

Towards Risk Scoring of Bitcoin Transactions

Malte Möser, Rainer Böhme, and Dominic Breuker

Department of Information Systems, University of Münster, Germany

Abstract. If Bitcoin becomes the prevalent payment system on the Internet, crime fighters will join forces with regulators and enforce blacklisting of transaction prefixes at the parties who offer real products and services in exchange for bitcoin. Blacklisted bitcoins will be hard to spend and therefore less liquid and less valuable. This requires every recipient of Bitcoin payments not only to check all incoming transactions for possible blacklistings, but also to assess the risk of a transaction being blacklisted in the future. We elaborate this scenario, specify a risk model, devise a prediction approach using public knowledge, and present preliminary results using data from selected known thefts. We discuss the implications on markets where bitcoins are traded and critically revisit Bitcoin’s ability to serve as a unit of account.

1 Introduction

Whenever a merchant receives a 100-dollar note, she is well advised to carefully check whether it is authentic. Bitcoin is a decentralized cryptographic currency with the ambition to take over the role of dollar notes at least in the domain of online transactions. As bitcoins¹ are mere references in a public ledger, the Bitcoin equivalent to checking the authenticity of conventional banknotes would be to rule out inconsistencies in the global system state which could nullify an incoming payment. The recommended and well-known defense against the so-called double spending risk is patience. The merchant has to wait until a transaction is sealed deep enough in the block chain to make revisions extremely costly, and hence unlikely [17]. But sooner or later, patience will not be enough.

The popularity of Bitcoin among criminals [13], allegedly for its anonymity and loose to absent regulation, has called for new approaches to fighting financial crime committed in or settled through Bitcoin. A promising strategy is to blacklist transaction prefixes to invalidate assets originating from criminal proceeds [25]. This strategy is effective and practical because the blacklists can be enforced at the services accepting bitcoins. Those are not decentralized and therefore cannot evade law enforcement in their jurisdiction of residence; and, by extension of mutual legal assistance, the set of internationally recognized provisions for the fight against financial crime. In fact, the ability to enforce such a blacklisting policy thwarts the very idea of a decentralized currency by projecting power of

¹ Convention: We capitalize Bitcoin when referring to the name of the system and use lower case for the monetary unit (like dollar, euro). BTC is shorthand for the unit.

the legal system into Bitcoin. This is why blacklisting practices are controversial among Bitcoin enthusiasts [6]. We leave this philosophical debate aside and concentrate on the effect of blacklisting policies on transactions in general. In practice, blacklisting is reality in Bitcoin [10] and new ventures seek to offer whitelisting services with similar effect [16].

This paper contemplates a future of Bitcoin where blacklisting of known bad transaction prefixes is common practice and the resulting blacklists are observed by all relevant parties where bitcoins can be spent. As a result, end users receiving payments in Bitcoin must screen incoming transactions as well. We can safely assume that suitable services and APIs will be offered by third parties.

However, even when payments appear benign, recipients can never be certain if a prefix of their incoming transactions will be blacklisted *in the future*. They have to accept a risk of invalidation while holding bitcoin. This specific risk is probably small compared to all other risks involved with Bitcoin for the time being, but the proportions may change as the currency gains popularity. Unlike other risks, this risk is idiosyncratic for the transaction history of the specific incoming transaction. For example, a transaction that forwards freshly mined bitcoins (so-called coinbase transactions) has less likely been involved in a crime than a transaction consisting of bitcoins that have changed ownership more often. This gives rise to the idea of *predicting the risk of blacklisting* to value incoming transactions and manage the spending risk.

This paper sets out to specify a risk model and outline a prediction approach using public knowledge from the Bitcoin block chain. We also present preliminary results for selected known thefts; although the low number of events and heterogeneity of data prevent us from actually calibrating and running the model. As an equally important contribution, we discuss the implications on the future of Bitcoin. The paper is organized as follows. Section 2 recalls essential features of the Bitcoin system and ecosystem with special emphasis on risks in general. Section 3 develops a model for the specific risk of transaction blacklisting. Section 4 presents our empirical findings, and Section 5 discusses implications. The paper concludes with an outlook on future work (Section 6).

2 Background

2.1 Bitcoin and the Real World

To reason about Bitcoin and its relation with the real world, it is useful to introduce some terminology. Our conceptual model in Figure 1 distinguishes the core Bitcoin system from a surrounding ecosystem. The core system consists of a protocol, implicitly specified by the reference implementation of the client software, and data representing the global consensus system state. This state is stored in the public block chain and continuously being updated by all clients participating in the Bitcoin peer-to-peer network. The core system is decentralized and designed with the aim to withhold control by central entities.

