Challenges and Opportunities Associated with a Bitcoin-based Transaction Rating System

David Vandervort
Xerox
david.vandervort@xerox.com

Abstract. It has been shown that seller ratings given by previous buyers give new customers useful information when making purchasing decisions. Bitcoin, however, is designed to obfuscate the link between buyer and seller with a layer of limited anonymity, thus preventing buyers from finding or validating this information. While this level of anonymity is valued by the Bitcoin community, as Bitcoin moves toward greater adoption there will be pressure from buyers who wish to know more about who they are doing business with, and sellers who consider their reputation a strong selling point, to allow greater transparency. We consider three different models by which a reputation/rating system could be implemented in conjunction with Bitcoin transactions and consider pros and cons of each. We find that each presents challenges on both the technological and social fronts.

1 Background

Bitcoin is electronic currency. This fact has consequences for how people think about it and how they regard doing business with it. They expect to do business in a businesslike way. The promise of Bitcoin is that transactions will be quick and frictionless. Unlike businesses using more established and controlled currencies, they may also be anonymous. The Bitcoin protocol provides for monetary and other transactions to be tied to addresses, not identities. The ability to generate and use new addresses provides a degree of anonymity that non-technically inclined users are unlikely to pierce. The Bitcoin community points to this anonymity, often referred to as pseudonymity because it is not absolute, as an asset, a way of circumventing surveillance and cumbersome regulatory regimes [1]. It is also considered a defense against the user profiling/data mining practiced by large merchants such as K-Mart.

The original description of Bitcoin described the features of non-reversibility of payments and cryptographic verification as substitutes for more traditional forms of ensuring trust [2]. However, the anonymity of Bitcoin transactions removes another source of trust in transactions: Knowing who you are doing business with. It opens the possibility of counter-party risk, the danger that the other person will not live up to their obligations once the money is received. For buyers, being able to know who they are buying from can increase trust and so reduce barriers to a successful purchase. Likewise, sellers whose business grows by word of mouth or who consider good will a valuable asset may wish to share information about successful transactions with the buying public.

Bitcoin allows different forms of knowledge and openness from traditional forms of transaction. In the pre-cryptocurrency market, transactions are mostly anonymous

and away from public view. When a buyer pays with cash, the record of the transaction contains only the amount and codes related to the items purchased. The use of a credit card creates more of a paper trail as the credit card service (i.e. Visa), the vendor and the purchaser retain records identifying both buyer and seller, as well as items and amounts. However, this information is generally not published for the world to see. Such financial records may be bought and aggregated in many ways, including scoring the credit worthiness of the buyer. They may be mined by the seller to learn much about their customers but, since virtually all companies consider financial information to be confidential, the results are not shared with customers or with other companies.

Studies have shown that buyer behavior can be influenced by knowledge of seller reputation, product quality, and experiences and sentiment of previous customers. Specifically, simulations have shown that reputation systems can improve the overall quality of an online market [3]. The Bitcoin protocol includes as a central feature, the blockchain, a public ledger of all transactions. Amounts paid, change received, and both input and output addresses are stored and cryptographically verified. The information contained therein, however, is sparse. An address is no more than a temporary identity. By intent it contains no direct links to more stable and human understandable identities. Likewise, the exact nature of the goods or services purchased is absent.

As Bitcoin and associated altcoins become more mainstream and their communities attempt the leap to general acceptance, the sparsity of information may become an issue. There is some likelihood that some buyers and sellers will seek ways to leverage and augment the information on the blockchain to learn more about each other. The ability to mine the blockchain for data is part of the promise of Bitcoin. This paper explores three ways in which such augmented information could be developed, looking specifically at the case where buyers seek information about sellers and sellers voluntarily cooperate. The three cases considered are site based systems, wallet based systems and coin based systems.

2 Characteristics of Rating Systems

For purposes of this discussion we define a rating system as a means for customers to provide feedback on a purchase that future potential customers may access and include in their decision making process. Other purposes, such as for sellers to track customer satisfaction or for regulatory bodies to monitor performance, while potentially possible are not considered here. Ratings may take different forms, such as a number of stars, a thumbs up/thumbs down, or a numerical rating. Some systems may aggregate multiple ratings. For example, a transaction may be given a separate rating in each of several categories such as price, speed of delivery and product quality. Free form product reviews are also common. Whatever method is used, the goal is to encode some concept of quality associated with the transaction, product or provider in a form that can be easily understood by future customers. This requirement does not rule out machine mediation of information, such as developing

an average of ratings or converting numbers into words (1 ~ 'Excellent') that are considered more readable by humans.

