TradeLayer: Peer-to-Peer Electronic Cash Backed By Decentralized Derivative Exchange

Patrick Dugan
desk@tradelayer.org
www.tradelayer.org

Abstract. Bitcoin trading has depended on custodial exchanges' corporate bank accounts, on-chain transaction volumes are dominated by batched withdrawals between exchanges, un-even fee structures pollute data-authenticity, and stable units of account suffer from reflexive instability or risk of regulatory capture. The Lightning Network was proposed in 2015 as a way to desegregate transaction data-footprint from the Bitcoin blockchain's capacity, it hinges on the guarantees of holding an unpublished transaction signed by two parties, related to spending an input from a 2-of-2 multisignature address, this has use for trade execution. The OmniLayer protocol was developed from the first ICO's proceeds in 2013, attempting to realize a vision of decentralized exchange on the Bitcoin blockchain, playing host to Tether USD which has issued many billions of dollars, yet has never completed an audit. The TradeLayer protocol completes the vision for the Bitcoin family of protocols by extending OmniLayer to support multisig-channel-based trading, and decentralized derivatives that clear in chains of bilateral participants. The two provide a fully collateralized, hedged basis for the issuance of audit-able, bearer currency redeemable for cryptocurrency. TradeLayer enables the economies of Bitcoin-like blockchains to sustain decentralized currency and trading systems independent from external price information and bank accounts.

1. Introduction

Commerce using Bitcoin has come to rely almost exclusively on financial institutions serving as trusted third parties to process cryptocurrency withdrawals. While the system works well enough for most transactions, it still suffers from the inherent weaknesses of the trust-based model.[1]

In order to make Bitcoin and its adjuvants a more perfect money, we must solve the key problems of decentralized exchange latency, make custody easier for individuals to manage, and create liquid markets in on-chain derivatives to enable native currency in users' chosen denominations.

There are four threads of research and development that came together to enable the protocol alternative presented in this paper: decentralized derivatives work started in 2014 by this author and extended by this author's colleagues in 2017/18, general potential designs for

OmniLayer [2] style transactions embedded into OP_Return payloads on Bitcoin-like blockchains, State Channels, and the Lightning Network

The introduction of leverage enables the issuance of currency. Currency historically is not power money, it is a derivative of money, a B-side, a receipt. Since the loss of the gold standard, the US dollar has been the primary currency used to purchase oil. In the 21st century, redeemable currency backed by various kinds of cryptocurrency will compete with bank deposits. These decentralized currency units can be used as collateral for on-chain peer-to-peer derivatives trades. It eliminates dependence on issued bank account receipts, such as Tether USD.

Decentralized currency also complements regulated currency tokens, as those generally pay little or no yield but carry regulatory guarantees the market may find attractive. As regulation tightens, the regime of regulated bank-deposit-backed tokens being able to transfer anonymously without respect to the Travel Rule will end, and issuers will require features like TradeLayer's modular identity parameters, enabling traders and issuers to whitelist potential counter-parties. As Bitcoin is a Sybil-prone, pseudonymous protocol, users can self-certify their exemption from regulation where they reside as a default option, taking on potential liability if that isn't true. The decentralized currency units do not require an identity assertion to transfer or trade for BTC/LTC but does require at least a self-certification to create/redeem, as one becomes a derivative counterparty at such time.

The problem with using a Bitcoin-like blockchain for settlements of trades is the matching of those trades becomes fraught with the fuzzy nature of time in the Bitcoin protocol. The only true timestamp is the Block Height, hence the original OmniLayer token exchange mechanism, relying on mempool timestamps, would always be fraught with miners re-ordering those fuzzy timestamps. Additionally, to the world of high frequency trading, waiting anywhere from 2 to 200 minutes for a Bitcoin block to maybe confirm your limit order, in maybe the priority you wanted, is a non-starter.

To solve this problem, while still having native decentralized exchange on the Bitcoin protocol, we use 2-of-2 multisignature addresses to create Trade Channels. Two parties can handshake a new Trade Channel in milliseconds, if they are comfortable taking credit risk, or wait for several block confirmations of a transaction committing the collateral. Then the two parties can co-sign trades together with microsecond latencies, gaining a strong assurance of trade finality. The parties can subsequently publish the transaction immediately, or hold it and defer settlement to save on fees, periodically updating the trade with a new Block Height expiration value.

