

Join Me on a Market for Anonymity

Malte Möser¹ and Rainer Böhme²

¹ Department of Information Systems, University of Münster, Germany

² Department of Computer Science, University of Innsbruck, Austria

Abstract. We present the first measurement study of JoinMarket, a growing marketplace for more anonymous transfers in the Bitcoin ecosystem. Our study reveals that this market is funded with multiple thousand bitcoins and generated a turnover of almost 8 million USD over the course of eight months. Assessing the resilience of the market against a well-funded attacker, we discover that in a typical scenario, a selective attack with 90% success rate requires an investment of 32,000 USD (which is recoverable after the attack). We formulate stylized economic models of supply and demand to explain the existence of this novel market for anonymity and underpin some theoretical arguments with empirical data.

1 Introduction

Anonymity and economics are an odd couple. Most microeconomic models assume agents without name, and fail to predict outcomes if agents become identifiable [3, 39]. Anonymity can be defined as the state of being “not identifiable within a set of subjects, the anonymity set” [33]. With this definition, a simple measure of the quality of anonymity is the size of this set, as the probability of successful identification by random guessing is inversely proportional to the set size.

If there exist situations where the state of being anonymous improves an agent’s wealth, one would expect a market for anonymity to develop. Yet anonymity is an unconventional economic good. To produce it, other agents must behave in an indistinguishable way to an observer of the agent who seeks anonymity. This is nicely summarized in the expression “anonymity loves company” [17]. In the language of economics, the production of anonymity generates positive externalities because all agents who contribute to the supply of anonymity also receive the good in demand. Production and consumption are hard to tell apart. So, what should the market price for anonymity be?

While the economics of privacy [34] and personal data [22] have been studied for long, and empirical research has estimated price information (see [23] and [1] for reviews), remarkably little is known about the price of anonymity. Acquisti, Dingledine and Syverson [2] study the incentives to participate in anonymous communication systems based on mix networks, such as Tor. Their analysis is comprehensive, includes attacker behavior and adoption dynamics, but remains theoretical. Spiekermann [38] interprets survey data collected by self-selection among the early-adopters of an academic anonymous communication system (JAP). Köpsell [21] uses technical measurements in an experimental setup of

the same system to approximate the value of anonymity by observing users' aggregate willingness to trade performance for anonymity. (Time-tradeoffs were used subsequently to quantify the value of security, e. g., [19]). All these works contribute interesting observations, but barely scratch the surface of the puzzles associated with markets for anonymity.

In this paper, we leverage the cryptographic currency Bitcoin and its ecosystem as a “social science laboratory” [10]. We present a longitudinal measurement study of a market designed to match supply and demand of anonymous value transfers. In principle, Bitcoin transactions can be traced and histories inspected for known identifiers that allow informed parties to associate the initially pseudonymous account numbers with real-world identities. JoinMarket, our object of study, offers a clandestine marketplace to arrange a special kind of transaction that mixes transfers of many different parties, thereby exponentiating the complexity of deanonymization attempts. All parties involved in such a transaction roughly form an anonymity set as used in the above definition of anonymity.

Our approach draws on a combination of methods to answer a number of research questions. We collected price information from the public order book of this market between June 2015 and January 2016. Moreover, we obtain volume information by matching changes in the order book with likely trades in Bitcoin's public ledger. To validate our method with ground truth, we participated on the supply side of the market at selected points in time using a minimal invasive trading strategy. This combination of methods allows us to quantitatively describe the market development over time. Our second contribution is on the economics of security. We study the theoretical possibility of an attacker participating in the market and estimate the cost of deanonymization over time. This cost is expressed in terms of capital employment and as a function of the targeted probability of success. We note that adversaries may even profit from launching attacks, generated from fees paid for the (then empty) promise of better anonymity. To better understand this and other anomalies of markets for anonymity, we formulate stylized economic models of supply and demand. The first model uses time preferences and the second model uses qualitative differentiations to explain the existence and price formation on this market. Where possible, we underpin the underlying hypotheses with data from JoinMarket.

The remainder of this paper is organized as follows. We offer the necessary technical background about Bitcoin, CoinJoin transactions, and JoinMarket in Section 2. Our measurement approach, descriptive results, and the hypothetical attack scenario are presented in Section 3. Section 4 discusses economic models devised to explain the observed anomalies. The paper closes with a brief discussion (Section 5) and concludes in Section 6.

2 Background

2.1 Bitcoin in a Nutshell

Bitcoin is the first, and to this date most popular, instance of a decentralized cryptographic currency [32]. A key characteristic of this first generation of

cryptographic currencies is their public ledger, replicated on every “full” node of a peer-to-peer network [10, 11]. This ledger, called blockchain, contains the records of all transactions that have ever taken place in the system. Since only transactions reassign ownership of bitcoins and each transaction references all relevant previous transactions, it is possible to validate the state of the public ledger by following the references backwards. A probabilistic consensus protocol resolves conflicting updates at the end of the ledger. Its design, incentive mechanisms, and security properties are vital for the system but irrelevant for this paper.

The Bitcoin protocol assigns value, denominated in bitcoins (BTC), to addresses. Bitcoin addresses serve like account numbers. They are derived from the public keys of an asymmetric encryption system. Ownership of accounts is controlled by the knowledge of the corresponding private keys. As anyone can generate fresh key pairs, Bitcoin users enjoy a degree of pseudonymity. However, with all addresses and the associated transactions stored in the public blockchain, an observer can identify relations between them. If this knowledge is enriched with auxiliary information on the real-world identities behind addresses, it becomes possible to deanonymize selected users [27, 35, 36].

Transactions in Bitcoin specify the origin and the destination of the value transferred. They include a list of inputs, which are references to existing funds, and a list of outputs that specify its new owners. Whenever a user transfers bitcoins, she has to spend the full value of the input and therefore returns any surplus as *change* to an address under her control. A common deanonymization technique is to cluster addresses that are associated with inputs combined in one transaction, as this behavior suggests common knowledge of all of the associated private keys [35, 36].

CoinJoin, a special convention for transactions, intentionally breaks this heuristic [25]. Multiple senders and recipients of funds combine their payments in a single joint transaction (cf. Figure 1). This is possible and secure because valid Bitcoin transactions require individual signatures for each public key associated with the funds used. If CoinJoin was default in Bitcoin, it would increase users’ privacy by rendering the aforementioned multi-input heuristic more difficult to apply. A limitation remains: individual values may still leak sufficient information to derive matching subsets [5, 26].