The Bitcoin *ecosystem* is the set of market operators leveraging the Bitcoin system. It includes Bitcoin-specific financial *intermediaries*, such as exchanges,

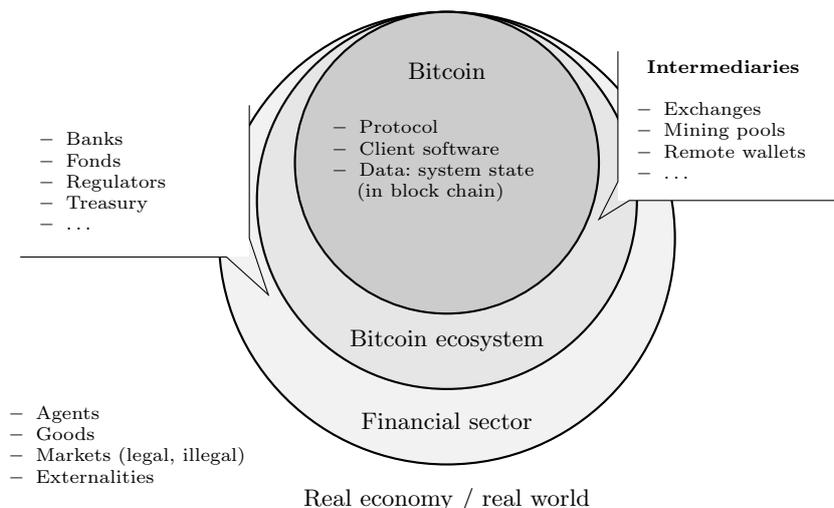


Fig. 1. Bitcoin's relation to the real world

mining pools, remote wallets, or transaction anonymizers. Some intermediaries are necessary to make Bitcoin usable as a global Internet currency, but unlike the core system, Bitcoin intermediaries are *not* decentralized. To avoid single points of failure and to discipline the intermediaries, competition between intermediaries offering substitutable services is desired and required.

The outer layers in Figure 1 reflect the conventional separation of the financial from the real sector. As some Bitcoin intermediaries, notably exchanges, interface with conventional financial intermediaries, notably payment systems, we can depict the financial sector as another layer shielding Bitcoin from the real world. The intersection of all layers at the top of the figure symbolizes the possibility to skip layers. For example, the externality of cycles burned to reach consensus via proof-of-work materializes in energy consumption and heat production in the real world without necessarily involving the layers in between [3].

2.2 Implications for Crime Fighters

The description in this section draws on our prior work [25]. Recall that markets in the real economy include legal and illegal activities alike. As criminals use the financial sector (and Bitcoin) to commit crimes and to launder criminal revenues, law enforcement has to take precautions. Fighting crime in the conventional financial sector lasts on the Know-Your-Customer (KYC) principle. The principle mandates financial intermediaries to verify the identity of clients before doing business with them. KYC was tightened in the US Patriot Act in order to strengthen efforts of anti-money laundering (AML) and combating the financing of terrorism (CFT). Many jurisdiction followed this US initiative. However, KYC

is only one cornerstone. It must be complemented by risk assessment, monitoring, reporting and enforcement. Once identities are established via KYC, they become the identifiers enabling all downstream activities. Standard procedures include suspicious activity reports filed with financial intelligence units (FIUs), or automatic cross-checks against blacklists maintained by financial crime fighters, such as the US Office of Foreign Assets Control. In simple terms, fighting financial crime in conventional payment systems relies on known identities and does not require a full picture of all transactions.

Bitcoin, by contrast, is designed with pseudonymous identities. Account numbers are public keys of a digital signature system. Account ownership is established by knowing the corresponding private key. Everyone with a computer can create valid key pairs from large random numbers and thus open one or many Bitcoin accounts. Although the relation between Bitcoin accounts and civil identities of their owners is a priori unknown, Bitcoin transactions are not anonymous. A simple abstraction for Bitcoin is to think of it as a *public* distributed ledger that records all transactions between valid Bitcoin accounts. In short, fighting financial crime in Bitcoin means dealing with imperfect knowledge of identities, but may exploit perfect knowledge of all transactions.

2.3 Risks of Holding Bitcoin

Individuals or organizations holding bitcoins are faced with several types of risks, some of which can be managed by taking appropriate precautions. Most prominently, there is exchange rate risk. Compared to ordinary currencies, Bitcoin is still very volatile. Within just four weeks in fall 2013, Bitcoin soared as the exchange rate increased from 200 USD to 800 USD, i. e., by 400 %. On the contrary, when the Bitcoin exchange Mt. Gox was hacked in June 2011, the perpetrators caused the exchange rate to drop from 17,50 USD to just a single cent [18]. (But the exchange rate recovered minutes after the event.) The uncertainty of the exchange rate is one of many factors that might impede businesses accepting bitcoin. A mitigation strategy is to regularly convert bitcoins into local currency and keep only a small transaction budget at risk. Commercial payment providers, for instance BitPay [1], offer services to automate this process.

Closely related to exchange risk is the risk of a systemic Bitcoin failure. This means that a catastrophic event dries out the market and lets the exchange rate plummet close to zero. One reason could be a major government intervention. Although no government can stop Bitcoin from existing, a coordinated action of large countries can nevertheless force the currency into the underground. While some jurisdictions appear to tolerate Bitcoin, others, such as Thailand [20], are more reserved. Moreover, any glitch in the implementation of the Bitcoin protocol could easily cause a failure, too. Namecoin, a special-purpose Bitcoin derivate, was affected by such a failure recently. The attempt to recover it was still ongoing at the time of writing [7]. Obviously, this kind of risk can only be managed by not holding more bitcoins than one can afford to lose.

Whenever users deal with intermediaries, they are exposed to counterparty risk. There are many reported cases where Bitcoin intermediaries closed their business

with their clients' deposits as loot. Whether the root causes were fraudulent motives or plain bankruptcy is of secondary interest in the absence of effective means of fund recovery. Moore and Christin [24] have empirically analyzed factors behind exchange failures and calibrated a prediction model for such events.