The sum of these functions is a decentralized exchange environment capable of:

1) pairing the home chain's native currency with any token on the layer (the native metacoin, any hedged dCurrency, issued securities, regulated stablecoins)

2) pairing any two tokens on the layer
3) trading decentralized derivatives that settle based on the above on-chain trade data
4) fast execution using native multisig

The derivatives clearing algorithm allows for leveraged trading and path-dependent bilateral settlement of contracts; trade in a channel, but have the fallback of a global clearing system that reconciles accounts every four hours. Several mechanisms reduce the risk of insolvent liquidation of leveraged positions and the risk that such liquidations affect the wider system. The apex of this liquidity is the creation and redemption of transferable units representing fully-funded hedged positions.

2. Transactions

Transaction inputs consist of one or more user-provided inputs, the outputs consist of: 1) a change address, 2) reference address with dust amount (optional to tx where external address is the target), 3) an OP_Return. OP_Return payloads are base58 encoded strings that, when parsed, segment into the following fields: 2 bytes -> protocol marker; 2 bytes -> tx type; n bytes -> parameter1 through parameterN. The OP_Return size allocation varies between Bitcoin-like blockchains. Large and rare transaction types can be encoded into concatenated series of tx emissions that all reach the mempool in sequential order, such as may be necessary to broadcast proofs. OP_Return's have continuity, and are parsed to arrive at a consensus hash about the state of the layer's balances, orderbooks, lists and other data structures, which enables consensus checks between nodes.

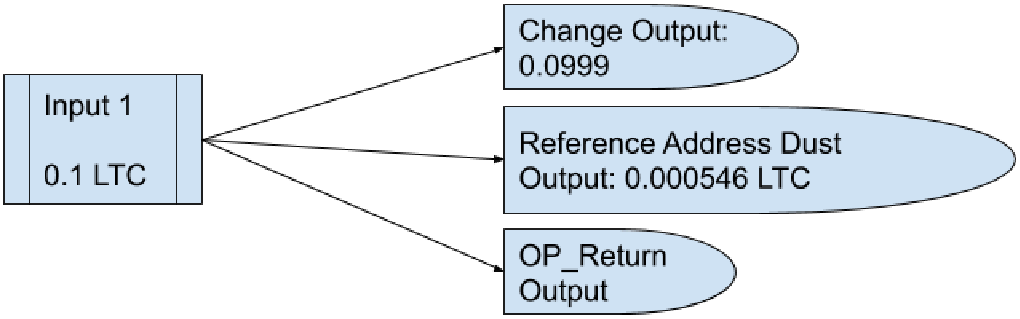

An input is needed to publish the transaction and produce a dust output, the dust threshold is a low number of units that must be spent to put a transaction on the Bitcoin blockchain. A 3rd party service provider could fund the miner fee component of the tx. Earlier dust inputs could be recycled for new dust outputs if fee difference is covered externally. If not then the user will want a change address to keep the surplus amounts from the inputs. Having published a dust output, the user is allowed one OP_Return output to encode a layer transaction payload. If sending or granting tokens, a tx may also color a reference address with another output, or for a simple send, the dust output serves this purpose.

3. Trade Channels

There are two mechanisms in the Bitcoin protocol that create a sense of truth: the proof-of-work Block Height and signature verification based on the mathematical asymmetry of public-private key cryptography. To deal with the trading bottleneck of the prior, we lean on the latter, using the power of 2-of-2 multisignature addresses to create co-signed consent.

Unlike depositing BTC to a 2-of-2 address, as in the Lightning Network, the OP_Return based token balances can be Committed while “on a string” and pulled dynamically with a Withdrawal transaction. The logic embedded into the protocol reflects the nonce-based system designed by Dr. Christian Decker for BIP 118[3], if the block height encoded into the two-signature trade transaction has not yet arrived, then the Instant Trade transaction will be valid if confirmed prior-to or by that block height. This allows parties to periodically sign a new transaction to defer on publishing, but also allows either party to publish the transaction with a calculated anticipation of expiration, or as a counter-move to an attempted Withdrawal that may be an attempt to welch on the trade. A window of 7 block-confirmations is used to complete a channel withdrawal, in the Bitcoin version of TradeLayer, and a window of 28 blocks in Litecoin.