2.2 JoinMarket

A major barrier towards the adoption of CoinJoin is to match users who are interested in creating a transaction. This opens up an opportunity for intermediaries. JoinMarket [20] is a platform for Bitcoin users wishing to participate in CoinJoin transactions. It has been operational since May 2015 [7]. In contrast to previous approaches [9, 16], JoinMarket does not aim at matching different users who all want to create a transaction at the same time. Instead, it divides users in two groups: (market) *makers* and *takers* of CoinJoin offers. Makers offer their bitcoins for use in takers’ CoinJoin transactions. The advantage of this approach is that users do not have to wait for partners when creating CoinJoin transactions. Instead they can choose from a list of offers by the market makers. To incentivize

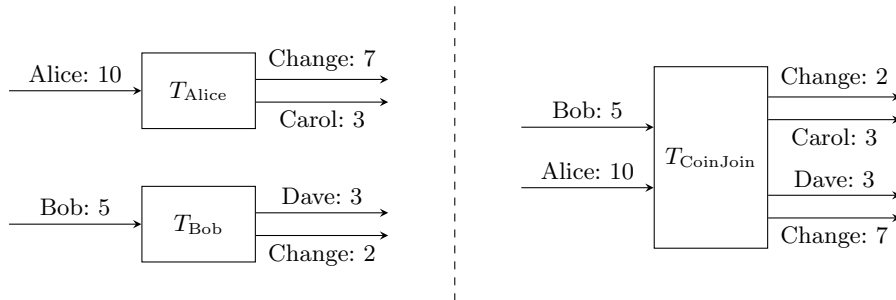


Fig. 1: Two or more individual payments (left) can be combined in a single CoinJoin transaction (right). The spending amounts need not be identical, but in JoinMarket they are.

participation, takers pay makers a small compensation (further referred to as *maker fee* to be distinguished from the general *miner fee* in Bitcoin).

The technical backend of JoinMarket is rather simple. It does not use mixing protocols such as Xim [8], CoinShuffle [37], CoinParty [43], or Mixcoin [12] that provide decentralization, Sybil-resistance, or at least warranties. In the current implementation, takers and makers communicate via a centralized Internet Relay Chat (IRC) channel. Whenever a maker joins the channel she announces her offers to the channel. JoinMarket has no central server that stores a list of available offers; the offer announcements over time form a public order book that every member can maintain locally³. Whenever a maker announces or updates her offers, or leaves the channel, the local database is updated. A maker wishing to participate in CoinJoins must stay connected to the IRC server, typically by running an IRC bot that automatically announces offers and updates.

Each offer in the order book is uniquely identified by the combination of a username and an offer identifier (*uid, oid*). Every offer record further contains a maker fee f of type $t = \{\text{rel, abs}\}$, a minimum amount v_{\min} , a maximum amount v_{\max} , and a contribution to the miner fee c that has to be paid in order for the transaction to be included in the blockchain. Relative fees (with $t = \text{rel}$) are specified in percent of the value a taker intends to send, not the sum of her input values. The miner contribution is intended to reimburse the taker for the increase of the miner fee due to a larger transaction size. It gives more flexibility in specifying the total fees for the taker.

JoinMarket does not automatically match orders. A taker interested in creating a CoinJoin transaction will join the IRC channel, request the current list of offers and receive them from each individual maker in a private message. She will then choose a set of offers with one of the following methods:

1. Random draw from a distribution that weighs offers based on their fee;
2. Deterministically choose offers with lowest fees; or
3. Manual selection.

³ Some websites allow to inspect the order book, e.g., <https://joinmarket.me/ob/>.

The first option is preferable over minimizing the fees because it avoids the risk that a single cheap offer dominates the market. After requesting a set of inputs and destination addresses from each participant, the taker constructs the CoinJoin transaction. Finally, she will publish the transaction on the Bitcoin network.

Two additional details are necessary to understand our analysis of JoinMarket. First, JoinMarket transactions are special cases of CoinJoin transactions with characteristics that allow us to identify them in the transaction graph generated from the public ledger. If n participants construct a transaction, it will have n outputs with the exact same value (we call this the *spend*) and usually also the same number of “change” addresses. Takers can also choose to sweep their wallets and send all of their funds to an address. In this case there will only be $n - 1$ change outputs. There must also be at least n inputs⁴, each associated with a different Bitcoin address.

Second, JoinMarket makers use a deterministic wallet as specified in BIP 32 [42]. Different wallet chains separate so-called mixing *depths*. Each spend is sent to an address belonging to the next depth, while the change stays in the current depth. This prevents the reuse of a spend/change address pair as common inputs in a future JoinMarket transaction. In the early days of JoinMarket, each maker would simply pick the wallet chain with the highest amount of bitcoins available. Nowadays many makers publish offers for each of the chains, distinguished by the *oid*. This raises interesting questions related to the pricing of individual offers. Makers offering large amounts, for example, could demand higher fees as long as there is sufficient demand. Depending on the distribution of the spending values, different mixing depths may also be priced differently in order to merge funds in such a way that they fit the distribution. These questions require a richer analytical model and answering them is beyond the scope of this paper.

3 Measurements

We now present what is to the best of our knowledge the first measurement study of JoinMarket.

3.1 Data Collection and Preprocessing

We used JoinMarket’s built-in order book watcher to monitor the available offers and stored a snapshot about every five minutes since the beginning of June 2015 until the end of January 2016. This gives us 288 order book snapshots per day, from which we extract all individual offers (6.16 million entries in total). Due to IRC disconnects and crashes of the order book watcher, we miss data for a few timestamps. In total, our dataset covers 98.06 % of the whole timeframe.

One issue when analyzing the order book is that it is impossible to verify whether stated offers are indeed genuine – makers could easily overstate the

⁴ Specifically, there must be n input subsets with a value greater or equal to the spend.

amount of bitcoins they offer, or offer low fees and then fail to deliver. In principle, makers could even serve the market without ever stating offers on the public channel (i. e. only make private statements to takers requesting the order book). Still, we assume that the majority of offers is genuine and visible in the public order book. For data cleaning, we decided to remove offers with the *uid* “fakeorder”, which claimed to offer up to 2.1 million bitcoins (i. e. 14 % of all bitcoins in circulation). We also removed offers with a maximum amount of zero. To accommodate the risk of outliers due to short-lived exaggerated offers in the order book, we calculate the median over a time interval of one day whenever we report aggregated values.

Besides taking snapshots of the order book, we also ran our own maker bot for multiple weeks. Running a maker bot allowed us to analyze the characteristics of real CoinJoin transactions without significantly influencing what we aim to measure. Of course, we cannot rule out the possibility that the transactions we attribute to normal users are the result of other researchers’ participation. In total, we participated in 503 CoinJoin transactions with spends below 0.5 BTC. This number was largely endogenous due to the behavior of other market participants. From the point of view of research ethics, we would have liked to keep this number lower. However, the market was bumpy at times and manual intervention would have compromised the reliability of the ground truth data. Note that participation as a maker (as opposed to being a taker) does not generate volume. If at all, it marginally increases the anonymity offered to market participants.

We then used one of the transactions we participated in as a starting point and traversed the transaction graph for other JoinMarket transactions between block heights 358,000 (end of May 2015) and 396,048 (end of January 2016). Our criteria for identifying JoinMarket transactions were n spending outputs (i. e. outputs with the same value) and n or $n - 1$ change outputs as well as at least n inputs with $n \geq 2$. We excluded some obvious false positives, such as transactions where all inputs belong to the same address or transactions where the largest input is necessary to create multiple spending outputs. With this technique, we identified 8,648 potential JoinMarket transactions, which correctly include all 503 ground truth transactions.

