COTI: A Distributed Trust and Mediation System for Online Payments

November 5, 2017

Abstract

A payment system for managing and pricing trade risk between online pseudonymous parties would unlock trade opportunities that currently fail to manifest due to the untrusting nature of electronic commerce. We propose a distributed trust metric derived from an automated analysis of an open transaction network's activity to provide a predictive behavior-based risk assessment that appropriately prices trade risk and self-learns. Accurate pricing of trade risk informs fee pricing to supply a sustainable amount of credit for instant settlement and the funding of dispute arbitration by the network. To avoid a trust-dependent arbitration scheme, we present a decentralized mediation system that relies on a staking coordination game among incommunicable profit-seeking mediators. A trust metric further serves to inform individual transaction fee pricing so as to require higher risk transactors to fund the costs of absorbing losses as a result of negative payment outcomes. This incentivizes honest behavior - rewarding low-risk users with little to no fee requirements - and minimizes overall transaction costs over time.

1. Introduction

The root problem for online commerce is how to create a relationship of trust between disparate transacting parties. Localized social enforcement of trust is insufficient in global e-commerce to account for the risks involved between parties who have never met and may never interact again. Given these challenges, often no trade happens, or costly outcomes for one or both parties results. The information gap in assessing transaction risk implies an online reputation system is needed to establish trust.

Blockchain on-chain solutions are insufficient alone to handle the matter of trust, as they fail to capture integral payment data that occur outside their networks without the help of oracles. Blockchain systems are only designed to cover half the transaction process, focusing only on fund transfers, but lack the visibility to ensure the delivery of goods or services. Furthermore, digital currencies like Bitcoin are not natively equipped to deal with protections and services that consumers and merchants have grown accustomed to with traditional financial networks, including credit, instant payments [1], and a mechanism for handling fraud or ensuring customer satisfaction. Centralized online reputation systems do exist, but are limited to individual marketplaces and do not scale across networks. Moreover, the established credit bureaus have proven poor custodians of identity data [2], and frequently demonstrate misevaluations of creditworthiness en masse.

To address these issues, this paper proposes a distributed trust and mediation system that efficiently manages and prices trade risk. The establishment of a fundamentally-derived trust indicator enables a sustainable credit ecosystem with highly sought-after payment features. By appropriately assigning and pricing risk, a plethora of untapped trade potential can be unlocked for the benefit of all parties.

2. Trust Scoring

Traceable transactions allow a network to build a fundamentals-based trust reputation scoring for honest transactors that informs buyers and sellers about the reliability and fluidity they can expect from the transaction experience. We propose an autonomous machine learning-based trust scoring system to measure the risk of a transactor to the network in order to price and collect fees to sustainably fund a reserve fund that credits instant settlements and covers arbitration expenses. Machine learning based trust scoring is used for varied measures by COTI to inform pricing but the main feature discussed here is the transaction trust score.

2.1. Transaction Trust Score

The predicted transaction trust score ($\widehat{y}_i$) is a value between 0 and 1 that represents the probability of a successful transaction ($P(ST) \equiv \widehat{y}_i$), or can be defined by a user not losing a mediation ($P(LM) = P(ST)'$). A trust score of 1 implies that a user will have a 100% chance of a successful transaction, meaning no mediation occurs or the user wins the case. The score is based on a set of features (x_i) that describe the user's activity, which include the user's transaction volume-adjusted network centrality, mediation history, and satisfaction ratings.

The network centrality will likely have a strong influence on the overall trust score given the topological nature of one's position relative to their peers and mimics the properties behind Google's well-known PageRank algorithm. To derive the feature of network centrality for constant use under computationally efficient circumstances, we propose a method of polling the network graph to accumulate a time series representation of each user's centrality history as follows:

```
def run(): {
    // Load network graph
    val g = graphDB.loadGraph()
    // Run for 5 iterations
    val g2 = g.eigenVectorCentrality.run(g, 5)
    // Loop through each graph vertex
    g2.vertices.collect.foreach { nodeRank =>
        // Set the timestamp
        val ts: Long = System.currentTimeMillis
        // Get the user id
        val query = MongoDBObject("id" -> nodeRank._1)
        // Get the rank and timestamp
        val rankAndTs = MongoDBObject("rank" -> nodeRank._2, "ts" -> ts)
        // Append timestamp and rank to user
        mongoDB.update(query, $push("nodeRank" -> rankAndTs), upsert = true)
    }
    run()
}
```

Fig. 1: Code describing network graph polling for user centrality scores. Quantized ranking of users in the graph is determined using a recursive accumulation process using Scala and MongoDB.

This produces a feature appearing below as JSON:

```
{
    "userId_Abc123": [
        {
            "rank": 0.1,
            "timeStamp": 1394064000000
        },
        {
            "rank": 0.2,
            "timeStamp": 1225411200000
        }
    ],
    "userId_Zyx987": [
        {
            "rank": 0.3,
            "timeStamp": 1394064000000
        },
        {
            "rank": 0.5,
            "timeStamp": 1225411200000
        }
    ]
}
```

See Appendix A for further details into how the volume-adjusted network centrality feature is calculated and propagates over time.

It should be noted that new feature additions are possible under the governance rules of the network, and the ongoing set of features is likely to grow to include more analytic and demographic data. For testing or use with a model that requires normalization, and to more easily interpret the data, features are scaled according to the following formula:

$$x_{norm} = \frac{x - min(x)}{max(x) - min(x)}. \tag{1}$$

To improve computational efficiency, the historic max and min are cached and a given feature is only renormalized when its latest value being interpreted is greater or less than the historical max and min. Each normalized feature on each update is used to train the model to set the feature weight parameters (θ_i), using a logistic loss function on every iteration:

$$L(\theta) = \sum_i [y_i + \ln(1 + e^{-y_i}) + (1 - y_i)\ln(1 + e^{-y_i})], \tag{2}$$

where y_i is the output with which the hypothesis is compared. The model's objective function can handle overfitting once more features are being considered by use of a regularization function added to the loss function to decay each parameter's influence.

As the initialization machine learning models, a boosted tree algorithm along with Bayesian optimization are proposed to set the parameters and hyperparameters respectively due to their scalability [3], success [4], and transparency. To account for historically-expected model error, the model evaluates a daily cross validation sampling test against a wide enough rolling window of data to determine the recent mean accuracy percentage of the model, which is used to discount the score's estimated uncertainty. The model's accuracy (defined as its success probability) is multiplied by the mean accuracy of the model's cross validation test (δ) to get a success probability that accounts for the accuracy discount ($P(ST)_{adj} = P(ST) \cdot \delta$). A forward-looking method to account for the unseen future accuracy would be an ideal improvement for a later upgrade.

The output distribution of scores is expected to follow a bimodal distribution, limiting the influence of the mean, due to the nature of individual behavior with respect to credit and a wealth of empirical analysis [5]. The bimodality of credit outcomes is part of the inspiration for developing a score and pricing mechanism that will alleviate low-credit-risk individuals from socializing the cost of high-risk transactors and allows high-risk transactors to participate fairly.

3. Network Fee Structure

The network has a transparent fee structure with a majority of fees allotted to provide network-enhancing services. The fee pricing required for a transaction includes fees that are paid by the buyer and seller; however, the buyer fees are included in the price of the product, while the seller fees are paid directly to the network. The seller determines the highest-risk buyer (lowest trust score) that they are willing to transact with, and the ticket price displayed reflects that:

$$\text{Display Price} = (\text{Price of Product}) + (\text{Lowest Allowed Trust Score Buyer Fees}). \tag{3}$$

A discount is presented to the buyer if their trust score is higher than the lowest allowed score,

$$\text{Discount } = (\text{Lowest Allowed Trust Score Buyer Fees}) - (\text{Actual Buyer Fees}). \tag{4}$$

Transactions between buyers and sellers include an overall transaction fee that is made up of four individual fees: reserve credit fund fee, mediator fee, market maker fee, and COTI network fee.

