blockchains, February 2019, available on arXiv

Raphael Auer, Cyril Monnet and Hyun Shin (2021), Permissioned Distributed Ledgers and the Governance of Money, February 2021, available on ssrn

Babich, V., and G. Hilary. 2020. OM Forum—Distributed Ledgers and Operations: What Operations Management Researchers Should Know About Blockchain Technology. Manufacturing & Service Operations Management 22:223–240.

Soumya Basu, David Easley, Maureen O'Hara, Emin Gun Sirer (2019), Towards a Functional Fee Market for Cryptocurrencies, accessed via arXiv

Biais, B., C. Bisiere, M. Bouvard, and C. Casamatta. 2019. The Blockchain Folk Theorem. Review of Financial Studies 32:1662–1715.

Blaettchen, P., A. Jagmohan, K. Ratakonda, and M. Franceschini. 2020. A information network for food safety: A simulation approach. Working Paper

Eric Budish (2018), The Economic Limits of Bitcoin and the Blockchain, NBER working paper

Cao, S., L. W. Cong, and B. Yang. 2018. Auditing and Blockchains: Pricing, Misstatements, and Regulation. Working Paper .

Catalini, C., and J. Gans. 2019. Initial Coin Offerings and the Value of Crypto Tokens. NBER Working Paper .

Jonathan Chiu and Thorsten Keoppl (2017), The Economics of Cryptocurrencies? Bitcoin and beyond, available on ssrn

Chod, J., and E. Lyandres. 2018. A Theory of ICOs: Diversification, Agency, and Information Asymmetry. Working Paper.

Chod, J., N. Trichakis, G. Tsoukalas, H. Aspegren, and M. Weber. 2019. Blockchain and The Value of Operational Transparency for Supply Chain Finance. Management Science

Ciaian, Pavel, d'Artis Kancs and Miroslava Rajcaniova (2021), Interdependencies between Mining Costs, Mining Rewards and Blockchain Security, available on arXiv

Cong, L. W., and Z. He. 2019. Blockchain Disruption and Smart Contracts. Review of Financial Studies 32:1754–1797.

Cong, L. W., Y. Li, and N. Wang. 2019. Tokenomics: Dynamic Adoption and Valuation. Working Paper.

Cui, Yao, Hu, Ming and Liu, Jingchen (2019), Values of Traceability in Supply Chains. Available on SSRN

Cui, Yao, Gaur, Vishal and Liu, Jingchen (2020), Blockchain Collaboration with Competing Firms in a Shared Supply Chain: Benefits and Challenges. Available on SSRN

Lin William Cong, Zhiguo He, Jiasun Li (2019), Decentralized Mining in Centralized Pools, available on ssrn

David Easley, Maureen O'Hara, Soumya Basu (2019), From mining to markets: The evolution of bitcoin transaction fees, Journal of Financial Economics,Volume 134, Issue 1,October 2019, Pages 91-109

Zahra Ebrahimi, Bryan Routledge, Ariel Zetlin-Jones (2019), Getting blockchain incentives right, CMU working paper

Joshua Gans and Neil Gandal (2019), More (or Less) Economic Limits of the Blockchain, December 2019, available on ssrn

Hanna Halaburda, Guillaume Haeringer, Joshua Gans and Neil Gandal (forthcoming), The Microeconomics of Cryptocurrencies, Journal of Economic Literature, forthcoming

Hanna Halaburda and Miklos Sarvary (2015), Beyond Bitcoin: The Economics of Digital Currencies, Palgrave MacMillan

Franz Hinzen, Kose John and Fahad Saleh (2019), Bitcoin's Fatal Flaw: The Limited Adoption Problem, NYU working paper

Gur Huberman, Jacob Leshno and Ciamac Moallemi (2021), Monopoly without a monopolist: An economic analysis of the Bitcoin payment system, available on ssrn

Garud Iyengar, Fahad Saleh, Jay Sethuraman and Wenjun Wang, Economics of blockchain adoption. Working Paper August 13, 2020

Mariana Khapko and Marius Zoican (2020), How Fast Should Trades Settle?, Management Science, Vol. 66, No. 10, October 2020, pp. 4359-4919

Alfred Lehar and Christine Parlour (2020), Miner Collusion and the Bitcoin Protocol, available on ssrn

Jacob Leshno and Philipp Strack (2020), Bitcoin: An Axiomatic Approach and an Impossibility Theorem, American Economic Review: Insights, Vol. 2, No.3, September 2020, pp. 269-86

Nikhil Malik, Manmohan Aseri, Param Vir Singh and Kannan Srinivasan (2021), Why Bitcoin will Fail to Scale? Economics of Collusion on Bitcoin, accessed via ssrn

Katya Malinova and Andreas Park (2017), Market Design with Blockchain Technology, accessed via ssrn

S. Narang, M. Byali, P. Dayama, V. Pandit and Y. Narahari (2019), Design of Trusted B2B Market Platforms using Permissioned Blockchains and Game Theory, 2019 IEEE International Conference on Blockchain and Cryptocurrency (ICBC), Seoul, South Korea.

Emiliano Pagnotta (2021), Decentralizing Money: Bitcoin Prices and Blockchain Security, Review of Financial Studies

Julien Prat and Benjamin Walter (2019), An equilibrium model of the market for Bitcoin mining, accessed via ssrn

Fahad Saleh (forthcoming), Blockchain without waste: Proof-of-Stake, Review of Financial Studies, forthcoming

Ioanid Rosu and Fahad Saleh (forthcoming), Evolution of Shares in a Proof-of-Stake Cryptocurrency, Management Science, forthcoming