The potential for the computer to repackage information is relevant to the case of Bitcoin transactions. In raw form, a Bitcoin transaction contains very little data that a human finds meaningful, particularly if that human is an average, non-technical user. A Bitcoin transaction is essentially little more than a hash code that is used to identify a collection of other hash codes, with numbers representing transaction amounts, attached. Software can easily break this down into a set of inputs and outputs, showing addresses Bitcoins are drawn from and sent to. There is little meta-data that would identify the purpose of the transaction or the entities involved (that is, the owners of the associated addresses) is not present. This is an efficient system for transmitting the required information. Any system for rating Bitcoin transactions must contend with this efficiency, preferably without compromising it.

Online rating systems are subject to several weaknesses that can be exploited by malicious parties (Josang, et al). One such weakness is falsified ratings. That is, friends or enemies of a seller may post fake ratings (possibly using multiple false identities) intended to boost or suppress a seller's reputation regardless of whether they had ever engaged in business with the seller, or of the outcome if they had done business. It may also be possible for malicious intruders in a system to tamper with authentic ratings, making them invalid (Yao, Ruohomaa and Xu). It is essential for any rating system, including one involving Bitcoin transactions, to address these weaknesses.

It should be noted that rating systems are often referred to in the literature as reputation systems. This terminology is not used here largely because the question of whether ratings can be reliably related to sellers is still an open one. Therefore in this paper we use the more limited terminology of rating systems.

3 Considerations for a Bitcoin-Based Rating System

The heart of the Bitcoin protocol is the blockchain, the ledger of all transactions that is maintained by the entire Bitcoin network. Cryptographic proof-of-work makes it extremely difficult to forge transactions and ensures a high degree of integrity to the blockchain. Because the blockchain "remembers" all Bitcoin transactions, it is large and continuously growing. In the context of transaction ratings this means that adding rating meta-data directly into the blockchain may cause enough increase in the storage and transmission size of transactions to be impractical.

A more salient point than data size is the speed at which transactions become "fixed" in the blockchain. One of the attractions of Bitcoin is that transactions are very quick, with confirmations coming from around the network within minutes. Confirmation means that some number of nodes around the network have verified the cryptographic proof-of-work associated with the transaction. It also means that, once confirmed, the transaction cannot be changed. Customers who provide feedback on a transaction may wait until a product is delivered, or until it is installed, or until it breaks. The amount of time between transaction and rating varies but in virtually all cases will be greater than brief period between creation of a transaction and

confirmation. The consequence of this is straight forward. Ratings of a transaction cannot be included in the data for that transaction.

It may be possible to create a new kind of transaction that references an older one and adds meta-data to it. However, there is no such proposal at this time and any scheme that depends on such a proposal being implemented should be considered difficult to implement. No such proposal will be considered here.

The anonymity-by-intent nature of the Bitcoin protocol is a major consideration. The official Bitcoin client software can generate any number of random addresses that can be used to send or receive Bitcoin payments.

Nothing about any given address links it to the entity using it, although there are techniques by which addresses may sometimes be de-anonymized. For example, if several addresses are inputs to the same transaction they are all known to be owned by the same entity. Re-use of any of the addresses therefore reveals information about all of them [4]. This has the effect on a rating system of making aggregation of ratings difficult or impossible. This is potentially a greater hurdle to a strong rating system than the immutable nature of the blockchain. Any attempt at a rating system must de-obfuscate the relationship between addresses and the parties in a transaction. A rating must attach to a seller and come from a buyer, not merely an address.

The above considerations are largely technical. There are also social considerations. The Bitcoin protocol does not merely allow anonymous transactions, it is intentionally built on them. The community up to now has seen this as a virtue. However, for a rating system to be meaningful, ratings must attach to some identity. Therefore it is a given that participating sellers must give up some privacy. In deference to the values of the community, then, any rating system must be opt in, meaning that sellers who do not wish to be rated are not required to accept it and buyers are never required to provide ratings. However, those sellers who do opt in, should be prevented or discouraged from being rated selectively, cherry picking only the good transactions, or otherwise subverting the quality of ratings.

Likewise, buyers who wish to anonymously submit ratings should be permitted to do so, within the constraints of the need to verify the authenticity of the rating. This means that any effective system must still be able to weed out bogus ratings as described above. Linking all ratings to an origin address corresponding to previous transactions may fulfill this requirement. However, the anonymity of ratings to sellers, who may have records of the buyer's identity (for example, in a shipping address) likely cannot be protected. Whether ratings might also be linked back to the buyer's identity by third parties who have access both to the ratings themselves and to the blockchain is a significant question that should be carefully explored before implementing any rating system.