The risks described above affect all users (or all users within a large group) alike, but there are also risks idiosyncratic to users. First, users face the risk of making mistakes when sending transactions. As Bitcoin transactions are irreversible, typos in the transaction amount require the recipient's active collaboration to undo that error. Fool-proof client implementations are necessary to mitigate the risk of making mistakes.

Careless users may lose the private keys, which are required to spend their bitcoins, e. g., due to a failure of the storage medium of their wallet. Nobody really knows to which extent users have suffered losses so far. Ron and Shamir identify large amounts of dormant coins, i. e., bitcoins which have not been used for a long time, in their transaction graph analysis and conjecture that these might be lost coins [28]. Regular backups of private keys reduce this risk.

Users may not only lose private keys by improvidence, but may also become victims of theft. Many Bitcoin users do not keep their private keys in their own domain of trust. Instead, they entrust online service providers with managing their wallets. Such providers are hacked quite regularly, which usually means their customers lose everything. Recently, the wallet provider "inputs.io" has been compromised and the bitcoin equivalent of 1.2 million USD has been stolen [21]. Replacing online services by personal devices is not necessarily a solution. For instance, wallets managed with Android devices have been found vulnerable to a weakness of Android's random number generator [2]. Hence, paying close attention to security is critical to mitigate the risk of theft.

Another risk that received considerable attention is double-spending [3, 17, for example]. Bitcoin's nature of a decentralized peer-to-peer system relying on proof-of-work to maintain the integrity of the global state puts individual clients at the risk of believing in a transaction that will be invalidated in the future. The specific risk of double-spending declines exponentially with the number of blocks after the inclusion of the transaction [26]. Hence, while double-spending occurs regularly [8], some patience when accepting Bitcoin payments is enough to avoid falling for it.

Similar to double-spending, blacklisting is another risk of receiving apparently valid bitcoins at one point in time, which become invalid at another. Although not extensively used these days, if blacklisting becomes common practice, it is in the users' best interest to account for the risk of blacklisting whenever accepting a Bitcoin payment. What is special about this risk is that whether bitcoins are blacklisted or not depends on their transaction history, i. e., on whether those transactions preceding the current one were involved in a crime. This calls for risk scoring based on the public information contained in the block chain. Although blacklisting has been a topic in the Bitcoin community for some time, we are not aware of any attempts to set up such a scoring model.

The collection of risks provided in this section is by no means exhaustive. More subtle risks exist as well, such as losing financial privacy if the association between a Bitcoin address and its owner becomes public. This paper focuses on the blacklisting risk specifically.

3 A Risk Arrival Model for Blacklisting Events

3.1 Blacklisting Policies: Poison and Haircut

To tackle the quantification of the risk of transaction blacklisting, it is important to specify what the consequences of blacklisting can be. Transactions are blacklisted with a certain probability if they are involved in a crime. Typical Bitcoin crimes include theft from popular online wallet providers or illegally earned proceeds from blackmailing, e. g., with ransomware such as CryptoLocker [15]. The goal of blacklists is to render the criminals' bitcoins useless, thereby lowering the incentives for this criminal activity. To achieve this end, governments could mandate all legitimate businesses not to accept transactions directly associated with blacklisted transactions.

There are several problems with this approach. First and foremost, criminals can create as many identities as they want [14]. Hence, they can send their dirty bitcoins through several fake addresses. They could repeat this procedure until it appears to a ingenuous observer that there is no connection to the criminal source. To avoid this, blacklisting has to propagate through the transaction graph to punish anyone, both fake identities of criminals and ordinary users, for accepting blacklisted bitcoins. Honest users can avoid undue punishment by obeying the blacklist preemptively.

Unfortunately, there will be a certain timespan between the point in time at which an illegal transaction takes place and the point in time at which it is added to the blacklist. Thus, honest users may accept a dirty bitcoin despite their best efforts to comply with the blacklist. These users, not knowing that they have accepted a dirty bitcoin, might combine three small amounts of bitcoins A1, A2 and A3 to create a large transaction B. With the propagation mechanism in place, B would also be affected by blacklisting if only one of its input transactions A1, A2 or A3 is dirty.

Consequently, it is important to specify how exactly B would be affected. Two basic blacklisting policies are conceivable. In the first, which we call "poison", B would be invalidated just as any other blacklisted coin. The poison policy implies that every transaction is invalidated that has at least one dirty predecessor, no matter how many generations above. Note that the propagation works on the level of transactions (not addresses) and requires the recipient to act. This prevents that saboteurs can destroy other people's bitcoin wealth by routing a blacklisted transaction to their publicly known address.