3.1. Reserve Credit Fund (RCF) Fee

The RCF fee (ϕ_r) is only implemented when the buyer or the seller opts for instant access to unsettled funds (discussed in more detail in Section 5). Sellers will use this, for example, to access a buyer's cash instantly, even while a product or service fulfillment is still underway. A buyer takes advantage of the instant settlement when they wish to move funds instantly that are otherwise slow to settle, like paying with Bitcoin or a bank wire.

The fee capitalizes the RCF to ensure there is enough money to pay the buyer for certain mediation outcomes where the seller is at fault, while also providing a backstop for the risk of default when credit is provided to the buyer trying to pay with a slow-to-settle method. The RCF fee is dependent on the trust score of the user and is determined by first assessing the zero-odds expected value of the credit event (a loss in funds to the rolling reserve) occurring. Given the two-state outcome, the binomial expected value of a successful transaction is:

$$\text{Expected Value} = P(S)_{adj} \cdot (\phi_r) + P(CE \mid LM)_{adj} \cdot (-\tau), \tag{5}$$

where τ is the transaction size, the worst-case negative cost, and both $P(S)_{adj}$, the probability of a successful transaction, and $P(CE \mid LM)_{adj}$, the probability of a credit event given losing mediation, are adjusted for the expected model error rate. With a zero-state expected value we can then determine the fee by the following formula:

$$\phi_r = \frac{P(CE \mid LM)_{adj} \cdot \tau}{P(S)_{adj}}. \tag{6}$$

3.2. Mediator Fee

Since mediation cases require a reward incentive, the mediation fee (ϕ_m) is designed to collect a sufficient amount over time to entice the required number of mediators to clear the caseload (mediation is discussed in more detail in Section 6). The fee is paid by both the buyer and the seller according to their respective trust scores, and it relies on the transaction size (τ) and $f_{transaction}$, a fractional value determined by a preset arbitrary commission to serve as the estimated fair initial expected fee for incentivizing arbitration.

$$\phi_{m,buyer} = P(ST)^{'}_{adj,buyer} \cdot (\tau \cdot f_{transaction}), \tag{7}$$

$$\phi_{m,seller} = P(ST)^{'}_{adj,seller} \cdot (\tau \cdot f_{transaction}). \tag{8}$$

Mediation cases require a minimum number of mediator participants to justify the consensus process, and given that the quantity of mediators willing to participate is dynamically variable, a market mechanism is required to ensure a sufficient supply relative to the caseload height. The adjustment mechanism is important to ensure a minimum predictable rate of output, assuring the network of its ability to meet its minimum mediation requirements. This system is similar to how Bitcoin targets one block every ten minutes using a difficulty adjustment mechanism to influence mining profitability, or how Uber uses surcharging to fill gaps in supply and demand. In the case of mediation, the target is for the network to have at least a sufficient supply of mediators relative to the case demand queue.

To help trend toward this equilibrium, a surcharge is applied depending on the forward curve prediction of the supply and demand for mediators. The surcharge is dependent on the ratio of the predicted number of cases (N_c) to the predicted number of mediators available (N_m), matching the expected mediation period start time to the forward curve's prediction. Since one mediator can handle multiple cases per time slot, we find the expected supply of available mediators by multiplying by the three-month simple moving average[1] of cases handled per mediator (μ), and normalize this value product by the minimum number of mediators required for a single case:

$$N_{m_{adj}} = \frac{N_m \cdot \mu}{N_{m,min}}. \tag{9}$$

Therefore, the surcharge is calculated by multiplying $1 - \frac{N_{m_{adj}}}{N_c}$ by a fixed-cap surcharge range (e.g. 0 to 5%) as a percentage of the transaction size, and is only applied when $N_{m_{adj}} < N_c$. When $N_{m_{adj}} > N_c$, a discount is applied based on a fixed discount range, which will reduce the default $f_{transaction}$ share of τ.

The predicted number of mediators and mediations is derived from a time series event forecasting algorithm. The model is trained on a set of features that includes but is not limited to categorical features such as mediator rewards, the number of mediation cases currently open, time to settle mediation, and number of mediators available during the mediation initiation period.

3.3. Market Maker Fee

The market maker fee (ϕ_n) depends on the currencies used in the transaction; the easier the foreign exchange routing, the fewer fees are required. The market makers' general duties include supporting exchange rates on an ongoing basis to minimize all conversion transaction costs (see Section 7 for more details on the market maker).

Buyers and sellers can both choose their preferred payment and receipt currencies, but the configuration required to settle on a per-trade basis will determine the fee required. The way to avoid any fees is to transact in the system's native currency, COTI (XCT). If it is accepted by the seller and offered by the buyer, no market maker fee is required. As the market maker is responsible for converting the relevant transaction fees to XCT, the overall fee is lower when part of the transaction is already in XCT, since less needs to be converted. In case of an exchange requirement due to a currency mismatch, the system's market maker will source or supply foreign exchange (FX) liquidity as efficiently as possible to settle the trade, and flow through its cost in fees to do so [2].

[1] A more optimized method here may be required, including possibly a machine learning-based approach.

[2] The FX exchange cost will be very low compared to the payment industry standard for relatively liquid pairs, as the market maker is tapped into near-bank-rate liquidity sources.

3.4. COTI Network Fee

A fixed percentage of the transaction, lower than the standard acquirer assessment domestic fee [6], is charged per transaction to fund the foundational network administration.

4. Cryptographic Escrow

An escrow system can fully protect buyers and sellers when transacting with a counterparty they don't know or trust. In this system, the money is only released once the buyer acknowledges that the delivered good or service is as promised. Unless the seller has a high enough trust score to justify opting into a credit pay system where the payment occurs instantly (explained further in Section 5), an escrow-based transaction flow based largely on a scheme designed by Satoshi Nakamoto will be offered [7].

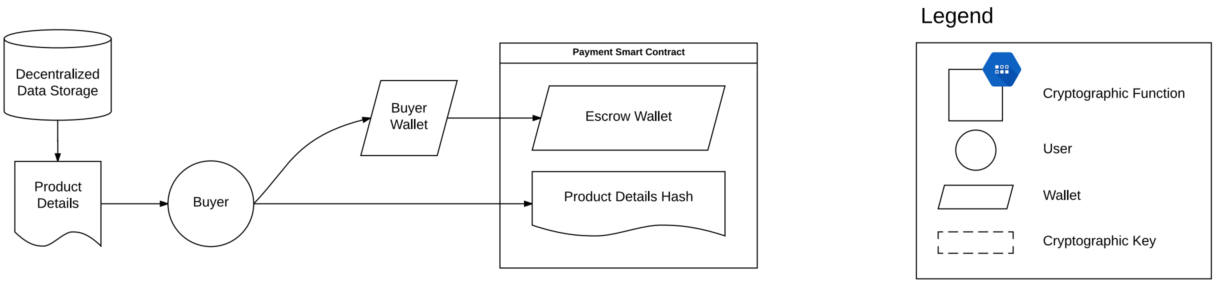

Fig. 2: Create a Payment. The buyer creates a payment by sending coins and product details reference to the payment smart contract.

To initiate an escrow purchase of a good or service, the system leverages Bitcoin's native smart contracting tools, with the buyer sending the purchase amount to a multisignature wallet that can only move money if at least two private key signatures are provided. The buyer and seller each retain one key and can leverage this over the other to incite honest behavior and encourage resolving the matter without mediation (see Figure 3).