Whenever we report USD values, we use an exchange rate of 400 USD per BTC, roughly the average exchange rate in early 2016, according to [14].

3.2 Market Overview

Figure 2 plots descriptive statistics of the JoinMarket order book over time. We report both the total amount of bitcoins available as well as the number of offers and makers. The thin lines connect daily medians. Thick lines are fitted smoothing splines with ten degrees of freedom. The total number of bitcoins available rose from initially a few hundred in mid 2015 to more than 2,000 in November, dropping to values between 1,000 and 1,500 in January 2016. Overall there has been a steady increase in the number of offers, while the actual number of makers stays relatively stable. Our interpretation is that makers adopt more advanced bots which offer bitcoins at different mixing depths.

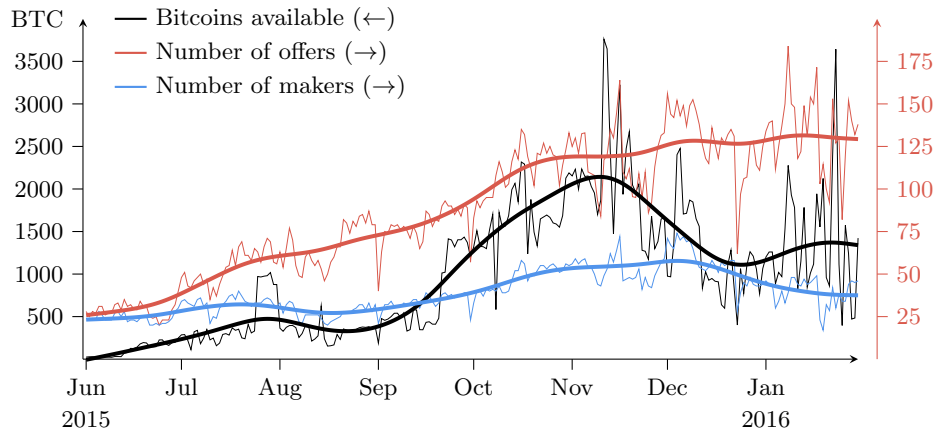


Fig. 2: Key indicators of the JoinMarket order book

While a four-digit figure of bitcoins available already suggests a large market (e. g., 2,000 BTC correspond to roughly 800,000 USD), it is more instructive to look at the maker fees at which those bitcoins are available. Figure 3 breaks down the available bitcoins by maker fees in percent of the transaction amount. Absolute fees are converted to relative terms using the maximum available amount. Observe that the majority of bitcoins is available for a maker fee below 0.01%. In comparison, centralized services offering coin anonymization often charge fees between 1–3% [30].

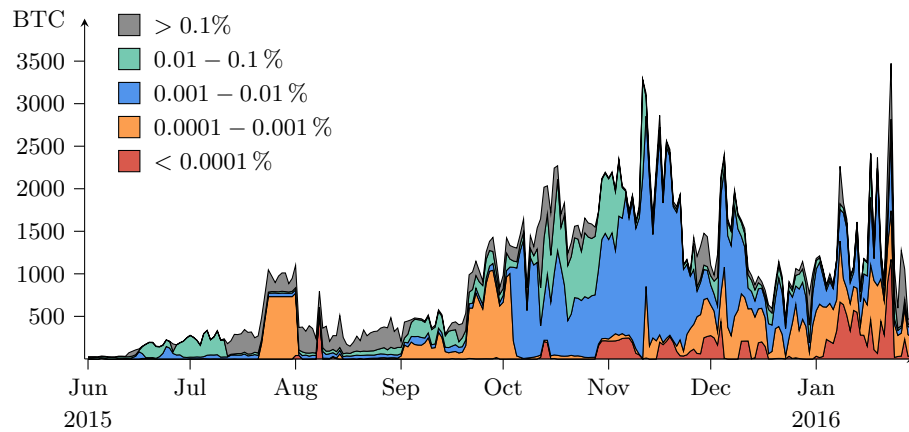


Fig. 3: Amount of bitcoins available on JoinMarket at a certain maker fee

3.3 Transaction Volume

It is not directly observable from the order book how many of the offers have been accepted. We present two estimates for the total number of JoinMarket transactions.

Our first approach is to identify transactions in the blockchain that are directly related with our ground truth JoinMarket transactions. This strategy is valid because the outputs of other makers are usually reused as inputs in other JoinMarket transactions. By traversing the transaction graph in both directions, following the inputs and outputs of each transaction and identifying potential JoinMarket transactions, we obtain a set of 8,648 JoinMarket transactions. The total spending value of these transactions amounts to 19,802 BTC, which corresponds to almost 8 million USD. Since we cannot rule out that the Bitcoin transaction graph contains other subgraphs separate from our own transactions in the relevant timeframe, this figure may serve as a lower bound.

Our second strategy is based on changes in the order book. Whenever a maker took part of a JoinMarket transaction, she updates her public offer as the bitcoins available in her wallet chains have changed. We count 39,290 changes of offers in the order book between consecutive timestamps. We aggregate this data by the timestamp because most CoinJoin transactions involve multiple makers. This introduces a small probability of error if we aggregate offers belonging to different concurrent CoinJoin transactions. Given the transaction volume measured in the first approach, we are confident that the frequency of our timestamps is high enough to keep this error negligible. This approach yields an estimation of 11,727 JoinMarket transactions.

Combining these two values, we get an rough estimate of 10,000 JoinMarket transactions over the course of eight months, i. e. about 41 transactions per day on average. Figure 4 compares the estimated daily transaction volume for both methods over time. We see that usage peaks between October and December at around 150 transactions per day. The substantial co-movement between both graphs suggests that our estimation heuristics are reliable. We conjecture that the slightly larger estimates from the order book updates are due to manual interventions of makers and delayed updates of offers, which make them appear as separate transactions.

3.4 “Pay Your Attacker”

A major limitation inherent to JoinMarket is that an attacker can conduct a Sybil attack [18] in order to deanonymize takers. In such an attack, the attacker would impersonate a large number of makers in order to be the sole other participant in a CoinJoin transaction. As she knows which of the inputs and outputs belong to herself, she can attribute the remaining inputs and outputs to the taker. This renders the transaction as traceable as without CoinJoin. She could then sell this information to Bitcoin intelligence firms, for instance.