The second, less drastic policy is one we call "haircut". Instead of invalidating a transaction entirely, it is devalued proportionally to the amount of blacklisted bitcoins in its inputs, again applied recursively. In the example above, if the

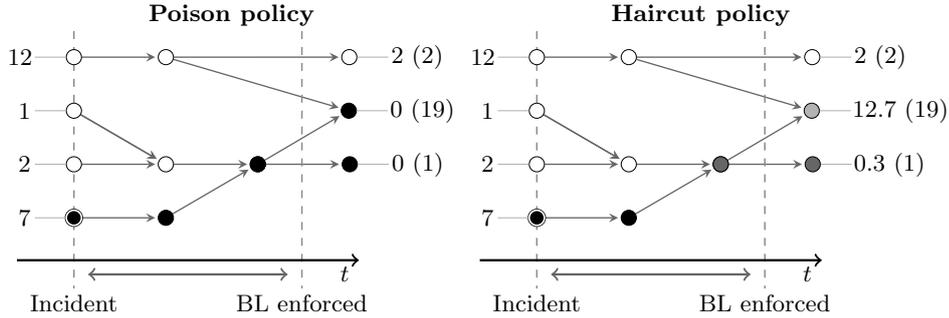


Fig. 2. Timing model of transaction blacklisting (BL) with different policies

three transactions A1, A2 and A3 were all worth one bitcoin and one of them was blacklisted, transaction B would be treated being worth 2 BTC (although nominally, in the block chain, it would be worth 3 BTC). It is easy to see how this policy propagates through the transaction graph.

Figure 2 shows an advanced transaction graph example of how the two basic policies affect transaction values. Nodes represent transactions, whereas arrows represent the flow of value between them (i. e., an output of a former transaction is used as an input in the successive transaction). The color of a node represents the state of blacklisting, where white represents clean coins and black blacklisted coins. In the poison scenario, an initial theft of 7 BTC leads to a total loss of 20 BTC, as blacklisted coins were combined with clean coins and thereby change their state. In the case of the haircut policy, different colors of grey illustrate the amount of a transactions devaluation. As the stolen 7 BTC are repeatedly combined with clean coins, the share of blacklisted value decreases (and the color gets brighter). In contrast to the poison scenario, the total amount of blacklisted value stays the same.

Of course, more advanced policies that look deeper into transactions are conceivable. Imagine a “FIFO” policy, where the order of the inputs determines which outputs are affected by blacklisting. If a transaction has two inputs worth 1 BTC and only the second input originates from a blacklisted transaction, the first output(s) will be considered “clean” until they add up to a value of 1 BTC and the remaining outputs will be partially or fully invalidated. Note that as a result, such granular blacklisting policies may make the internal structure of transactions subject for negotiations between sending and receiving parties.

3.2 Risk Arrival and Impact Analysis

As discussed, blacklisting exposes all users to the risk of accepting a bitcoin which is blacklisted in the future. Rational users desire to keep this risk small. In particular, at some point in time, a user will be presented with a transaction created by another user and he will have to decide whether or not to accept it; or,

more precisely, how much the bitcoins being transferred are worth to him. The user reduces the nominal value, i. e., the BTC amount specified in the transaction, by an appropriate risk premium reflecting the risk of blacklisting.

This risk is idiosyncratic for the transaction as it depends on the history of all inputs. For simplicity, we consider only one input noting that the generalization to multiple inputs is straightforward if the blacklisting policy is known. A rational user will analyze the history using suitable predictors. By observing previous thefts and the traces criminals leave behind in the block chain, he could identify characteristic patterns and search for these patterns in the given history. This information could be used to estimate the expected loss associated with accepting the transaction at face value.

For both policies, poison and haircut, users must first estimate the probability that a transaction will be affected by blacklisting as a function of time. Second, users must estimate how long they keep (parts of) this transaction in their own accounts. This is non-trivial as users may prioritize spending of risky coins or spend bitcoins faster in general, which has monetary implications not further detailed here. Third, they have to estimate their loss in case of blacklisting. This depends on the policy.

For the poison policy, the third part is straightforward. As all affected bitcoins are void, users multiply the blacklisting probability with the transaction value. Things are more complicated with a haircut policy. But the haircut policy also has advantages. Imagine the poison policy was in place and a user has accepted a transaction. When he combines this transaction with others, e. g., to consolidate his funds in cold storage, he also puts his other funds at risk. In such a policy, combining transactions would increase risk, effectively lowering the value. As a result, users would avoid doing so if not absolutely necessary. With the haircut policy, by contrast, the total value at risk does not change when transactions are combined.

Transaction histories grow substantially over time. Apart from technical challenges with analyzing such histories (e. g., finding efficient algorithms and data structures), this also causes hard-to-calculate risk. There will always be a very small but positive probability of very old incidents becoming generally known. Especially with a poison policy, this could affect very large numbers of transactions potentially causing systemic instability of the currency. Hence, some form of statutory period after which no blacklisting is done seems reasonable.

4 Prediction Approach

4.1 Model and Data

The heart of any risk scoring model is the set of predictors that separates the dangerous coins from the harmless. To test such predictors, data from real thefts is needed in order to distinguish normal from criminal behavior. While law enforcement agencies might have a comprehensive list of incidents at their hands, the only source of public information we are aware of is a list of major thefts

and losses that happened from 2011 until today (November 2013). The list is maintained by users of the “Bitcoin Forum” [4]. There are a number of aspects limiting the applicability of this dataset:

- The list contains only large thefts, between 922 and 263,024 BTC.
- Only few thefts include a list of the relevant transactions.
- It is difficult to determine when a theft was “officially” announced, yet the exact point in time is needed to determine when official blacklisting could have taken place.

Out of nine thefts that contain a list of relevant transaction, we can only assign a concrete blacklisting timestamp to six incidents. We used the time of first announcement in the Bitcoin forum as the assumed time of blacklisting. Furthermore, only three of these thefts show block chain activity between the theft and the blacklisting timestamp, rendering the other three useless for our purposes (cf. Table 1). Nonetheless, we show exemplary results from our attempts to evaluate our predictors using the information from these three incidents.