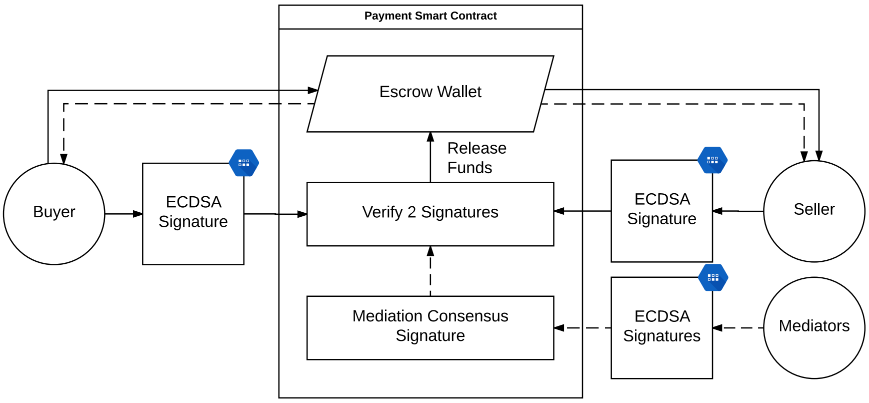

Fig. 3: Escrow System. Funds sit in an escrow wallet within the payment smart contract. The smart contract acts as a multisignature wallet. The solid arrows point in the direction of capital flow and represent a transaction with the escrow system that does not require mediation; both the buyer and the seller agreed to the transaction. The dotted arrows also point in the direction of capital flow, but they represent a transaction that requires mediation; thus, the capital can either be returned to the buyer or continue to the seller (see Figure 2 for legend).

Should either party believe funds are indefinitely locked without merit, mediation can be initiated by either party. In the event of mediation consensus, a third key is generated and automatically exercised in favor of the winner (explained further in Section 6). Contracts are built with timelocks based on a fair duration after the imputed delivery times, and given no settlement or mediation within the preset timeframe, the contract will expire and permanently send the funds to the seller.

5. Reserve Credit Fund (RCF)

The overarching goal of payment industry players is to avoid any delays in settlement. Not all transactions require an escrow safety mechanism because of the nature of the transaction or the trust between buyer and seller, and when possible, these transactions should settle immediately. The challenge to such an ideal is that most settlement networks that do not rely on credit have significant delays between funds being sent and received.

Supplying credit, however, requires an appropriate understanding of each user's creditworthiness, which is already a key feature of the COTI network - the trust score. The COTI network leverages these known trust scores to supply credit for instant payments for sellers, and to backstop any losses with a network-maintained RCF.

Traditionally, a rolling reserve is required of merchants, providing a margin to account for possible chargebacks by the merchant's clients. RCF in COTI's case makes possible instant payments, as the reserve buffer exists to deal with cases where such instantaneous resolution has proved imprudent or unfortunate (see Figure 4).

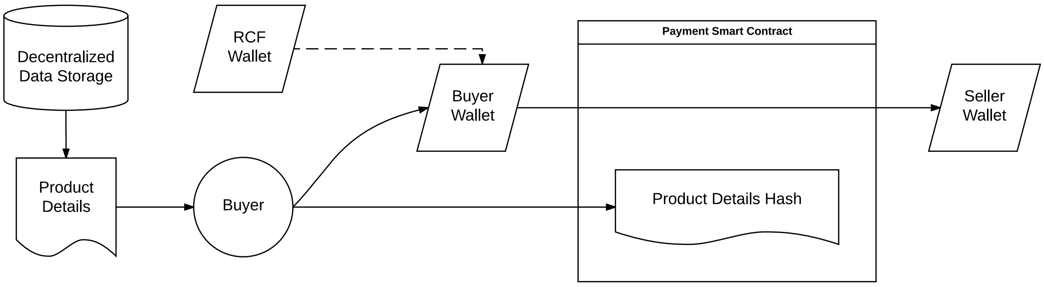

Fig. 4: Create Payment with Instant Settlement. When a seller opts in for instant settlement, funds are transferred from the buyer to the seller without sitting in an escrow wallet. The dotted line represents a buyer winning a mediation case, where they receive funds directly from the RCF (see Figure 2 for legend).

6. Mediation

Mediators act as data oracles, as they are incentivized to honestly validate real-world information in disputes, and their consensus triggers smart contracts that enforce the efficient and fair flow of commerce.

The proposed source of mediation is a decentralized distributed system that reaches consensus on disputes using crowdsourced mediators. Just as Bitcoin miners receive Bitcoin for their efforts, COTI's mediators are rewarded with COTI's native currency; however, COTI rewards are not paid out of inflation, but rather in part from the premine, and in the long run from ongoing fees.

The network's mediators are a special class of COTI holders who choose to engage in mediation by running a validator node and bonding whatever stake of their COTI balance they decide to wager in a proof-of-stake like fashion.

6.1. Mediation Process

A transaction occurring on the COTI network requires both parties to sign off on a sales agreement at the start, which governs the terms of the transaction including mediation. The signed agreement and details of the transaction are encrypted, hashed, and stored on a public blockchain[3]. The platform also provides encrypted communication features[4] with the ability to upload new information that is appended to the case for consideration by possible mediation.

[3] There are numerous storage options. The base protocol as of this time is agnostic to which is chosen, but a decentralized, highly redundant option is ideal.

[4] A connectable solution like Bitmessage may be used to handle secure communication. A basic model for this was introduced by a now-abandoned project known as Orisi [8].

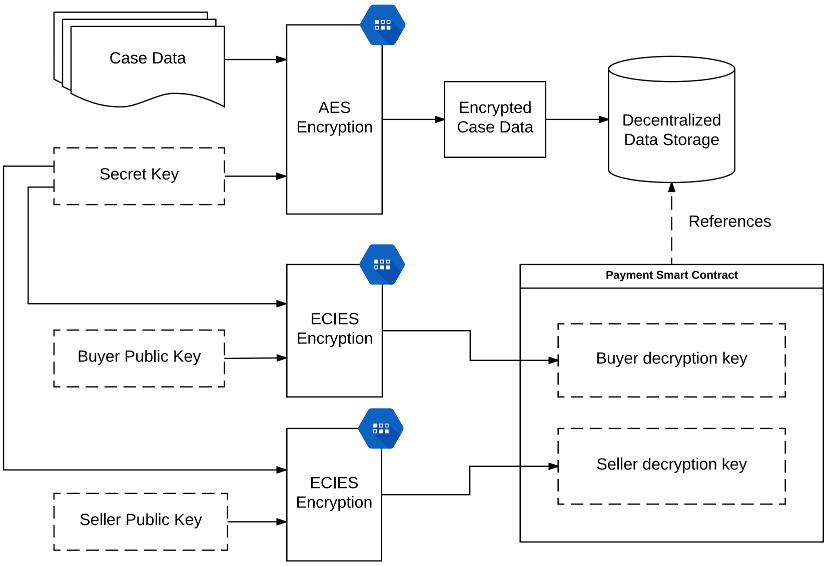

Fig. 5: Encrypting Purchase Data. When the buyer or the seller adds case data, they use a randomly generated secret key to encrypt the new case data (using AES encryption), as well as generate an Elliptic Curve Integrated Encryption Scheme (ECIES) encrypted version of the secret key, which can only be decrypted by the other party (see Figure 2 for legend).

In the case of the escrow system, when either party believes their funds have been locked without merit, or in the case of instant settlement where the buyer is not satisfied with the goods or services, they can broadcast their case to a queue managed by the COTI network's smart contract worker. The worker then reads the queue and randomly distributes[5] caseloads to active qualifying mediator nodes[6], granting them temporary access keys to broadcast votes with.

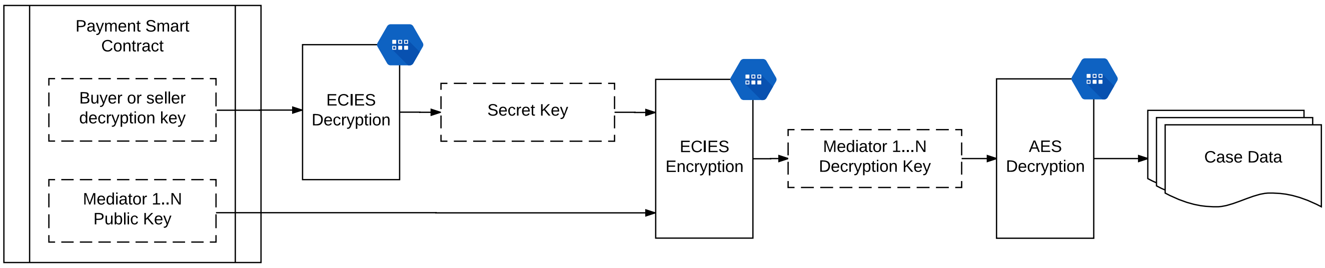

Fig. 6: Decrypting Purchase Data. Buyers or sellers decrypt the case data secret key using their Integrated Encryption Scheme (IES) private key, then re-encrypt the secret key using the mediator IES public keys so the mediators can access it using their IES private keys (see Figure 2 for legend).