JoinMarket does not actively prevent Sybil attacks but relies on the market mechanism to make such an attack costly. Whenever a taker creates a CoinJoin

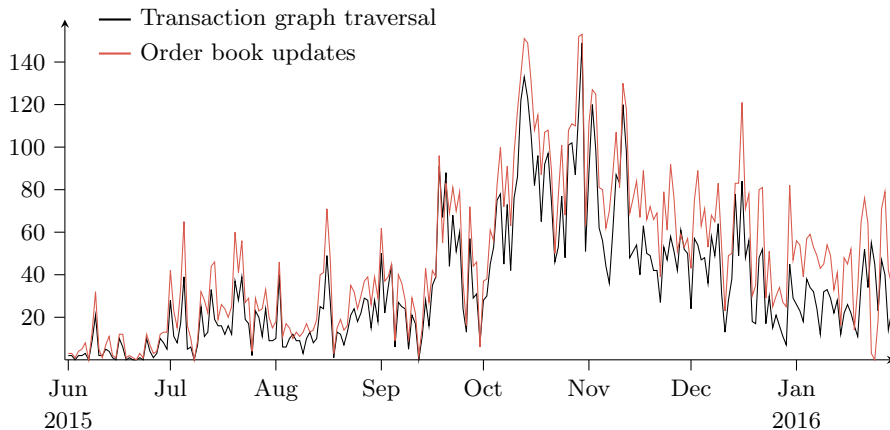


Fig. 4: Estimated daily JoinMarket transaction volume based on transaction graph traversal and changes in the order book

transaction, she can choose one of the three offer selection methods mentioned in Section 2.2. The default option, random draw with weighted probability function, takes multiple rounds depending on the number of makers chosen. First, for each maker the cheapest matching offer is selected such that only one offer per maker can possibly be chosen. This prevents makers with multiple offers from gaining advantage in the selection process. Next, the offers are sorted by their total fee (i. e. adjusted maker fee minus the miner fee contribution) and each offer receives a probability based on an exponential function parametrized to take into account the number of participants and the distribution of fees. After selecting an offer based on these probabilities, the process is repeated, now excluding all offers of the previously chosen maker(s). This bounds the success rate of a Sybil attacker by the number of coins at her disposal in order to outnumber all other offers for each potential spending amount.

We can estimate how many of the offers an attacker would have needed to control in order to be the sole participant of a user’s CoinJoin. To this end, we first explore typical behavior by tabulating the choice variables from the JoinMarket transactions identified in the blockchain:

- the number of makers takers choose to join with and
- the distribution of spending values.

Figure 5 shows that most takers join with two makers, likely because this is the default. It is somewhat alarming to see that most users choose a low number of participants as this makes it easier to deanonymize the subsets [26]. Informed by this empirical distribution, we decided to calculate scenarios with 2, 3 and 5 makers. The spending values of the CoinJoin transactions follow a log-normal distribution (see Figure 8 in the appendix), from which we choose scenarios at the 25%, 50% and 75% quantiles, which correspond to 0.035 BTC (14 USD), 0.254 BTC (102 USD), and 1.4 BTC (560 USD).

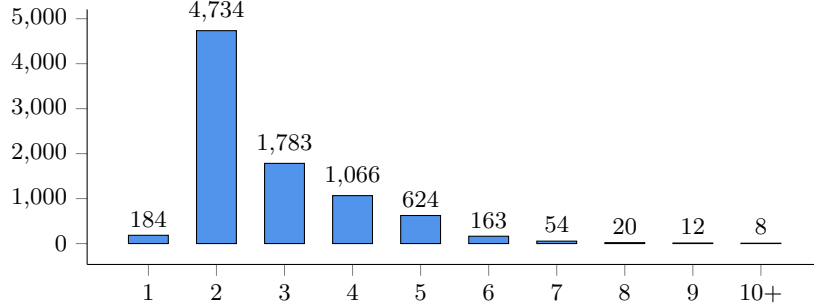


Fig. 5: Number of makers chosen by the takers (default = 2)

We compute probabilities for all combinations of spending values and numbers of makers. To speed up the calculation, we only use a subset of our data, i. e. one snapshot per hour. Then we select the set of cheapest offers that, in total and across multiple rounds, yield a success rate for an imaginary attacker of 75 % and 90 %, respectively.

Table 1 reports the number offers that are needed in order to reach the specified success probability in a given scenario defined by the taker’s choice of spending value and number of participating makers m . For example, to be the sole other participant in a transaction that seeks three makers at a spending value of 1.4 BTC, the attacker would need to control (on average) the eleven cheapest offers for a success rate of 75 %, and thirteen for a success rate of 90 %. In general, we see that the number of offers needed grows with m . There is almost no difference between attacking transfers of 0.035 BTC and 0.254 BTC, but the number of offers drops for the scenario with 1.4 BTC.

Next, we also extract the cumulative value of these offers, which gives us an indicator for the amount of bitcoins an attacker would need (cf. Table 2). For example, an attacker who wants to be the sole other participant in a CoinJoin transaction with $m = 3$, a value of 0.254 BTC and $p \geq 0.9$ would require about 80

Table 1: Number of offers needed to be the sole other participant in a CoinJoin transaction with probability p , depending on the number of participating makers m and the value transferred by the taker

(a) Number of offers needed for $p = 0.75$ (b) Number of offers needed for $p = 0.9$

Value (quantile)	m			Value (quantile)	m		
	2	3	5		2	3	5
0.035 BTC (25 %)	8	12	19	0.035 BTC (25 %)	10	14	21
0.254 BTC (50 %)	8	12	19	0.254 BTC (50 %)	10	15	21
1.4 BTC (75 %)	8	11	17	1.4 BTC (75 %)	9	13	19

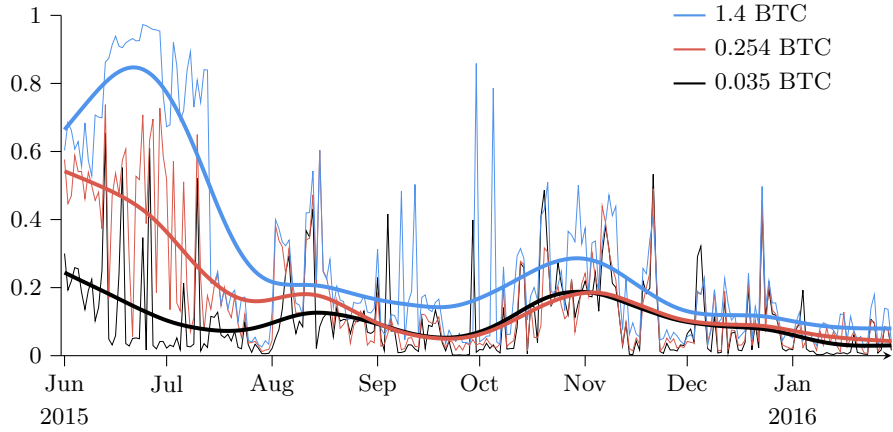


Fig. 6: Share of volume present in 90 % of all joins with two makers

BTC (around 32,000 USD) (cf. Table 2b). Note that this amount is not consumed but rather generates a profit during the attack. The actual cost of an attack is only the cost of capital. An attacker who operates in USD mainly faces exchange fees and possibly a risk premium for holding Bitcoin instead of USD.