Table 1. List of known incidents with transaction data

Incident	Severity	Reporting Time	Coins used
Allinvain Theft	25000 BTC	2011-06-13 20:47	yes
Linode Hack	46653 BTC	2012-03-01 21:43	yes
Betcoin Theft	3172 BTC	2012-04-13 12:19	no
May 2012 Bitcoinica Hack	18548 BTC	2012-05-11 13:16	no
Bitfloor Theft	24086 BTC	2012-09-04 17:08	yes
Cdecker Theft	9222 BTC	2012-09-28 08:10	no
Mass MyBitcoin Theft	4019 BTC	unclear	–
2012 Trojan	3500 BTC	unavailable	–
Bitcoin Syndicate Theft	1852 BTC	unavailable	–

4.2 Candidate Predictors

We briefly discuss a few potential predictors for our risk model. First of all, predictors need to be powerful and efficiently measurable. As they are based on public data, criminals could try to outsmart them, hence predictors should ideally be hard to manipulate. Our list includes public information that can be gathered from the block chain only. It is conceivable to include other public information, such as references to Bitcoin addresses on the web. If the risk prediction is offered as a service, it is also possible to include private information, for instance collected by exchanges and other intermediaries; and enriched with behavioral profiles acquired from social media or search engines. It is unclear if such information is of any help to detect indications of criminal activity. Very speculatively, a potentially powerful predictor using semi-public information would be whether

transaction prefixes have been seen at exit nodes of the Tor network. If this information helps, one could even pay the operators of exit nodes for sharing such knowledge. But such ideas are clearly out of scope of our initial prediction attempts.

In the following we describe the predictors considered in this study.

Transaction Value Ordinary users can be expected to transfer ordinary amounts of bitcoins between addresses for ordinary purposes such as changing bitcoins for local currency at a Bitcoin exchange. The number of bitcoins being transferred will be rather small. Any activity involving very large sums of Bitcoins would be suspicious, e.g., because a criminal hacked an exchange and exfiltrated all bitcoins at once. Hence, a straightforward indicator is to observe transaction volumes and to pay close attention to outliers from an expected distribution as they may stem from a major theft.

Obfuscation Patterns If a thief is clever enough to steal a large amount of bitcoins, he is probably also clever enough to avoid leaving obvious traces of his crime behind. Instead of simply combining his entire haul to a huge lump sum to spend on a Ferrari, he may also use carefully designed obfuscation patterns, possibly involving numerous fake identities, to stop law enforcement from tracing him. Assuming that those taking on such efforts are more likely involved in criminal activities than others, risk scoring models could search for typical obfuscation strategies as indicators of blacklisting risk. Examples of such patterns have been observed in the Bitcoin transaction graph before, e.g., in the context of analyzing Bitcoin mixing services, which make extensive use of peeling chains [25]. Other studies also identified various characteristic patterns used by larger organizations in the Bitcoin ecosystem. Most importantly, these patterns involve aggregations, foldings (i. e., combining blacklisted transactions with clear transactions), splits and also peeling chains [22]. Hence, peeling chains are not necessarily associated with Bitcoin mixing but may still increase the risk.

Frequency of Usage A thief trying to obfuscate the origin of his bitcoins may construct a complex web of transactions between multiple fake addresses to make others believe these Bitcoins were involved in ordinary business. What a thief might also want is to launder stolen coins as soon as possible. As he knows his bitcoins were involved in a crime, he is fully aware of the blacklisting risk and holding them for a long time increases his risk of loss. Hence, the thief must construct his fake transactions fast, resulting in a short time span between transactions involving these coins. This motivates measuring transaction frequency to use it as an indicator of risk. Whenever frequency increases above average values, there are reasons for suspicion.

Change Addresses and Multi-Input Transactions The predictors presented so far make use of characteristics of the transaction graph, but do not

take the addresses associated with the transactions into account. Heuristics such as the detection of change addresses [22] or the combination of addresses from multi-input transactions [27], which are likely to belong to the same user, can detect connections between apparently unrelated transactions. This allows to possibly link public keys of transaction prefixes to other blacklisted transactions. Furthermore, it can provide evidence whether obfuscation patterns or a high frequency stem from a natural origin or are constructed by a single entity.

Coinbase Transactions Coinbase transactions are a special type of transaction of which there is one per block in the block chain. These transactions do not have inputs, i. e., they create new bitcoins without using up the bitcoins from other transactions. At the moment, their value is equal to 25 bitcoins plus the sum of all transaction fees associated with the other transactions in the block. The primary purpose of coinbase transactions is to provide an incentive for users to participate in the creation of proof of work, which is necessary to ensure manipulation resistance of the global state. Also, they solve the problem of bootstrapping the network [9]. Because coinbase transaction have no history attached to them, their risk of being involved in a theft is obviously very low – except for the case in which a thief is able to control a miner or mining pool directly and thus the coinbase transaction itself would be blacklisted. If the thief colludes with a miner to launder his coins through large transaction fees, one could apply the blacklisting policy also to transaction fees – where the haircut policy seems to be the more appropriate policy as it leaves the fixed reward of the coinbase transaction unaffected. To measure the portion of “clean” coins, we compare the value of coinbase transactions $C_i(t)$ to the value of all $V_i(t)$ transaction i steps away from the transaction of interest t , and then sum up the individual values for each i in a degressive fashion. This reduces the influence of transactions further away. One can limit i to $i = 10$, as larger values for i will influence the score only marginally:

$$X(t) = \sum_{i=1}^{10} \frac{1}{2^{(i-1)}} \cdot \frac{C_i(t)}{V_i(t)}.$$

Larger values of $X(t)$ imply a larger amount of value stemming from coinbase transactions and thus reduce the risk of blacklisting.