4 Rating System Models

The constraints and capabilities of the Bitcoin protocol result in three major ways for a rating system to be built. These are at the website level, the wallet level and the coin level. The three methods are not mutually exclusive. The following sections will describe what each of these means and the strengths and weaknesses of each.

4.1 Site Based Systems

Site based rating systems are already common on the Internet, with the prototypical example being the seller ratings on eBay. These systems evolved to allow a measure of trust between non-local persons doing business with strangers (Resnick). The basic mechanism is therefore known. The question becomes, is it possible for such systems to extend to payments via Bitcoin? In one sense, Bitcoin is identical to any other currency. Marketplace websites ordinarily force payments to be funneled through the site interface. The alternative would be for checks to be mailed or funds to be transferred outside the system. This introduces the possibility that the seller will fail to deliver goods after payment is received or even to notify site and buyer of receipt. Keeping payments inside the site's control removes this friction.

Facilitating Bitcoin payments within a market requires that buyer and seller each establish a Bitcoin wallet (aka an account) in the site and stock it with funds, just as they would if paying with dollars. Several sites do exactly this, including localBitcoins.com, a Bitcoin trading platform and bidinBitcoins.com, a merchandise auction site. The downside of this solution is that it is not portable, meaning that a seller could easily engage in bad behavior on one site, receive bad ratings, and migrate to a different site where the ratings will not be available. This problem is independent of Bitcoin and is a result of the general anonymity of the Internet.

Web site ratings systems are usually not anonymous. They are instead tied to buyer identity on the site, which is also typically tied to an email address. Bitcoinary.com, a web site where users buy and sell Bitcoins, is an example of a site based rating system that goes to extra lengths to verify that users have some real identity beyond the site itself by linking accounts to social media profiles such as Twitter, Facebook and LinkedIn. In order to preserve anonymity, however, the site does not reveal the details of those profiles to other users. It merely indicates that they exist.

Identity verification protects the site and its users from fake ratings but creates the potential for exposure due to careless programming or malicious attack on the database. The site is a target for such attacks simply by virtue of the presence of currency. Exposure of user information does not directly compromise their complete Bitcoin portfolio, if they maintain a separate wallet outside the market site. The Bitcoin addresses used within the site remain separate from those generated elsewhere. It is possible, though, to develop a site requiring user's to input addresses generated by an external wallet for use in transactions. If the addresses so used were ever used from the wallet, links would be created from the user's identity on the site to the rest of their wallet. Further, any such compromise of identity would link that user's ratings of others to transactions outside the site. This partial identity compromise should be considered when designing market sites or when deciding whether to open an account on one.

4.2 Wallet Based Systems

A Bitcoin wallet is software that stores Bitcoins, generates addresses and sends and receives Bitcoin transactions. Wallets can reside on a computer hard drive, on

specialized hardware or on a mobile device such as a smartphone. Users can manage their own or can sign up for online wallet services (such as coinbase.com) that manage the details for them. The software that comprises the wallet could be written in a way that maintains meta-data about addresses and transactions as well as those items themselves. This development is already taking place. The official Bitcoin wallet allows users to tag Bitcoin addresses in its address book and to refer to the tags when creating a transaction. The Electrum wallet (electrum.org) has a slightly more sophisticated graphical user interface that resembles a check register, with a field for the user to enter the payee and a description line. This allows linking of sellers to their addresses, even when the address changes with every transaction. Adding a field for a transaction rating to this type of interface is trivial. What is done with the ratings afterwards is of more interest.

This wallet level rating suffers from an even worse form of the siloing problem than site based systems in that ratings are not shared with other users at all. This does not necessarily make them useless, as they can still be an aid to a buyer's memory. Their usefulness is still limited in that they provide no information about sellers previously unknown to the current user. Likewise, if the wallet software instituted a peer-to-peer exchange of such information, it would end up being shared only with other users of the same client, making them little different from self-contained web systems. As the payee and description lines are (currently) hand-entered, spelling and typos become an issue that makes aggregation of scores more difficult unless commonalities of seller Bitcoin addresses can be used to resolve them.

One development that may work to ease the problem of aggregating seller information even across clients is in the proposals related to the Bitcoin payment protocol. Specifically, Bitcoin Improvement Proposal number 72 (https://github.com/bitcoin/bips/blob/master/bip-0072.mediawiki) would create a new link type that would be embedded on a web page or in the signature of an email. When clicked, one of these links would initiate a Bitcoin payment. The web address of the seller would be included in the meta-data on the link. This proposal is currently in draft state and has not been enabled in the Bitcoin protocol. Developments like this reduce seller anonymity while increasing ease of use and also provide data that could be used by payment systems.