Table 2: Cumulative value (in BTC) of the offers needed in order to be the sole participant in a CoinJoin transaction with probability p , depending on the number of participating makers m and the value transferred by the taker

(a) Cumulative value of offers for $p \geq 0.75$ (b) Cumulative value of offers for $p \geq 0.9$

Value (quantile)	m			Value (quantile)	m		
	2	3	5		2	3	5
0.035 BTC (25 %)	17.63	37.64	90.42	0.035 BTC (25 %)	24.77	55.01	104.59
0.254 BTC (50 %)	35.86	57.35	149.13	0.254 BTC (50 %)	45.23	80.01	163.57
1.4 BTC (75 %)	61.38	133.09	197.24	1.4 BTC (75 %)	81.73	161.70	207.16

These figures are aggregates over our period of observation. Figure 6 presents a longitudinal breakdown in relation to the overall value available on JoinMarket for the most common scenario with two makers and a success rate of 90 %. Although the share of capital necessary to perform the attack has dropped significantly from the early days, it still remains between 4–10 % in early 2016. At market volumes of 1,000–1,500 BTC, this yields a capital requirement in bitcoin equivalent to a medium 5-figure USD amount.

4 Economics of JoinMarket

In the last section, we have described the properties of this market for anonymity and estimated the capital needed for an economic attack against the implied matching mechanisms. In this section we draw on economic theory to explain why and under which condition the market exists. We compare empirical facts against stylized models to validate our explanation attempts.

4.1 Excess Supply

A market brings together demand and supply. Trades happen if agents on both sides differ in their preferences for bundles (q, v) composed of at least one tradable good (assumed homogeneous, divisible and with perfectly measurable quantity q) and monetary instruments of value v . Preferences change as trades happen, a relation typically expressed in demand and supply functions, which intersect at a market price p . Anonymity is a special good. Specifically, the service of participating in CoinJoin transactions increases q —think of the size of the anonymity set—for all involved parties. There is no simple explanation on why agents appear on both sides of a market. Why should some agents pay for the same service that others get paid for?

Indeed, the supply side of JoinMarket, that is where agents get paid, is more attractive. The red line in Figure 7 visualizes the maximum concurrent demand per day. We compute the maximum demand based on the transactions extracted from the transaction graph. We aggregate the concurrent demand as the sum of spend values multiplied by the number of chosen makers per block, and then select the largest value for each day. The graph shows that at no point in time the demand came close to the total supply (the same series shown in Figure 2 as Bitcoins available). This confirms our initial doubt.

The excess supply corresponds well to the observable race to the bottom on maker fees (cf. Figure 3). The following subsections model reasons why agents would appear on the demand side at all.

4.2 Time Preferences

Agents may differ in their time preferences. In a stylized model, agents can be divided into two types. Type 1 wants to make an anonymous payment soon, for example in exchange for goods and services. Type 2 has a longer time horizon and wants to generate some return on her capital in Bitcoin. Type 1 is willing to pay a premium for immediate service, hence agents of this type appear on the demand side.

To underpin this theory with empirical evidence, we use a proxy for both types. We assume that agents of Type 1 use coins to fund a CoinJoin transactions faster than agents of Type 2, which rely on the market to request their funds for use in a transaction. Because takers first have to send their funds to JoinMarket’s internal wallet in order to create a CoinJoin transaction, we compare the time difference between this funding and the spending transaction for a taker with the

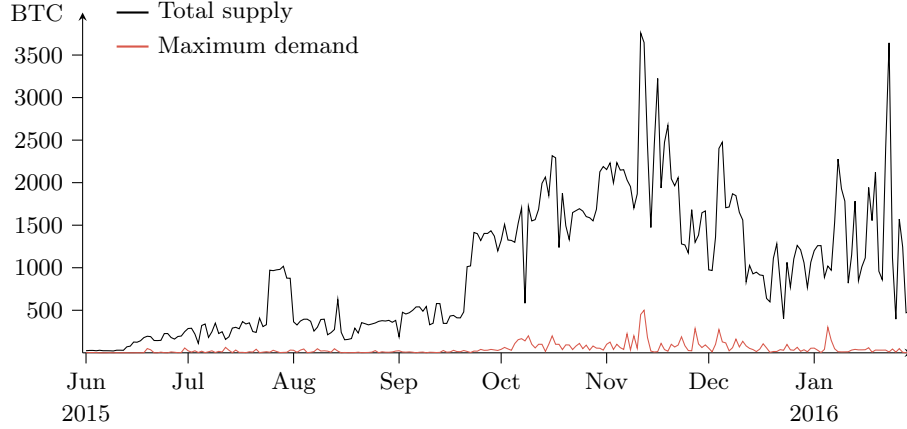


Fig. 7: Supply and demand on JoinMarket

time gap to the previous transactions for the maker. To tell apart takers' inputs from makers' inputs, we use subset matching (cf. [5, 26]). We could unanimously identify the taker's subset of inputs and outputs for 5,788 transactions (67%). In these cases, the taker's subset is the combination that loses value. The makers' subsets usually increase as they collect a maker fee from the taker.

Table 3 shows the time gap between the previous transaction and the CoinJoin transaction for both types of agents. The gap is measured in block intervals, Bitcoin's built-in time scale. One unit corresponds to just below 10 minutes on average. For most months, the average time until coins in JoinMarket's internal wallet are included in a CoinJoin transaction are significantly lower for Type 1, which lends some support to our explanation that takers are willing to pay for faster processing. The average time varies quite substantially between months, reflecting the unsteady environment in which JoinMarket operates. This calls for heavy econometrics if one wants to statistically test this hypothesis across time.

In general, the importance of differences in time preferences should be inversely proportional to the liquidity. The higher the frequency of trades the more predictable is the time until a reasonably priced offer is taken, and the more attractive becomes participation on the supply side.

Table 3: Idle time (in blocks) of inputs before entering JoinMarket transactions

	Jun'15	Jul'15	Aug'15	Sep'15	Oct'15	Nov'15	Dec'15	Jan'15
Type 1 (taker)	547	221	345	194	187	385	532	533
Type 2 (maker)	733	630	679	416	297	408	557	917
Diff. (maker – taker)	186	409***	334***	222***	110*	23	25	384***

p-values of two-sided *t*-tests: **p* < 0.05, ***p* < 0.01, ****p* < 0.001

4.3 Qualitative Differentiation

Time preference is not the only explanation. JoinMarket may indeed reflect the fact that bitcoins are not fungible. Each individual bitcoin can be traced back to the unique block in which it was created (“mined” in jargon). This allows market participants to tell bitcoins apart. Arguably, bitcoins can have different value depending on their history (e. g., [24]). For example, increasing adoption of risk scoring may make it harder to spend funds that can be associated with a specific activity (e. g., criminal offenses) or whose transaction history contains patterns suggesting such use in the past [31]. Bitcoins are thus differentiable on a quality dimension, denoted in our model by $z \in [0, 1]$. Even if this quality is not directly observable, the perceived quality along with agents’ expectations about the market valuation of coins of different quality can explain the existence of JoinMarket.