4.3 Preliminary Results

Although not adequate to calibrate our model, we calculate indicator scores for all three thefts and compare the results to a control group, where we randomly choose a transaction for each affected transaction. As n (the number of affected transactions) is low for the Allinvain and Betfloor theft, we increase n in the control group to 112. Transactions of the control groups are drawn from the same block of the blacklisted transaction to ensure comparability. (Finding the right sampling and bootstrapping approach for this purpose is deferred to

Table 2. Predictors for thefts and control groups

Predictor	Allinvain		Bitfloor		Linode	
	Incident	Control	Incident	Control	Incident	Control
n	7	112	8	112	82	112
Average	2.7743	1.2664	2.8249	0.3669	1.2715	0.8287
Median	3.2355	1.3573	2.8489	0.4848	1.1491	1.0553
SD	1.1687	1.1495	0.8206	1.0718	1.3929	1.0087
Variance	1.5936	1.3332	0.7695	1.1592	1.9641	1.0265
Peeling Chains	0.2857	0.6639	0.6250	0.3571	0.3415	0.3482
Duration (h)	611	218	68	25	28	93
Coinbase Score	0.1027	0.0575	0.0056	0.0103	0.0181	0.0948

specialized future work.) If there are not enough transactions available, we draw from nearby blocks. The first four predictors analyze transaction values. Due to the small number of transactions, very large values may bias the results. We therefore log-transform all transaction values. The other predictors are calibrated as follows. For peeling chains, we look if a transaction comes from a peeling chain with a minimum length of 6 transactions. To measure the portion of coinbase transactions, we use the formula stated above. However, we reduce the depth i from 10 to 8 in order to increase performance.

The results (cf. Table 2) reflect our previous observation that these three thefts constitute a very special case of theft: they affect long-term user or intermediaries of the Bitcoin system, who have a large amount of bitcoins available, in parts presumably from own mining activity. The average and median transaction size in the control group is always smaller than in the incident set. There is not only heterogeneity between the particular incident and control groups, but also between the members of both incident and control groups. As a result, without better and more general data, we are currently not in a position to derive an adequate set of predictors for a risk scoring model with acceptable predictive power.

5 Market Implications

In this section we discuss how markets are affected by a regime that strictly enforces blacklists of transaction prefixes. We assume that some (imperfect) prediction models for spending risk are available, either based on public information or using proprietary information against a small fee.

5.1 1 BTC \neq 1 BTC

Bitcoins are not alike. Each transaction is a descendant of a unique transaction history, which is readily available in the public block chain. Therefore, markets participants can, in principle, scrutinize the history and become selective in which

transactions they accept; or, with more granularity, how much they value it. The fact that most participants do not differentiate for the time being is hard to justify with economic rationality. A necessary consequence of differentiation is that market prices reflect the information encoded in the transaction history. Dealing with bitcoins of two kinds (e. g., black and white, under the poison policy) may be manageable, essentially at the cost of lower liquidity in both market segments.² Pricing *every history individually* poses new challenges to the design of market mechanisms, for example at exchanges; but it also affects every small merchant who accepts bitcoins in exchange for goods or services.

As price differences reflect spending risk, we may follow the model of credit markets and introduce intermediaries who publish commonly accepted risk ratings of all unspent transactions. This comes with two issues. First, transaction rating agencies are a new kind of intermediaries that need to be paid. Second, the fact that people must rely on them conflicts with the idea of decentralization. If we try to decentralize ratings to resolve these issues (for example by trivial replication), the rating model must be confined to public information, or use non-trivial homomorphic encryption. Both options raise new questions. For example, can we generate meaningful risk ratings based on public information with public algorithms and still remain game-proof? Game-proofness is an important property that discourages whitewashing of transaction histories. If one could bounce transactions between own accounts to predictably increase their value, everybody would do it, resulting in a choked up network and block chain.

The credit rating analogy has limits. Here is one important difference: conventional credit ratings (allegedly) have an information advantage. They aggregate *private* information that is not readily available to all market participants. In the absence of better information, market participants rely on ratings as common proxy. If market participants had access to the disaggregated information—machine-readable and at no cost, like for Bitcoin transactions—some would prefer to aggregate the information using customized models. And it would be naive to assume that buyers and sellers agree on the same model, let alone on its parameters. The multitude of private valuation functions calls into question the conventional order book approach followed by popular Bitcoin exchanges. Instead, we need new efficient mechanisms that reveal and match market participants' private valuations of all transactions on a marketplace.

One thing that is conceivable under the haircut policy is risk pooling, using insurance markets and CoinJoin [5] as models: risk-averse bitcoin owners can reduce the variance of their spending risk by forming large transactions with many others, thereby distributing the impact of bad transactions equally. Of course, such schemes require coordination effort and they are vulnerable to adverse selection as parties who know that they possess dubious transactions have higher incentives to participate in the pool. This is a known issue of insurance markets.