6.2. Voting Rules

The mediator nodes must review the case within a prespecified time and sign a sealed vote in favor of the buyer (vote = 0), seller (vote = 1), or neither (vote = 0.5), implying more information is required (see Figure 7 and 8). Votes provided by mediators are assessed in a vote matrix with rows of mediators and columns of case votes. A preset quorum of mediators, comprising an 80% supermajority, is required to exercise the mediation key. Different M of N key combinations will be offered, and these quorum quantities are blind to voters. A mediator that does not submit his vote within the prespecified time will lose their bond and after each instance will exponentially increase the minimum bond required to participate.

[5] Specialization branching for mediator selection is a future optimization worth considering to improve settlement efficacy and efficiency.

[6] Mediators may not qualify for a case due to scoring or connection conflicts, such as previous history with the subjects involved. A path distance requirement may be enforced to handle further degree connection conflicts.

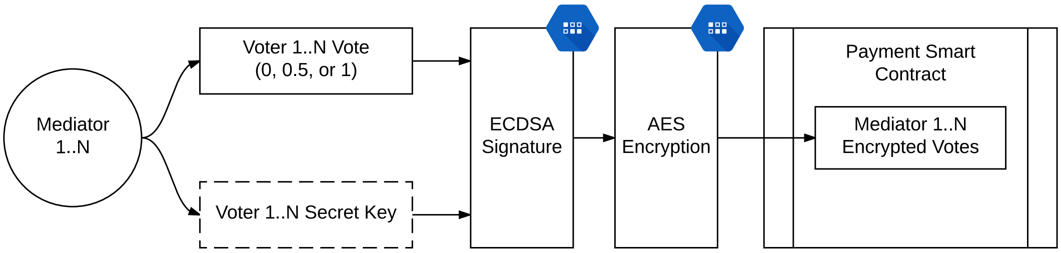

Fig. 7: Mediation Vote Casting. Each mediator signs their vote (using a ECDSA Signature), encrypts their vote with a randomly generated secret key (using AES encryption), and uploads the encrypted vote to the payment smart contract (see Figure 2 for legend).

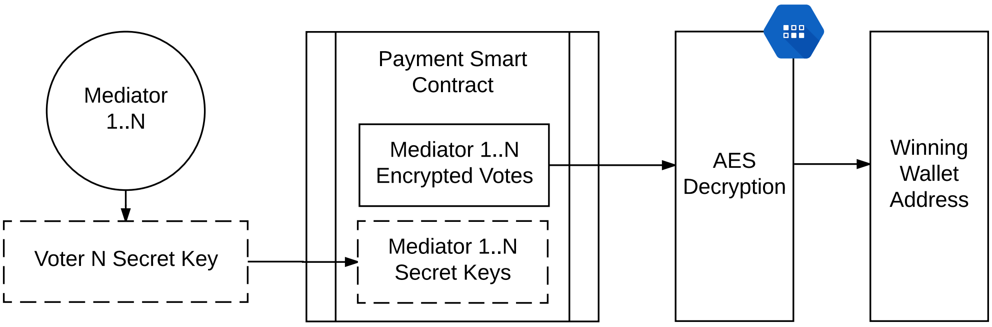

Fig. 8: Revealing Mediation Winner. Once N voters have voted, each voter publicly reveals their secret key. All votes are now able to be decrypted and verified. The winning party is determined and the funds are released (see Figure 2 for legend).

6.3. Consensus Mechanism

A focal point (Schelling point) coordination game is played amongst mediators in order to find consensus of who won or lost a dispute. The mediators must reach similar conclusions without coordinating, and in doing so are likely to search for focal points of which the singular truth serves as a glaringly-accessible example. To incentivize the search for truth in disputes, the network makes available a portion of the transaction fees, encouraging mediators to vote on cases and reduce the caseload queue.

While participating in mediation, a mediator can retain or increase their overall stake based on conformist and participatory actions, or stand to lose at least part of their bond from the opposite behavior. The blinded decision making process will have an added benefit of incentivizing users to make larger bets, giving them more skin in the game, and incentivizes the completion of participating in the case once agreeing to begin. Minimum bonds will be enforced to ensure some meaningful cost to those who fail to vote.

To start reviewing a case, before casting votes, mediators are required to issue blind bets relative to each other that estimate their own ability to ultimately be part of a successful quorum of mediators. Quorum details and information on the other mediators are also blind to each participant.

	Bond	Vote	Incentive Distribution %
Mediator 1	100 XCT	-	71.43%
Mediator 2	10 XCT	-	7.14%
Mediator 3	10 XCT	-	7.14%
Mediator 4	10 XCT	-	7.14%
Mediator 5	10 XCT	-	7.14%

Assume a four of five quorum is needed, and after the vote, Mediator 5 is outside this quorum, with the rest reaching consensus. Mediator 5 loses his role in this mediation and relinquishes his stake to be evenly distributed amongst the other mediators as shown in the table below.

	Bond	Vote	Incentive Ownership %
Mediator 1	107.69 XCT	1	76.92%
Mediator 2	10.77 XCT	1	7.69%
Mediator 3	10.77 XCT	1	7.69%
Mediator 4	10.77 XCT	1	7.69%
Mediator 5	~~10 XCT~~	0	0%

The percent share at this point of an individual mediator's stake relative to the pool of all participating mediator stakes will dictate the relative transaction fee reward distributions to each mediator. The amount awarded to mediators for a successfully mediated case (λ) is proportional to the transaction value (τ) of the disputed case:

$$\lambda = (\tau \cdot f_{transaction}). \tag{10}$$

The amount will be distributed among the participating mediators proportional to their relative bonds and will have a monetary cap of a predetermined value.

6.4. Collusion Prevention

An assumption of the Schelling point-based game is that users cannot coordinate, and therefore, collusion attempts have to be thwarted to maintain the system's integrity. The design is setup to impose costs on active coordination such that the cartel startup and operating cost is greater than the effort required to discover and report the truth. A key starting measure of collusion prevention is the anonymous distribution of case data. Mediators are granted temporary keys and are never provided information on other participants in the vote. In addition, mediators are disqualified from reviewing cases that involve the same transacting parties or the same mediation partners more than once within a sufficient timeframe.

6.5. Privacy

Identifying user data can be partially hidden or obscured at the point of collection, but to best equip mediators with pertinent case data, certain details of the transaction ought to be shared. Different approaches to privacy may be discovered and implemented.

6.6. Hung Jury

In the event that a supermajority is not reached, the second round of mediation is performed. If a supermajority is reached during this round, the incentive and bond are distributed between the mediators who voted correctly in both rounds. If a supermajority is not reached after the second round, the case is opened to a network supervisor for review.

In the example below, the mediators can't reach a decision after the first round.

Round 1	Bond	Vote	Incentive Ownership %
Mediator 1	10 XCT	1	-
Mediator 2	10 XCT	1	-
Mediator 3	10 XCT	1	-
Mediator 4	10 XCT	0	-
Mediator 5	10 XCT	0	-

In the second round, there is an 80% majority in favor of the buyer.

Round 2	Bond	Vote	Incentive Ownership %
Mediator 6	10 XCT	1	-
Mediator 7	10 XCT	0	-
Mediator 8	10 XCT	0	-
Mediator 9	10 XCT	0	-
Mediator 10	10 XCT	0	-

The incentive and bond is distributed between the mediators who voted correctly in both rounds.