In the presence of qualitative differentiation, CoinJoin transactions, in particular those matched on open platforms like JoinMarket, suffer from adverse selection. Agents with known “good” coins ($z = 1$) must expect that joining with k random coins in the system, the expected quality (over the randomization of the selection) of the outgoing funds z' is

$$z' = \frac{1 + \bar{z} \cdot k}{1 + k}, \quad (1)$$

where \bar{z} is the average quality of all coins in circulation. This expectation is very optimistic, because it assumes that the average quality of the coins available in JoinMarket equals the average in the population. In practice, the decision to offer coins on JoinMarket is endogenous. Agents are more likely to offer known bad coins to bet on the chance of getting better ones. More precisely, an agent strictly improves his wealth if $z < z^*$, where z^* is the agents’ expectation about the average quality of the subpopulation of coins on JoinMarket. As other agents anticipate this behavior, a race to the bottom of z^* is triggered, leading to a collapse of CoinJoin platforms for good coins [4]. The only way out for a platform is to enable signaling (and commitment) to a minimum coin quality as part of the offers. JoinMarket did not support this in the study period.

In fact, the absence of signaling may indeed explain why agents appear on both sides. Besides impairing traceability, agents on the demand side may be willing to pay a premium for raising the average (or expected) quality of their funds. Agents on the supply side offer better coins and have (or expect) fewer difficulties in exchanging coins with lower quality for goods and services. This difference may be enough to outweigh the transaction costs of trade and hence justify the existence of JoinMarket.

We considered underpinning this theory with empirical data. However, we did not find a good proxy in our data to elicit private information about the (expected) quality from the participants. An avenue for future research is to test

this theory with data generated by Bitcoin risk scoring and intelligence firms, such as Elliptic⁵, Chainalysis⁶ or Scorechain⁷.

4.4 Other Explanations

The true reasons for trade activity on JoinMarket may be a combination of factors, including mundane ones. For example, operating a maker bot on JoinMarket’s supply side requires more technical sophistication than taking offers on the demand side. For some agents this transaction cost may be higher than the fees paid for receiving the service as taker. Related to this, the risk of operating a “hot wallet”⁸ may not be offset by the low maker fees. Moreover, information asymmetries, in particular a lack of understanding of the anomalies of information goods, may keep users on the demand side. To some extent, information disregarding the adverse selection problem discussed above adds to the information asymmetries:

“In JoinMarket, normal legitimate investors in bitcoin just want to earn their coinjoin fee. They probably bought their coins from Coinbase.com or bitstamp and only want to earn an extra few % over time. These are the people you’re mixing with when you create a coinjoin transaction. This makes JoinMarket unique, unlike any other private enhancer (DarkWallet, BitcoinFog, etc.) where you only mix with others who also want to improve their privacy.”

Source: [41]

5 Discussion

JoinMarket offers an innovative way of solving the matching problem for CoinJoin transactions. The current implementation is suboptimal in many respects.

On the technical side, the software is still in a premature state. It provides little protection against Sybil attacks, Denial-of-Service or privacy-invasive behavior of market participants. A malicious taker could, for example, repeatedly ask for funds from maker bots without using them in CoinJoin transactions. Knowing which outputs belong to the same taker then enables her to deanonymize other takers’ CoinJoin transactions by telling apart the funds of all participating makers [6]. There is also little protection against makers unwilling to sign transactions. Even if all market participants follow the protocol, the resulting CoinJoin transactions are identifiable in the block chain with ease, as demonstrated in Section 3.1.

The architecture is centralized around a single IRC server, with all known disadvantages, and quite in contrast to Bitcoin’s principle of decentralization.

⁵ <https://www.elliptic.co>, retrieved on 2016-03-04

⁶ <https://chainalysis.com>, retrieved on 2016-03-04

⁷ <https://scorechain.com>, retrieved on 2016-03-04

⁸ Hot wallets are installations where the private key is entrusted to a device that is connected to the Internet 24/7. Hot wallets are worthwhile targets for cybercriminals.

As people entrust this market tens and hundreds of bitcoins, it may only be a matter of time until serious issues arise; possibly adding another data point to the list of incidents where Bitcoin users lost money (cf. [28]). It remains an open research question to design a robust, accountable, decentralized matching market for CoinJoin transactions which offers some principled privacy guarantees.

On the economics side, JoinMarket leaves us puzzled on why trades exist on a matching market where makers and takers improve their anonymity alike. Moreover, many details of the market mechanism seem to follow ad-hoc approximations, for example the default offer selection method. It would be interesting to study the special properties of markets for anonymity through the lens of mechanism design.

A relevant cost factor in anonymizing Bitcoin transactions are miner fees. They internalize the externalities of securing the public ledger and must be paid for each transaction in order to be included in the blockchain. Whether or not the system can maintain a fee level low enough for consumer transactions is a hot debate at the time of writing (cf. [15, 29]). On JoinMarket, takers bear all of the miner fees, which often greatly exceed the maker fees. If miner fees rise, JoinMarket's future is uncertain because anecdotal evidence suggests that takers often use multiple CoinJoin's, hoping to increase the size of the anonymity set.

Finally, we note that participation in decentralized systems is often motivated by political or altruistic beliefs. One user described their motives in JoinMarket's public IRC channel as: "I am not really trying to make a profit. I just want my coins to do some good work and maybe help the joinmarket network." These non-economic factors are hard to capture in an economic framework, but might be necessary in order to explain otherwise seemingly irrational behavior.

6 Conclusion

This paper documents the first study on markets for anonymity, an interesting and rather unexplored phenomenon, which we found in a niche (JoinMarket) of a niche (cryptographic currencies).

Our longitudinal measurement study reveals a growing market for anonymous bitcoin transfers, funded with multiple thousand bitcoins at an average fee below 0.01%. We find evidence for JoinMarket transactions worth almost 8 million USD in eight months. The interpretation of this number is not straightforward because takers may go through multiple JoinMarket transactions subsequently to increase their anonymity. The total amount anonymized is unknown, but likely only a fraction of the headline figure.

We propose and investigate several economic explanations for the existence of such a market for anonymity. In accordance with our model, we observe excess supply and willingness to pay for faster anonymous payments on the demand side. It remains unclear if the agents are aware of qualitative differences in coins being traded, and some public statements suggest the contrary [41].

Anonymity in society has a Janus face. Undoubtedly, cryptographic currencies facilitate some forms of criminal activity [10, 13]. In this sense, JoinMarket is the

cheapest way to launder money in Bitcoin we are aware of. It inherits the security of CoinJoin transactions against fraudulent counterparts and intermediaries. At the same time it inherits the lack of trust of markets over hierarchies [40], which backfires if the anonymity of transaction flows becomes a design goal next to security. We have shown that it is possible and affordable for realistic attackers to “buy” themselves into the CoinJoin transactions matched on JoinMarket. If law enforcement agencies consider this a viable and legitimate option, they should be careful to coordinate among each other. The attack scenarios described in this paper apply to cases with a single attacker only. If multiple attackers compete on the market, they may thwart each other’s effort pretty effectively. With the data at hand, we cannot exclude that the relatively high headline figure of available offers in the order of 800,000 USD is already inflated by such operations.