4.3 Coin Based Systems

Bitcoin is a protocol as well as a currency. Many other currencies have been created using variations of the protocol. Not all of these variants (called altcoins) are intended for use as currency. Namecoin, for example, uses a blockchain to store arbitrary name data, used to create an alternative domain name system for finding sites on the Internet. A system has been proposed for extending Namecoin to the registration of certificates, using the value of the currency as a proxy for trust [5]. "Colored coins" is a proposal to add a meta-data layer to Bitcoins that would convert them into some other type of asset, such as a stock or bond [6]. Inserting ratings directly into a Namecoin or as coloration to a Bitcoin, with the hash code of the transaction referred to included in the data, could provide several benefits over wallet and site based systems. The first is that, like the Bitcoin blockchain itself, it would be public record, accessible to anyone with an Internet connection and the right software. Another

benefit is that, again like Bitcoin transactions, ratings so recorded would be immutable. Once confirmed, a transaction is a part of the blockchain forever (barring a blockchain fork, which is a rare occurrence so far).

A type of cryptocoin dedicated entirely to storing transaction ratings on its own blockchain could be designed and integrated into existing wallets (including site-based wallets). This would remove the need to update the Bitcoin protocol to accommodate the new data but would add the problem of keeping the rating blockchain in sync with the Bitcoin blockchain. This is both an opportunity and a technical hurdle. The public nature of the Bitcoin blockchain means it can be used not only to verify that a transaction has taken place but that the buyer address and seller address referred to in a transaction rating were actually also involved in the transaction. This can greatly curb fake ratings. Unfortunately, it means that sellers and buyers who use multiple addresses can frustrate the system. Some method of tying addresses to identities could mitigate this problem but would run directly counter to the Bitcoin design philosophy.

The coin based rating system suffers from technical problems in coin generation and distribution. Bitcoins themselves are created by software that solves cryptographic problems and is rewarded with currency. There is a strict upper limit to the amount of Bitcoins that can be created, meaning there will come a time when there can be no more. Would rating coins have any economic value? How would someone who wishes to rate someone acquire the coins? If they have value, then acquiring them will either require mining them as Bitcoins are mined, buying them, or being paid in them. These activities burden the blockchain with multiple types of transactions requiring multiple types of processing as well as begging the question, how much money is a rating worth? Add these questions to the need for 2-step verifications of ratings and the entire system may become unwieldy.

A coin based system also contains no defense against sellers who use multiple addresses for transactions. By associating a rating with a transaction, only the addresses used in that transaction are marked. Sellers wishing to aggregate their ratings to show their good reputation would need some mechanism to register all the addresses they use as associated with a single identity. Meanwhile, those sellers who wished to hide from their ratings would have a ready method for doing so, by simply switching addresses frequently and not relating them to their own identity, or any other.

5 Comparison

A summary of the characteristics of the different approaches is displayed in table 1. In the table, verifiability refers to the ability to ensure that ratings are allowed only for real transactions. The public blockchain makes this easily enforceable. A flaw in this view is that Bitcoin is highly divisible into units of only 0.00000001 Bitcoin, an amount so small as to be virtually without value at current exchange rates. Nothing is to stop an attacker from sending very tiny payments to an address known to be associated with a seller in order to gain the ability to provide ratings.

The information sharing entries in the table refer to how public ratings are once created and how easily they can be aggregated. This is a difficult property to achieve in a system where a new payment address can be created for every transaction. This second issue is why coin based systems are listed as having only limited information sharing. While transactions are completely public, aggregation can be made much more difficult by the use of multiple addresses.

The write-once entry in the table refers to whether rating entries are immutable once created. This can be seen as a measure of data integrity. A database entry is entirely mutable if compromised by malicious or dishonest entities, therefore site and wallet based systems are seen as lacking this property. A blockchain entry, as in a coin based system, once confirmed, is immutable because of the cryptographic proof involved.

The entries for buyer and seller verifiability are references to the ability to verify that buyer and seller are real entities, rather than bogus identities created solely for the purpose of corrupting the system with fake ratings or even fake transactions. This is one of the most difficult properties to verify because of the difficulty distinguishing casual or infrequent users from those whose intent is not to participate at all.