Round 1	Bond	Vote	Incentive Ownership %
Mediator 1	~~10 XCT~~	1	0%
Mediator 2	~~10 XCT~~	1	0%
Mediator 3	~~10 XCT~~	1	0%
Mediator 4	16.7 XCT	0	16.7%
Mediator 5	16.7 XCT	0	16.7%

Round 2	Bond	Vote	Incentive Ownership %
Mediator 6	~~10 XCT~~	1	0%
Mediator 7	16.7 XCT	0	16.7%
Mediator 8	16.7 XCT	0	16.7%
Mediator 9	16.7 XCT	0	16.7%
Mediator 10	16.7 XCT	0	16.7%

6.7. No Winner

In the event that mediators determine by consensus a vote for neither party (a supermajority with a vote of 0.5), the case will move to the second round of voting. If the consensus remains, rewards will be distributed according to the normal redistribution and incentive allocation rules. The case determination will move to the network administrators after the second round to resolve with finality, and the decision will not impact the mediator rewards.

6.8. Mediation Settlement

Resolved mediation results in the activation of a key to unlock funds, and will return to the buyer or seller depending on the system used. When an escrow is used, the funds from the escrow will be sent to the winner of the mediation (see Figure 9). When instant credit is used, in case the buyer wins mediation, the RCF will send the money back to the buyer; however, the seller winning results in no change of funds (see Figure 10).

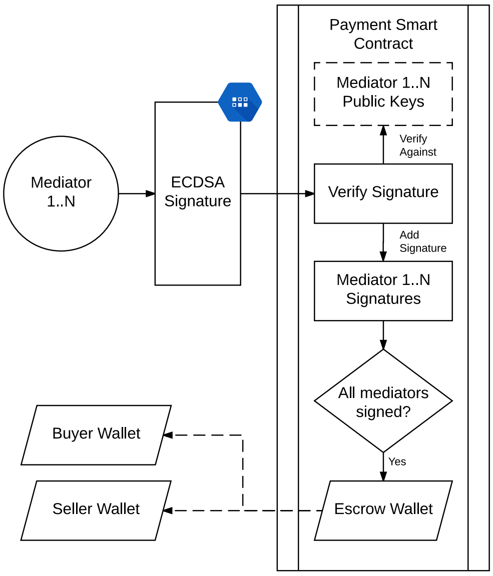

Fig. 9: Escrow Mediation Settlement. Each mediator adds their ECDSA signature to the payment smart contract. The payment smart contract can verify signatures based on the mediator IES public keys which exist in the contract. When mediators have met consensus, the escrow wallet are released to the winning side using smart contract scripting (see Figure 2 for legend).

6.9. Mediation Partial Settlement

The responsibility is placed on the mediation initiator to state the gap in service or goods rendered, and mediators will be required to validate this amount. The initiator may need to be warned that misrepresenting the amount mediated upon can lead to a lost mediation.

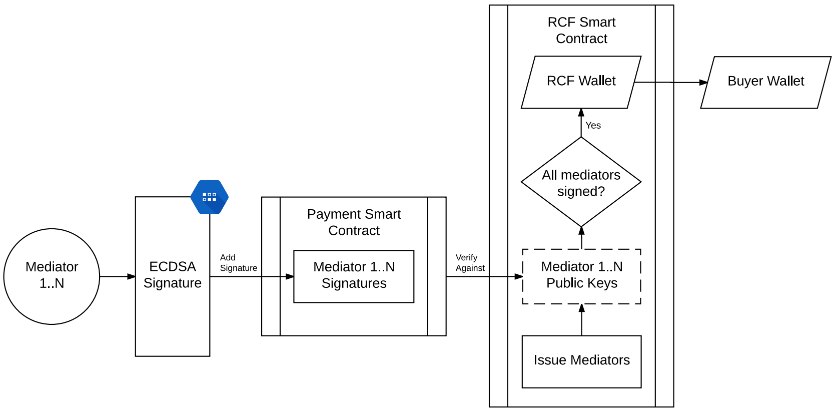

Fig. 10: Credit Mediation Settlement. When the seller opts in for instant payment, an escrow wallet is not used. Instead, payments go directly from the buyer to the seller. If the mediator's consensus is on the buyer side, they will receive funds directly from the RCF. The RCF wallet exists within the same smart contract that issues mediators (see Figure 2 for legend).

7. Making Efficient Markets

7.1. COTI Routing

A challenge of transacting across borders is the foreign exchange facilitation often required given a diversity of national currencies. A Canadian attempting to buy goods from India, for example, where both buyer and seller prefer their native currencies, causes issues because there is no easily available liquidity for all foreign exchange (FX) pairs. Only three currency pairs contain the majority of all FX trade volume globally, according to the Bank for International Settlements [9], as seen in the table below, leaving many FX pairs like CAD/INR without an easily accessible liquid direct exchange relationship.

Pair	Trade Percentage
USD/EUR	23.1%
USD/JPY	17.8%
USD/GBP	9.3%

To help route transactions in the COTI network, such that any combination pair can efficiently settle, COTI's network will make markets with XCT as the common denominator:

XCT / INR
XCT / USD
XCT / EUR
XCT / JPY
XCT / GBP

XCT / BTC
XCT / LTC
XCT / XRP
XCT / ERH
XCT / n...

In this way, there is always a path between any supported currency via XCT, making it an important means of exchange. An automated market making process is utilized to seed and supply liquidity across markets (described in 7.3).

7.2. FX Hedging

To support a system where buyers and sellers both can independently select the currencies they wish to pay and receive with, and under the premise that payment finality is not instantly determined, the RCF, buyer, and seller all need to defend against unwanted FX movements.

In the case of an escrow transaction, given an asymmetry in preferred payment and receipt currencies, the buyer's currency is held in escrow as an asset to the network, while the system expects a liability to the seller. Assume the buyer wants to pay with Canadian dollars (CAD) and the seller wishes to receive Indian rupees (INR); to manage the risk, the system market maker will automatically issue a short CAD/INR, effectively going long INR and short CAD to offset the liability to the seller and the retained asset of the buyer.

When the seller opts for credit, the liability from the system moves to the buyer in case the RCF needs to refund the buyer. The market maker must then hedge the system's risk relative to the RCF's most exposed holding, while still optimizing against the current state of the various FX markets.

7.3. Automated Market Making

A key network function is the automated market making system, which is responsible for converting appropriate fees to the COTI currency and supporting COTI's role as a routing mechanism for FX settlement. The success of the market maker depends on its ability to achieve the goals of increasing liquidity and narrowing spreads, thereby reducing potential slippage costs across all markets priced in COTI.

For information and trade execution access, the COTI market maker will be connected by the network administrators to an increasing number of liquidity sources, and will consider a range of execution opportunities when helping seed the order books for COTI vs. other assets. Third-party sources will be encouraged to connect in without the help of the network administrators; therefore, over time, new competing sources can emerge and the reliance on the network administrators can be removed.

7.3.1. Model

The market maker must be equipped with a model that is able to make decisions even given the absence of any market information, and upon receiving new data, make appropriate updates to account for what has been learned. To achieve this end, this paper proposes a zero-profit Bayesian market maker (BMM), improved on from Das and Magdon-Ismail's 2008 version [11], as the mechanism for pricing orders on crosses with limited liquidity, helping to improve pricing execution for all users in the system.

Das and Magdon-Ismail provide approximations for how to achieve a BMM by optimizing the marginal updating of the market maker's beliefs with incoming trader data. These beliefs are then used to derive pricing for a zero expected gain system and sustainably capitalized BMM.

The market maker is responsible for providing a price quote p_t for Q units. To account for trade sizes, the market maker offers prices based on a volume weighted average price (VWAP). Beliefs are updated successively based on trade action, as trades are either accepted or rejected, and the market maker attempts to be self-aware of its own validity enabling it to adapt to high magnitude belief updates. The acceptance and rejection of trades can happen on market exchanges supplied by COTI or by a user switching currencies to execute a product or service purchase with which the FX fees pricing proved unmoving.

7.3.2. LMSR comparison

The Bayesian approach is generally preferred largely for scalability and computational efficiency over market scoring rules, such as Robin Hanson's logarithmic market scoring rule (LMSR), which suffers computationally under large combinatorial circumstances like the infinite state continuous double auction markets the COTI network will be required to trade in. Moreover, the BMM in practice displays better convergence around the equilibrium price with the LMSR, conversely, demonstrating volatile behavior even when the true price remains stable [12].