² Intentionally “colored” coins have been proposed to deal with virtual goods of different value using Bitcoin as an infrastructure [29].

5.2 YouMoney or Bloodcoin ?

Taking the uniqueness and identifiability of Bitcoin transactions beyond the question of pricing offers interesting new insights. Precious metals or official fiat currencies are designed as homogeneous goods. This ensures fungibility: quantities are exchangeable and divisible, a precondition to fulfill the monetary function as *unit of account*. Bitcoin transactions, by contrast, are heterogeneous goods, differentiated on a quality dimension. The valuation of this quality is subject to individual preferences. This threatens the function as unit of account, as detailed above in Section 5.1.

On the upside, however, recipients of payments could apply ethical standards on what money they accept. This is best comparable to the Kimberly Process Certification Scheme [30], a set of international resolutions established with the aim to suspend the trade of blood diamonds, which are mined in war zones and sold to finance arms. One might wonder how well certain industries fared if the money they accept could be traced back with the ease offered by Bitcoin. Although market participants still use Bitcoin like a fungible good, being more selective in what one accepts stands to reason. At least those who have a brand or reputation to lose have every reason to be afraid of negative publicity linking their profits to violence, with evidence publicly accessible in the block chain.

More generally, the public transaction history turns Bitcoin into *personalized money*. Because the origin of a transaction matters for spending and reputation, accepting bitcoins from one party implies a trust relationship from the payee to the payer. This trust must be established and signaled outside the Bitcoin protocol, for example by linking Bitcoin identifiers to reputation or social networking systems. Like the rating analogy above, this is another interesting feature where Bitcoin markets resemble credit markets more than foreign exchange markets: Bitcoin recipients take over the sender's spending risk in a similar vein as creditors bear their debtors' credit risk. In rough circumstances, a recipient of Bitcoin transactions may even ask for a security to cover potential losses. So at second sight, Bitcoin is not so dissimilar to systems of decentralized debt obligations, like iOwe [19].

This collection of unfinished thoughts indicates that understanding the full implications of perfectly and publicly traceable payments remains a major task for interdisciplinary research.

5.3 Alternatives

What can we do about it? In a decentralized system with a public global state, the only way to make cryptographic tokens homogeneous, and the virtual good they encode fungible, is to make transaction histories untraceable. Zerocoin [23] promises this option. Other solutions involve hierarchical structures of fast and anonymous cryptographic cash [12, 11] issued by competitive intermediaries and backed with a slower and less anonymous cryptographic reserve currency, such as Bitcoin. However, all these options make it harder to fight crime by following the money. Ideally, we would like to see compromises like systems offering practically

untraceable transactions for small amounts; and the computational effort needed to trace entities decreases gently as the amounts involved grow. This turns the access to transaction histories from a global binary property to a variable transaction cost, again with implications for all market participants. At the time of writing, we do not know a technical solution for this set of requirements. We conjecture that it will be hard to realize when identities are cheap.

The economics of adoption are crucial in this context [9]. All conceivable alternatives require coordinated effort to switch from Bitcoin to the new regime. This revolution is hard to achieve against vested interests. Blacklists, by contrast, are evolutionary. They can emerge without changing the core system and thus are hardly avoidable if Bitcoin remains relevant.

6 Future Work

Directions for future work include observing blacklisting practices in the Bitcoin ecosystem, collecting data for more incidents, and finding better tailored predictors to estimate and eventually validate a risk scoring model.

For simplicity, we have assumed that one blacklist is obeyed globally. In practice, synchronizing this blacklist is another problem that is prone to race conditions. Moreover, as blacklists are enforced by national law enforcement agencies, it is likely that the world will see at least one blacklist per country. So the spending risk also depends on the spending conventions. Despite people say the Internet has no borders, the would-be Internet currency becomes as messy as international trade.

Note that the authors do not approve or disapprove blacklisting of Bitcoin transactions. Our mission is to reason about the consequences of foreseeable developments. A then relevant topic completely out of the scope of this work is to explore how the governance of blacklists can be put under decentralized control.

References

1. BitPay. <https://bitpay.com/>.
2. Android Security Vulnerability. <http://bitcoin.org/en/alert/2013-08-11-android>, 2013.
3. Jörg Becker, Dominic Breuker, Tobias Heide, Justus Holler, Hans Peter Rauer, and Rainer Böhme. Can We Afford Integrity by Proof-of-Work? Scenarios Inspired by the Bitcoin Currency. In Rainer Böhme, editor, *The Economics of Information Security and Privacy*, pages 135–156, Berlin Heidelberg, 2013. Springer.
4. Bitcoin Forum. List of Major Bitcoin Heists, Thefts, Hacks, Scams, and Losses. <https://bitcointalk.org/index.php?topic=83794>.
5. Bitcoin Forum. CoinJoin: Bitcoin privacy for the real world. <https://bitcointalk.org/index.php?topic=279249.0>, 2013.
6. Bitcoin Forum. Mike Hearn, Foundation’s Law & Policy Chair, is pushing blacklists right now. <https://bitcointalk.org/index.php?topic=333824.0>, 2013.
7. Bitcoin Forum. Namecoin was Stillborn, I Had to Switch Off Life-Support. <https://bitcointalk.org/index.php?topic=310954>, 2013.