Table 1: Comparison of rating system types

Characteristic	Site Based	Wallet Based	Coin Based
Verifiability	Yes	Yes	Yes
Information sharing	Limited	Limited	Limited
Write-once data	No	No	Yes
Buyer Verifiability	Yes	N/A	No
Seller Verifiability	Yes	No	No
Distributed control	No	Possibly	Yes

The final row in the table considers distributed control of ratings. Systems in which one entity or a small group of entities stores or aggregates ratings for the whole network have centralized, rather than distributed, control. A distributed system, where all entities on the network are equally vested in the rating system would be most in keeping with the values of the Bitcoin community, however it may also be the most difficult to implement. Site and wallet based systems, in general, are seen as examples of single entity control, though it is possible for them to cooperate in a distributed fashion. No such cooperation can be assumed, however.

All of the given characteristics of a rating system are desirable if the system is to be of highest value. It can be seen that none of the considered architectures possesses all of these characteristics, thus none is a total solution. Social measures may address some of the weaknesses. For example, a standard could be created by which sellers agree to publish all of their addresses. This would aid in aggregation of results to provide something like a true reputation system. Software to manage transactions could verify transaction addresses against the published addresses as well. Creating a common method of publishing that is discoverable by all buyers would need to be very carefully crafted, however.

Business solutions may also be possible. Sellers might engage third parties to audit their practices and validate that they are adhering to policies of identity transparency. Public posting of audit results could give buyers some confidence that sellers are not

gaming a rating system. This sort of heavy handed method would be expensive, however, and is clearly contrary to the values of the Bitcoin community. There is room for doubt that it would be widely accepted.

No system to police sellers can provide control of buyers. Protecting a rating system from abuse of sellers with fake ratings is an equally important issue. A coin based system, where ratings actually carry some small cost, in conjunction with minimum costs for the transactions that can be rated, may at least provide a disincentive for abuse. Such a pay-to-play system would require a balancing act between discouraging bad actors and encouraging good ones.

6 Conclusion and Future Research

Each of the systems described has strengths and weaknesses. None is a complete solution to the need for information about sellers. The structure of the Bitcoin protocol intentionally makes this difficult. The ease with which addresses are created and discarded makes it relatively simple for buyers and sellers, both, to maintain several identities for transactions, or to simply disappear into the blockchain, their Bitcoin balances known but never their habits. Therefore in each of the three systems the question of seller identification was touched on. The fact remains that honest sellers do have an incentive to allow themselves to be rated. Particularly within the context of an anonymous system like Bitcoin, the simple fact of submitting to ratings demonstrates a degree of good faith.

As described above, web-based rating systems for Bitcoin transactions already exist. Web sites, however, maintain their own infrastructure for doing so, without using the strengths of the Bitcoin blockchain in verification of transactions and in anchoring ratings to specific transactions. The strengths of the blockchain are in its public nature and in the strong cryptographic proof that its contents are valid. These strengths should also make a strong rating system. A coin-based system is most directly designed to capitalize on these strengths. One potential area of future research would be in solving the problems of multi-blockchain synchronization. Another is in designing transactions that coordinate the potential monetary value of colored coins, Namecoins and others, with entirely different uses such as transaction ratings.

The question of anonymity is one that none of the technologies considered is well able to handle. A web-based system is in the best position of any of the three to impose standards on sellers but may find enforcement extremely difficult. Transactions taking place outside the system, even if still with Bitcoin, will be unknown to the system and difficult or impossible to trace back to the specific seller.

A hybrid system, that combines two or all of the systems described in this paper is worthy of further investigation. Could a site based rating system interact with external wallets? Could a coin-based system be designed that would interact with both? At a higher level, what is the minimum level of adoption by sellers or buyers for a rating system to make it truly useful? Is this level higher, lower, or the same with wallet based systems or coin based systems than for web based systems? This question is especially interesting with a coin based system because of the peer-to-per nature of the Bitcoin protocol. Bitcoin is dependent on pure computing power to guaranty the

integrity of its blockchain. Bad actors that may attempt to spend coins they do not possess or otherwise corrupt transactions are frustrated by the enormous computational capacity of the more honest part of the network. Altcoins, including a hypothetical transaction rating coin, will almost certainly have smaller networks behind them and so may be less stable. The problem might be somewhat mitigated, however, by the act of reading against the Bitcoin blockchain. A better understanding of these factors would be helpful in building new features and services for the Bitcoin ecosystem.

It has been shown in this discussion that grafting features requiring some degree of identity validation onto the Bitcoin protocol is a difficult task. It seems also to show that technology alone cannot drive a complete solution. The cooperation of sellers and buyers is also key. Changes to the way the system works, whether at the local or the global level, will require careful consideration of the incentives for all parties 1033(s)-9.0951(s)3./R58 0 0 cm BT 7(l)8551(f)]TE946Tjel ll