7.3.3. Initialization

Since the market maker at times may have to make bets blind, where trader beliefs are unknown, an initial distribution must be assumed. Typically a Gaussian distribution is adopted and evolves away from this state after the learning process of the model adapts its expectations. Given that markets do not follow normal distributions, and rather, contain more data in the tail ends of a distribution than the Gaussian approach predicts, we prefer to use a better representative distribution to initialize with. A cumulative distribution function of the generalized extreme value (GEV) distribution is chosen to serve this role given its common use in finance for these properties [13].

The value of the market as the initial state is $U : GEV(\mu, \sigma, \xi)$ where $\mu \in \mathbb{R}$ is the location parameter and $\sigma > 0$ represents the scale, and $\xi \in \mathbb{R}$ is the shape parameter.

7.3.4. Pricing

The market maker uses its initial belief to then set prices in favor of the goal of zero-profit. To calculate bids and asks the market maker uses the GEV mean and variance. The spot price is the mean price M_t,

$$\begin{cases} \mu\sigma(g_1 - 1)/\xi & \text{if } \xi \neq 0,\ \ \xi < 1, \\ \mu + \sigma\gamma & \text{if } \xi = 0, \\ \infty & \text{if } \xi \geq 1 \end{cases} \tag{11}$$

where $g_k = \Gamma(1-k\xi)$, and γ is Euler's constant. The variance V_t,

$$\begin{cases} \sigma^2(g_2 - g_1^2)/\xi^2 & \text{if } \xi \neq 0,\ \xi < \frac{1}{2}, \\ \sigma^2 \frac{\pi^2}{6} & \text{if } \xi = 0, \\ \infty & \text{if } \xi \geq \frac{1}{2} \end{cases} \tag{12}$$

measures the uncertainty of the information set. If a finite mean or variance is incalculable the market maker reverts to a Gaussian belief system $U : N(\mu_t, \sigma_t^2)$ which is maintained in tandem and used in practice as a backup in case of a GEV noncomputable output.

We utilize the heuristic introduced by Brahma *et al.* [12] to account for share quantities where independent orders are of a fixed size α. The initial state starts at $M_1 = M_t, V_1 = V_t$ with an arbitrary expectation of k fragmented orders ordered sequentially. These orders are hypothetically processed one at time with the market maker quoting an ask for a given buy (the same is true for a sell in reverse with a bid order given alternatively)

$$\text{ask} = M_i + V_i Q(p_i)\sqrt{1+p_i^2} \tag{13}$$

and updates its belief at each iteration

$$M_{i+1} = M_i + V_i \frac{B}{A} \tag{14}$$

$$V_{i+1} = V_i(1 - \frac{AC + B^2}{A^2}) \tag{15}$$

for the full set of orders where A, B, and C are functions of the mean, standard deviation, learning updates and price as explained in the Guassian context by (Das and Magdon-Ismail 2008). The VWAP of this procedure is used to quote a price for Q units in practice,

$$\text{ask} = p(Q) = \frac{1}{Q}\sum_{i=1}^{k} \alpha_i p_i \tag{16}$$

To detect and adapt to paradigm shifts we must update the framework to handle a moving window of observations. We use the belief updates z^+ and z^- which are learned from trade action in addition to the market state to produce a probability that the sequence of trades will appear in the trade window of size W,

$$L(\mu, \sigma) = \int_{-\infty}^{\infty} GEV(v, \mu, \sigma, \xi) \prod_{i=1}^{a} \left(\Phi(z_i^+, v, \sigma_\epsilon) - \Phi(z_i^-, v, \sigma_\epsilon)\right) dv \tag{17}$$

where v is the value of the traded item. The consistency history is then defined as

$$C(\text{history}) = L(M_t, 2V_t) - L(M_t, V_t) \tag{18}$$

to increase the uncertainty in a way that is not sensitive to the choice of W or the trade sequence [12].

7.3.5. Risk Management

While the BMM has many of the appropriate features the COTI network's automated market maker requires, an unfortunate gap however is the possibility of bankruptcy, which can only be avoided by employing a risk management strategy. Other models can offer an ingrained guarantee against intolerable losses like the LMSR and its well-known b parameter, which sets the maximum loss bound the system is willing to tolerate. A dynamic pari-mutuel market making approach could entirely prevent risk to the network as well despite some initial subsidization being required in practice to kick-start liquidity, but requires a wide range of heuristics to infer unpredictable market participation.

The COTI BMM takes a statistical approach to managing risk by leveraging its beliefs to determine appropriate bet sizing and risk tolerance given its capital level. The aim is to minimize its risk of ruin while optimizing its positive impact.

7.3.6. Funding

The market maker will be partially funded by the premine (described in Section 7) and a percentage of the fees (described in Section 3) for non-COTI to COTI transactions.

8. Scaling Blockchain Payments

Trust scoring and the RCF combined with system-specific risk mitigation tactics can provide the ability to scale slow-to-settle payment systems, both traditional and blockchain alike. Blockchain currencies like Bitcoin and Ethereum face particular challenges given significant on-chain scaling limits that cause congestion and delays to sending funds across their networks. Throughput is a major issue for these cryptocurrencies since they boast at best double-digit transactions per second (TPS) at peak performance [14], while Visa and Mastercard alone, for example, commonly demand tens of thousands of TPS.

As a payment processor, COTI's network can provide credit facilities to blockchain transactions by monitoring and participating in a public blockchain's network dynamics to mitigate risk and sustainably provide off-chain scalability. A prime example in a blockchain context is demonstrable with current Bitcoin features enhanced by COTI's credit assessment and issuance services. The aim of such a service is to allow a confident acceptance of zero-confirmation transactions, which implies near-instant payments scalably for Bitcoin users.

Smart contract figures and explanations are trivialized in this paper to increase clarity. Today's smart contract platforms would be too costly to use in the manners described. Data stored within smart contracts will have to be minimized as much as possible by reducing storage to what is only absolutely necessary. Much of the data could be stored as references to off-chain data stores such as IPFS, Swarm or Filecoin. A P2P messaging protocol such as Whisper or a custom solution could be used to increase scalability by reducing communication overhead and blockchain access between buyers, sellers, and mediators.

The method for minimizing the network's risk requires the maintenance of a well-connected node that broadcasts transactions it credits across its network, while listening for a double-spend. If discovered, the COTI node will notify its peers of the invalid transaction as fast as possible. The propagation of the original, valid transaction will likely outweigh the influence of any double-spend, drastically minimizing the odds of the original transaction being improperly credited. This happens because the non-linear progression of nodes that become aware of the valid transaction will most likely quickly outpace nodes recognizing the double-spend.

The history of these credits can be used as a feature to help a network-run machine learning model predict the likelihood of a double-spend, and can therefore be used to charge a fee appropriate to the risk while still offering instant Bitcoin credit. The above scheme is largely based on an early proposal by Bitcoin founder Satoshi Nakamoto, who estimates the costs of such credit to be significantly lower than traditional credit card loss from fraud. The RCF can be used to fund any losses that may arise, and will maintain calibration with the risk involved by actuarially pricing fees. With transaction malleability resolved since the activation of segregated witness on Bitcoin, the risk of being fooled at the zero-confirmation state is further diminished.

8.1. Sidechain scaling

Various proposals are being developed to help alleviate mainchain congestion and provide scalable, secure state updating. Proposals like the lightning network [15] and extension blocks [16] are key examples of how these goals can be achieved. To the extent COTI is deployed as part of the Bitcoin blockchain, COTI plans to leverage parts of these designs, specifically channel state updating similar to the lightning network to enable large off-chain transactions to occur at high throughput while maintain on-chain enforceability.

The real power of COTI however is with the credit based model which provides buyers and sellers with credit alike, and therefore the slow-to-settle base infrastructure becomes significantly less limiting so long as COTI has confidence in eventual settlement.

9. Token Mechanism

9.1. Initial Coin Distribution

All COTI coins are premined, and the fixed money supply can never be tampered with. Coins will be distributed according to a public auction to ensure starting coin holders are those who value the coin's future potential the most. A premine is required to support the initial funding of system features, such as the RCF and the market maker, which are designed to self-sustain on fees thereafter.