Acknowledgments The authors thank Chris Belcher for his feedback on an earlier version of this manuscript and Christian Rückert for useful discussions. This work is funded by the German Bundesministerium für Bildung und Forschung (BMBF) under grant agreement No. 13N13505.

Bibliography

- [1] Alessandro Acquisti, Laura Brandimarte, and George Loewenstein. “Privacy and Human Behavior in the Age of Information”. In: *Science* 347.6221 (2015), pp. 509–514.
- [2] Alessandro Acquisti, Roger Dingledine, and Paul Syverson. “On the Economics of Anonymity”. In: *Financial Cryptography and Data Security*. Ed. by Rebecca N. Wright. Vol. 2742. Lecture Notes in Computer Science. Springer Berlin Heidelberg, 2003, pp. 84–102.
- [3] Alessandro Acquisti and Hal R. Varian. “Conditioning Prices on Purchase History”. In: *Marketing Science* 24.3 (2005), pp. 367–381.
- [4] George A. Akerlof. “The Market for ”Lemons”: Quality Uncertainty and the Market Mechanism”. In: *Quarterly Journal of Economics* 84.3 (1970), pp. 488–500.
- [5] Kristov Atlas. *Weak Privacy Guarantees for SharedCoin Mixing Service*. 2014. URL: <http://www.coinjoinsudoku.com/advisory/> (visited on 2016-03-04).
- [6] Chris Belcher. *Bad Faith Taker Spy Not Filling Orders So That It Learns Which UTXOs Belong to Which Maker, Allowing Future Unmixing*. 2015. URL: <https://github.com/JoinMarket-Org/joinmarket/issues/156> (visited on 2016-02-22).
- [7] Chris Belcher. *JoinMarket release on mainnet*. 2015. URL: https://www.reddit.com/r/joinmarket/comments/358dlv/joinmarket_released_on_mainnet/ (visited on 2016-03-04).
- [8] George Bissias, A. Pinar Ozisik, Brian N. Levine, and Marc Liberatore. “Sybil-Resistant Mixing for Bitcoin”. In: *Proceedings of the 13th ACM Workshop on Workshop on Privacy in the Electronic Society*. ACM, 2014.

- [9] Blockchain. *Shared Coin*. URL: <https://www.sharedcoin.com/> (visited on 2016-03-04).
- [10] Rainer Böhme, Nicolas Christin, Benjamin Edelman, and Tyler Moore. “Bitcoin: Economics, Technology, and Governance”. In: *Journal of Economic Perspectives* 29.2 (2015), pp. 213–238.
- [11] Joseph Bonneau, Andrew Miller, Jeremy Clark, Arvind Narayanan, Joshua A. Kroll, and Edward W. Felten. “Research Perspectives and Challenges for Bitcoin and Cryptocurrencies”. In: *2015 IEEE Symposium on Security and Privacy*. San Francisco, CA, USA, May 2015.
- [12] Joseph Bonneau, Arvind Narayanan, Andrew Miller, Jeremy Clark, Joshua A. Kroll, and Edward W. Felten. “Mixcoin: Anonymity for Bitcoin with Accountable Mixes”. In: *Financial Cryptography and Data Security*. Ed. by Nicolas Christin and Reihaneh Safavi-Naini. Vol. 8437. Lecture Notes in Computer Science. Springer Berlin Heidelberg, 2014, pp. 486–504.
- [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*. Rio de Janeiro, 2013, pp. 213–224.
- [14] CoinDesk. *Bitcoin Price Index*. URL: <http://www.coindesk.com/price/> (visited on 2016-02-23).
- [15] Kyle Croman, Christian Decker, Ittay Eyal, Adem Efe Gencer, Ari Juels, Ahmed Kosba, Andrew Miller, Prateek Saxena, Elaine Shi, Emin Gün Sirer, Dawn Song, and Roger Wattenhofer. “On Scaling Decentralized Blockchains”. In: *3rd Workshop on Bitcoin and Blockchain Research*. Barbados, 2016.
- [16] *Darkwallet*. URL: <https://www.darkwallet.is/> (visited on 2016-03-04).
- [17] Roger Dingledine and Nick Mathewson. “Anonymity Loves Company: Usability and the Network Effect”. In: *Proceedings of the Fifth Workshop on the Economics of Information Security (WEIS 2006)*. Ed. by Ross Anderson. Cambridge, UK, June 2006.
- [18] John R. Douceur. “The Sybil Attack”. In: *Peer-to-Peer Systems*. Ed. by Peter Druschel, Frans Kaashoek, and Antony Rowstron. Vol. 2429. Lecture Notes in Computer Science. Berlin Heidelberg: Springer, 2002, pp. 251–260.
- [19] Serge Egelman, David Molnar, Nicolas Christin, Alessandro Acquisti, Cormac Herley, and Shriram Krishnamurthi. “Please Continue to Hold: An Empirical Study on User Tolerance of Security Delays”. In: *Proceedings of the 9th Workshop on the Economics of Information Security (WEIS 2010)*. Cambridge, MA, 2010.
- [20] *JoinMarket-Org/joinmarket: CoinJoin Implementation with Incentive Structure to Convince People to Take Part*. URL: <https://github.com/JoinMarket-Org/joinmarket> (visited on 2016-03-01).
- [21] Stefan Köpsell. “Low Latency Anonymous Communication – How Long Are Users Willing to Wait?” In: *Emerging Trends in Information and Communication Security (ETRICS 2006)*. Ed. by Günter Müller. Vol. 3995.