8. Blockchain.info. Double Spends. <https://blockchain.info/de/double-spends>.
9. Rainer Böhme. Internet Protocol Adoption: Learning from Bitcoin. In *IAB Workshop on Internet Technology Adoption and Transition (ITAT)*, 2013.
10. Vitalik Buterin. Mt.Gox: What the largest exchange is doing about the Linode theft and the implications. <http://bitcoinmagazine.com/mtgox-the-bitcoin-police-what-the-largest-exchange-is-doing-about-the-linode-theft-and-the-implications/>, 2012.
11. D. Chaum, A. Fiat, and M. Naor. Untraceable electronic cash. In Shafi Goldwasser, editor, *Advances in Cryptology — CRYPTO' 88*, volume 403 of *Lecture Notes in Computer Science*, pages 319–327, New York, 1990. Springer.
12. David Chaum. Security without identification: Transaction systems to make big brother obsolete. *Communications of the ACM*, 28(70), 1985.
13. Nicolas Christin. Traveling the Silk Road: A Measurement Analysis of a Large Anonymous Online Marketplace. In *Proceedings of the 22nd International World Wide Web Conference*, pages 213–224, Rio de Janeiro, 2013.
14. John R. Douceur. The Sybil Attack. In Peter Druschel, Frans Kaashoek, and Antony Rowstron, editors, *Peer-to-peer Systems*, volume 2429 of *Lecture Notes in Computer Science*, pages 251–260. Springer, Berlin Heidelberg, 2002.
15. Dan Goodin. You're infected—if you want to see your data again, pay us USD 300 in Bitcoins. <http://arstechnica.com/security/2013/10/youre-infected-if-you-want-to-see-your-data-again-pay-us-300-in-bitcoins/>, 2013.
16. Kashmir Hill. Sanitizing Bitcoin: This Company Wants To Track 'Clean' Bitcoin Accounts. <http://www.forbes.com/sites/kashmirhill/2013/11/13/sanitizing-bitcoin-coin-validation/>, 2013.
17. Ghassan O Karame, Elli Androulaki, and Srdjan Capkun. Two Bitcoins at the Price of One? Double-Spending Attacks on Fast Payments in Bitcoin. In *Proceedings of the ACM Conference on Computer and Communications Security (CCS)*, 2012.
18. Mark Karpeles. Clarification of Mt. Gox Compromised Accounts and Major Bitcoin Sell-Off. https://www.mtgox.com/press_release_20110630.html, 2011.
19. Dave Levin, Aaron Schulman, Katrina LaCurts, Neil Spring, and Bobby Bhattacharjee. Making Currency Inexpensive with iOwe. In *Proceedings of the Workshop on the Economics of Networks, Systems, and Computation (NetEcon)*, San Jose, 2011.
20. Andrew Saks McLeod. Thailand Bans The Bitcoin! National Foreign Exchange Department Rules Bitcoin Illegal, Trading Suspended. <http://forexmagnates.com/bitcoin-binned-thailands-foreign-exchange-department-rules-bitcoin-illegal-trading-suspended/>, 2013.
21. Robert McMillan. USD 1.2M Hack Shows Why You Should Never Store Bitcoins on the Internet. <http://www.wired.com/wiredenterprise/2013/11/inputs/>, 2013.
22. Sarah Meiklejohn, Marjoir Pomarole, Grant Jordan, Kirill Levchenko, Damon McCoy, Geoffrey M. Voelker, and Stefan Savage. A fistful of bitcoins: Characterizing payments among men with no names. In *Proceedings of the ACM Internet Measurement Conference (IMC)*, pages 127–140, New York, 2013. ACM.
23. Ian Miers, Christina Garman, Matthew Green, and Aviel D Rubin. Zerocoin: Anonymous distributed e-cash from bitcoin. In *IEEE Symposium on Security and Privacy*, pages 397–411, San Francisco, 2013. IEEE.
24. Tyler Moore and Nicolas Christin. Beware the Middleman: Empirical Analysis of Bitcoin-Exchange Risk. *Financial Cryptography and Data Security*, (June 2011), 2013.
25. Malte Möser, Rainer Böhme, and Dominic Breuker. An Inquiry into Money Laundering Tools in the Bitcoin Ecosystem. *Proceedings of the APWG E-Crime Researchers Summit*, 2013.

26. Satoshi Nakamoto. Bitcoin: A Peer-to-Peer Electronic Cash System, 2008.
27. Fergal Reid and Martin Harrigan. An Analysis of Anonymity in the Bitcoin System. In Y. Altshuler, Y. Elovici, A.B. Cremers, N. Aharony, and A. Pentland, editors, *Security and Privacy in Social Networks*, pages 197–223. Springer, New York, 2013.
28. Dorit Ron and Adi Shamir. Quantitative Analysis of the Full Bitcoin Transaction Graph. In Ahmad-Reza Sadeghi, editor, *Financial Cryptography and Data Security*, volume 7859 of *Lecture Notes in Computer Science*, pages 6–24. Springer, Berlin Heidelberg, 2013.
29. Meni Rosenfeld. Overview of colored coins. <http://bitcoil.co.il/BitcoinX.pdf>, December 2012.
30. Wikipedia. Kimberley process certification scheme. http://en.wikipedia.org/wiki/Kimberley_Process_Certification_Scheme, 2013.