9.2. Coin Value

All transaction fees, exempt of the market maker[7], are converted into COTI and support the market value of COTI relative to those currencies. The long-run value of the coin can then be determined by the following formula:

$$\text{COTI Market Capitalization} = (\text{Transaction Fees}) \cdot (\text{Transaction Volume}) \cdot (\text{Net Present Value Multiplier}). \quad (19)$$

[7] The market maker may require holding funds in alternate currencies, and is not necessarily mandated to convert fees.

9.3. Governance

The network requires initial parameters to be set, and some are arbitrarily assigned by the system designers. These parameters include but are not limited to the chosen:

- Machine learning models and assumptions[8]
- Vote time periods
- Mediation quorum minimums
- Reliance on network administrators to resolve indeterminable cases
- Mediator qualifications

The network administrators also play a crucial role in the upstart and management of the system, but aim to increase the reliance on decentralized governance over time. A target implementation of the COTI network would be on the Bitcoin network via a sidechain, using Paul Sztorc's drivechain concept [10] as a governance model for connecting to the mainchain. The sidechain's parameters would then be set as a balance between the COTI clients and holders and the Bitcoin miners. At the genesis of the sidechain, a migration would take place from the existing COTI network, where COTI holders would be given a fixed-priced conversion into the sidechain, which would have an equal amount of COTI created. The support mechanisms behind the old network will then expire, and the network will continue with the Bitcoin blockchain at its core security and data integrity foundation.

10. Conclusion

In this paper, we proposed a means of conducting trade without trust and extending traditional blockchain attributes to include advanced payment features. We outlined a sustainable network-based payment feature set, including instant credit, risk-based pricing, decentralized mediation, and network-based market making to ensure a smooth and vibrant payment ecosystem. The ecosystem maintains incentives such that behaviorally-driven high-risk users bear their fair share of the costs, supporting the ability for low-risk users to transact with unusually low fees.

References

1. "Average Confirmation Time - Blockchain," https://blockchain.info/charts/avg-confirmation-time, 2017.
2. R. Miller, "Equifax data leak could involve 143 million consumers," https://techcrunch.com/2017/09/07/equifax-data-leak-could-involve-143-million-consumers, September 2017.
3. T. Chen, C. Guestrin, "XGBoost: A Scalable Tree Boosting System," https://arxiv.org/abs/1603.02754, 2016.
4. "XGBoost - Machine Learning Challenge Winning Solutions," https://github.com/dmlc/xgboost/tree/master/demo#machine-learning-challenge-winning-solutions, 2017.
5. S. Hlawatsch, S. Ostrowski, "Simulation and Estimation of Loss Given Default," FEMM Working Papers, 2010.
6. "Visa Canada Standard Acquiring Network Assessment Fees," https://www.visa.ca/dam/VCOM/regional/na/canada/Support/Documents/visa-canada-standard-acquiring-network-assessment-fees-feb2017.pdf, 2017.
7. "A proposal for a semi-automated Escrow mechanism," http://satoshi.nakamotoinstitute.org/posts/bitcointalk/325/#selection-21.3-21.431, August 2010.
8. "Orisi White Paper," https://github.com/orisi/wiki/wiki/Orisi-White-Paper, November 2014.
9. "Bank for International Settlements," https://www.bis.org.
10. "Drivechain: Enabling Bitcoin Sidechains," http://www.drivechain.info/, 2017.
11. S. Das, "The Effects of Market-Making on Price Dynamics," https://pdfs.semanticscholar.org/8b2a/2bea33ac3b563f1e3a1680ef595f6ef92b72.pdf.
12. A. Brahma, M. Chakraborty, S. Das, A. Lavoie, M. Magdon-Ismail, "A Bayesian Market Maker," https://www.cse.wustl.edu/ mithunchakraborty/papers/predMarkets.pdf.
13. D. Guégan, B.K. Hassani, B.K "A mathematical resurgence of risk management: an extreme modeling of expert opinions," Frontiers in Finance and Economics, 2014.

[8] Assumption bias may be mitigated by use of genetic algorithms.

14. "Youtube: Vlad Zamfir of Ethereum talks Scalability, Extraordinary Claims, PoS, & Responsibility - Part 1 of 5," https://www.youtube.com/watch?v=bj_Qpvx-GPM

15. "The Bitcoin Lightning Network: Scalable Off-Chain Instant Payments," https://lightning.network/lightning-network-paper.pdf, January 2016

16. "[Bitcoin-development] soft-fork block size increase (extension blocks)," https://www.mail-archive.com/bitcoin-development@lists.sourceforge.net/msg08005.html, June 2015

Appendix A

Eigenvector centrality measures the influence of nodes in a network by considering the number of incoming connections. In a node-edge graph that represents users as the nodes and transactions as the edges (weighted by the cumulative size of the transactions), the graph displaying the centrality values of the buyers has the arrow pointing toward the buyer, while the graph showing the centrality values of the sellers has the arrow pointing toward the seller. The seller centrality score describes the influence of a user as a seller, while the buyer centrality score describes the influence of a user as a buyer.

The seller centrality score increases if a high-scoring seller buys from that user. The buyer centrality score increases if a high-scoring buyer sells to that user. Buying from a high-scoring seller or selling to a high-scoring buyer does not increase a user's centrality score; for example, Amazon would be an example of a user with a high-scoring seller centrality score. Purchasing from Amazon does not make a user a more trustworthy buyer, as many users buy from from Amazon. But if Amazon bought from a user, it would increase that user's credibility as a seller, and subsequently increase their seller centrality score.

An increase in a user's centrality score, regardless of if it is their buyer or seller centrality score, correlates with an increase in the user's trust score. The amount that the score is influenced by the centrality score is determined by the machine learning algorithm that will take into account a variety of other user qualities, such as their satisfaction rating and mediation history.

The following example shows how a user's centrality scores propagate over a series of transactions. User 4 joins the network with three other users that have already transacted with each other. For clarity of the example, after Round 0, the only transactions that occur are with User 4; however, typically transactions will be occurring between all users all the time.

Round 0: User 3 and User 1 buy from User 2, User 2 buys from User 3

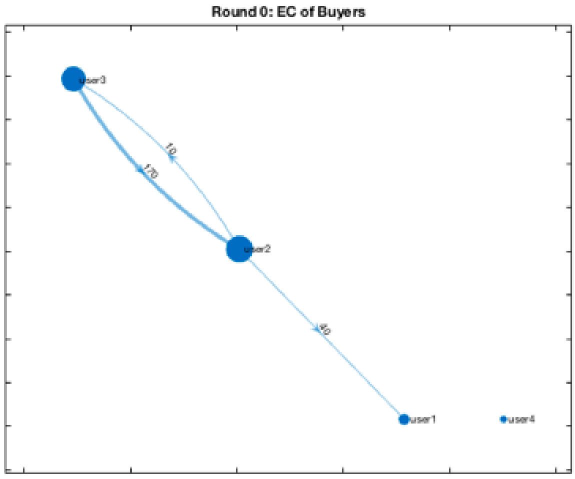

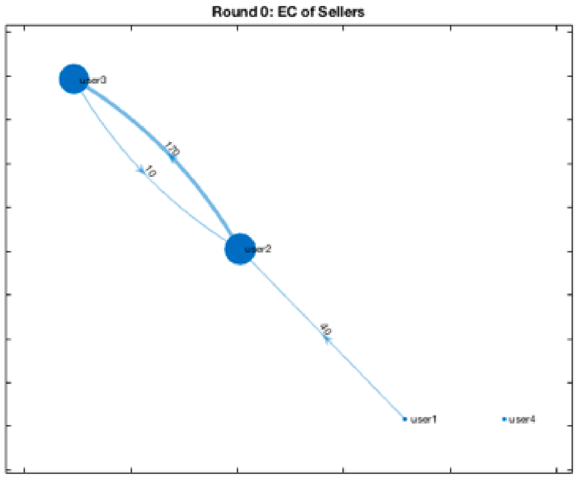

Initially, User 3 and User 1 have bought from User 2, and User 2 has bought from User 3. The centrality scores of the buyers and the sellers are as follows:

User	Buyer Centrality Score	Seller Centrality Score
1	0.1531	0.0476
2	0.3963	0.4634
3	0.3617	0.4414
4	0.0889	0.0476

Due to the number of users in this example, User 4's initial centrality scores before transacting with any other users is not equal to 0. As more users are added to the network, this initial centrality score approaches 0.