- Lecture Notes in Computer Science. Springer Berlin Heidelberg, 2006, pp. 221–237.
- [22] Kenneth C. Laudon. “Markets and Privacy”. In: *Communications of the ACM* 39.9 (1996), pp. 92–104.
 - [23] Michael Lesk. “The Price of Privacy”. In: *IEEE Security & Privacy* 10.5 (September 2012), pp. 79–81.
 - [24] *Looking to buy an old 50 BTC block. Where to buy?* 2015. URL: https://np.reddit.com/r/Bitcoin/comments/3sg8vm/looking_to_buy_an_old_50_btc_block_where_to_buy/ (visited on 2015-03-04).
 - [25] Gregory Maxwell. *CoinJoin: Bitcoin Privacy for the Real World*. 2013. URL: <https://bitcointalk.org/index.php?topic=279249.0> (visited on 2016-03-04).
 - [26] Sarah Meiklejohn and Claudio Orlandi. “Privacy-Enhancing Overlays in Bitcoin”. In: *Financial Cryptography and Data Security, 2nd Workshop on BITCOIN Research*. Ed. by Michael Brenner, Nicolas Christin, Benjamin Johnson, and Kurt Rohloff. Vol. 8976. Lecture Notes in Computer Science. Berlin Heidelberg: Springer, 2015, pp. 127–141.
 - [27] 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: *USENIX ;login:* 38.6 (2013).
 - [28] Tyler Moore and Nicolas Christin. “Beware the Middleman: Empirical Analysis of Bitcoin-Exchange Risk”. In: *Financial Cryptography and Data Security: 17th International Conference, FC 2013, Okinawa, Japan, April 1-5, 2013, Revised Selected Papers*. Ed. by Ahmad-Reza Sadeghi. Vol. 7859. Lecture Notes in Computer Science. Springer Berlin Heidelberg, 2013, pp. 25–33.
 - [29] Malte Möser and Rainer Böhme. “Trends, Tips, Tolls: A Longitudinal Study of Bitcoin Transaction Fees”. In: *Financial Cryptography and Data Security, 2nd Workshop on BITCOIN Research*. Ed. by Michael Brenner, Nicolas Christin, Benjamin Johnson, and Kurt Rohloff. Vol. 8976. Lecture Notes in Computer Science. Berlin Heidelberg: Springer, 2015, pp. 19–33.
 - [30] Malte Möser, Rainer Böhme, and Dominic Breuker. “An Inquiry into Money Laundering Tools in the Bitcoin Ecosystem”. In: *Proceedings of the APWG E-Crime Researchers Summit*. San Francisco: IEEE, 2013, pp. 1–14.
 - [31] Malte Möser, Rainer Böhme, and Dominic Breuker. “Towards Risk Scoring of Bitcoin Transactions”. In: *Financial Cryptography and Data Security, 1st Workshop on BITCOIN Research*. Ed. by Rainer Böhme, Michael Brenner, Tyler Moore, and Matthew Smith. Vol. 8438. Lecture Notes in Computer Science. Berlin Heidelberg: Springer, 2014, pp. 16–32.
 - [32] Satoshi Nakamoto. *Bitcoin: A Peer-to-Peer Electronic Cash System*. 2008. URL: <https://bitcoin.org/bitcoin.pdf> (visited on 2016-03-04).
 - [33] Andreas Pfitzmann and Marit Köhntopp. “Anonymity, Unobservability, and Pseudonymity – a Proposal for Terminology”. In: *Designing Privacy*

- Enhancing Technologies*. Ed. by Hannes Federrath. Vol. 2009. Lecture Notes in Computer Science. Berlin Heidelberg: Springer, 2001, pp. 1–9.
- [34] Richard A. Posner. “The Economics of Privacy”. In: *The American Economic Review* 71.2 (1981), pp. 405–409.
 - [35] Fergal Reid and Martin Harrigan. “An Analysis of Anonymity in the Bitcoin System”. In: *Security and Privacy in Social Networks*. Ed. by Y. Altshuler, Y. Elovici, A.B. Cremers, N. Aharony, and A. Pentland. New York: Springer, 2013, pp. 197–223.
 - [36] Dorit Ron and Adi Shamir. “Quantitative Analysis of the Full Bitcoin Transaction Graph”. In: *Financial Cryptography and Data Security*. Ed. by Ahmad-Reza Sadeghi. Vol. 7859. Lecture Notes in Computer Science. Berlin Heidelberg: Springer, 2013, pp. 6–24.
 - [37] Tim Ruffing, Pedro Moreno-Sanchez, and Aniket Kate. “CoinShuffle: Practical Decentralized Coin Mixing for Bitcoin”. In: *ESORICS’14*. Proceedings of the 19th European Symposium on Research in Computer Security. Vol. 8713. Lecture Notes in Computer Science. Berlin Heidelberg: Springer, 2014, pp. 345–364.
 - [38] Sarah Spiekermann. “The Desire for Privacy: Insights into the Views and Nature of the Early Adopters of Privacy Services”. In: *International Journal of Technology and Human Interaction (IJTHI)* 1.1 (2005), pp. 74–83.
 - [39] J. Miguel Villas-Boas. “Dynamic Competition with Customer Recognition”. In: *The RAND Journal of Economics* 30.4 (1999), pp. 604–631.
 - [40] Oliver E. Williamson. “Markets and Hierarchies: Some Elementary Considerations”. In: *American Economic Review* 63.2 (1973), pp. 316–325.
 - [41] *With JoinMarket, You Mix With Clean, Untainted Bitcoins*. 2015. URL: https://www.reddit.com/r/joinmarket/comments/38fi5m/with_joinmarket_you_mix_with_clean_untainted/ (visited on 2015-03-04).
 - [42] Pieter Wuille. *Hierarchical Deterministic Wallets*. 2012. URL: <https://github.com/bitcoin/bips/blob/master/bip-0032.mediawiki> (visited on 2016-03-04).
 - [43] Jan Henrik Ziegeldorf, Fred Grossmann, Martin Henze, Nicolas Inden, and Klaus Wehrle. “CoinParty: Secure Multi-Party Mixing of Bitcoins”. In: *Proceedings of the 5th ACM Conference on Data and Application Security and Privacy*. San Antonio, Texas, USA: ACM, 2015, pp. 75–86.

A Appendix

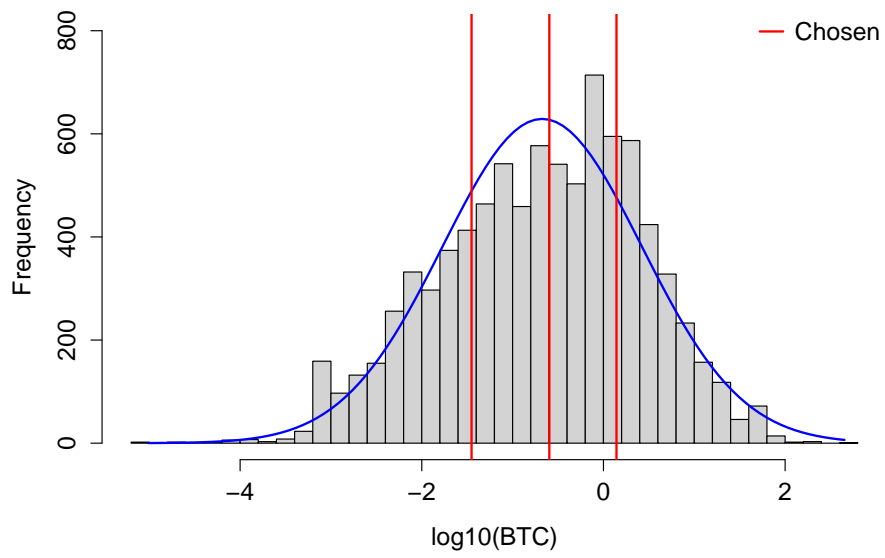


Fig. 8: Spend values chosen by the takers follow a log-normal distribution