User 1 has the same seller centrality score as User 4, since User 1 did not sell during this round. The seller score for Users 2 and 3 is similar - even though User 3 sold more than User 2, User 2 sold to more users, and thus their centrality values come out to be comparable. Even though User 1 bought slightly more than User 3 ($40 compared to $10), User 3's centrality score is higher than User 1 because User 2 bought from User 3, increasing the importance and subsequently the trust score of User 3 compared to User 1.

Round 1: User 4 buys from User 2, User 1 buys from User 4

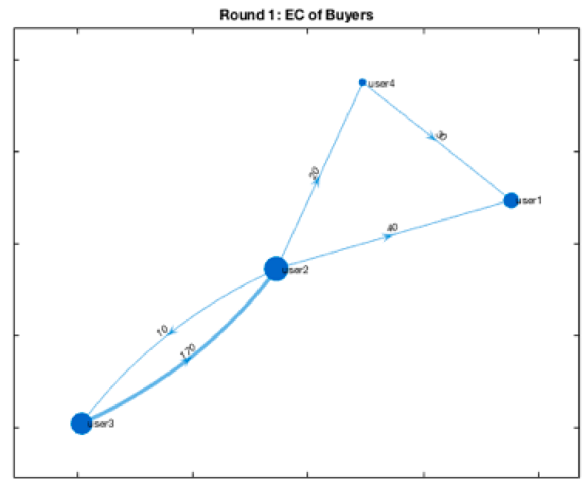

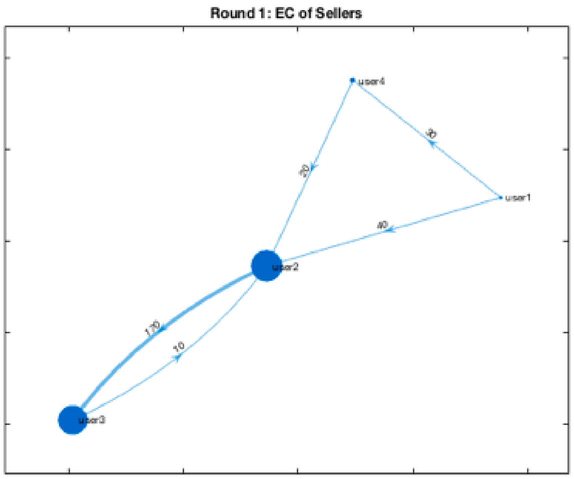

In round 1, User 4 buys from User 2 ($30), while User 1 buys from User 4 ($20). Note that the edge of the weight is the additive cost of the transaction over time. For example if User 4 buys a second good or service for $60 from User 2, the weighting of that edge would increase to $90. The resulting centrality scores are as follows:

User	Buyer Centrality Score	Seller Centrality Score
1	0.2245	0.0375
2	0.3569	0.4658
3	0.3196	0.4334
4	0.0990	0.0633

Compared to the initial round, User 4's buyer centrality score stays relatively constant, while the other three users' scores increase slightly. However, User 4's seller centrality score increases while User 1's decreases, and the other two users remain relatively unchanged.

Round 2: User 2 buys from User 4

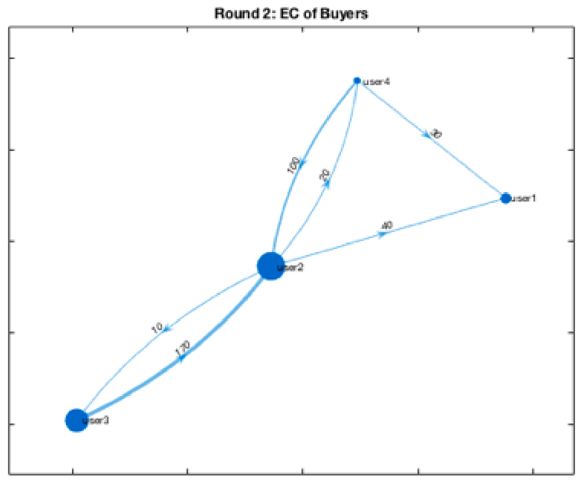

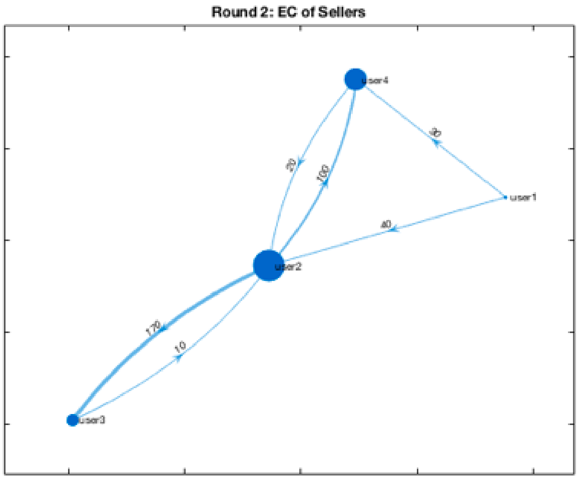

In the second round of transactions, User 2 bought from User 4 ($100). The centrality scores are as follows: User 4's buyer centrality decreases slightly, but their seller transaction score increases significantly, as User 2, a high-scoring seller, bought from User 4.

User	Buyer Centrality Score	Seller Centrality Score
1	0.1511	0.0375
2	0.4186	0.4658
3	0.3445	0.1695
4	0.0858	0.3272

Round 3: User 4 buys from User 3

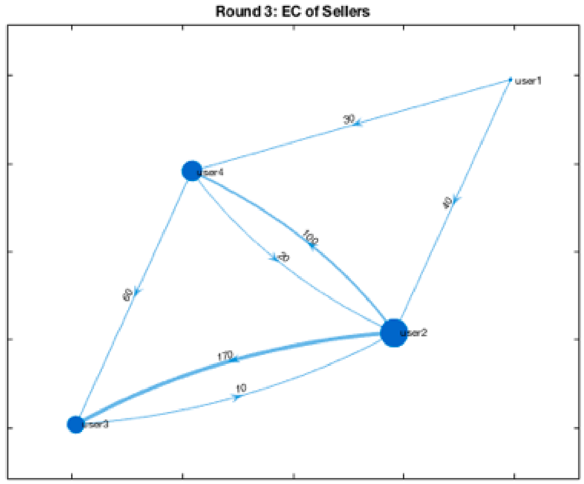

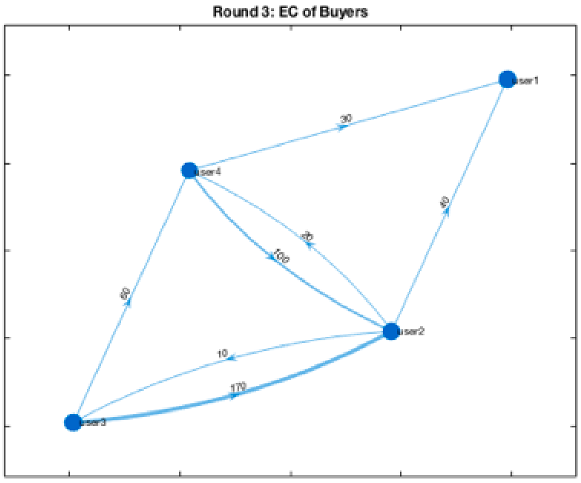

In round 3, User 4 bought from User 3 ($60). The centrality scores after this transaction are as follows:

User	Buyer Centrality Score	Seller Centrality Score
1	0.2547	0.0375
2	0.2538	0.4144
3	0.2584	0.2500
4	0.2332	0.2981

User 3's seller score increases, while the other users' scores remain relatively unchanged. However, all of the users' buyer centrality scores become relatively equal after this transaction, with a large increase in User 4's score.