Transfers of pledged collateral between Trade Channels can be done via a transfer transaction, though it must be jointly signed so the transfer is consensual. This allows a pool of capital already committed to Trade Channels to be moved around efficiently, making a dealer’s inventory model much easier when managing many Trade Channels. As with a Commit, a Transfer transaction with a distant time horizon can be accepted the way a bank accepts a cheque, or published to effectuate the Transfer, accepting a Transfer tx as collateral allows for capital to move around on very low latencies.

A PNL Settlement transaction can allow for two parties to batch a series of smaller, unpublished transactions into a single transaction and close a channel only having published four transactions on-chain (the initial commits, the settlement, and the total withdrawal of any one party). Trade Channels that have not published any transactions within 24 hours worth of blocks will expire, returning committed collateral to both counterparties.

4. Counterparty Value Adjustment

The trustworthiness of a counterparty isn’t a requisite with multisig trades, however it does make trading more efficient. Consider a transaction that is signed and has an expiration 3 blocks away. When one party signs it, the other party essentially holds an option to not co-sign the trade. This can be abused, so the history of a dealer with a known set of addresses and channels completing trades quickly helps build a reputation and makes the perceived cost of trading with that dealer lower. The value of the partially-signed transaction can be quantified using a methodology called Counterparty Value Adjustment [4].

After the 2008 Global Financial Crisis (GFC) and the chancellor’s second bailout for banks, quantitative analysts working for financial corporations had to invent a methodology for trading derivatives with other financial corporations while assuming that another GFC could happen at any time and threaten cascading defaults. What they realized is that an under-collateralized transaction must have an adjusted price similar to the value of an option. While the trade might be settled quickly, minimizing time-value, it must be accounted for in the pricing of the trade or the other party is being given a free option to maybe deliver, maybe not, depending on if price movements are favorable. In addition to the time-value, the extent of a CVA price adjustment would be sensitive to market volatility and the credit-risk of the counterparty. A formula came to be adopted combining elements of the Black Scholes model with the default-risk estimation used in pricing Credit Default Swaps.

Counterparty Value Adjustment also applies to Atomic Swaps between Bitcoin and Litecoin, the time-until-expiration of the timelock channel used for that trade determines the time value. In TradeLayer Trade Channels, time value is determined by the Block Height expiration of a signed-unpublished transaction. The historical efficacy of the counterparty to co-sign trades at low latencies can lower the perceived default risk.

While decentralized exchange has no inherent overheads other the ongoing operation of its users and the host blockchain, and thus can have fees approaching zero, centralized exchanges provide guarantees that a high-fidelity dealer in a Trade Channel environment can only approach. That guarantee’s quantifiable value can account for a sustainable difference in fees between these venues.

Compounding the cost of CVA optionality in multisig trading, is the uncertainty around block times and mempool backlogs making confirmation on Bitcoin more expensive and uncertain in a shockingly convex manner. Students of financial history know that shocking convexity is the arch-nemesis of financial systems. Because proof-of-work uses a deterministic Difficulty parameter to calibrate a fundamentally stochastic mining algorithm, you never know when blocks will show up; could be two minutes or two hundred. Hence the Block Height used for a Trade Channel trade must be at least 3, so you have time to react in case they come in fast.

When the mempool is full, fees can jump from a few satoshis per byte to hundreds, a brutal convexity, akin to the value of bribes from wealthy people trying to board full lifeboats on the Titanic. To be safe, Trade Channel trades should be constructed as replace-by-fee transactions. Even then the costs of trading in a decentralized environment such as TradeLayer will become expensive in times of blockchain congestion and/or high price volatility, giving more centralized liquidity venues such as sidechains and hosted exchanges a long-term moat. Having a Litecoin version helps in the medium-term. High frequency arbitrageurs who provide vital liquidity to panicking and liquidating contract holders in the occasionally fast price moves of >10% in <15 minutes, may pull their bids or reserve their Trade Channel liquidity to only last-look quotes.

In the global foreign currency markets, last-look quoting (where the dealer reserves the right to decide to reject the order on low-latencies) has increased liquidity and tightened spreads, so there is hope this could follow for Bitcoin derivatives and other assets on TradeLayer. The low-frequency institutional or retail customer of a high-fidelity dealer would then, in these mempool-backlogged, high-volatility contexts, depend on the dealer to not burn their own reputation by treating the order flow as a free option and instead hedge off on another venue and widen their spreads. In practice there is wiggle room on sub-100ms latencies for algorithms to reason on microsecond latencies about what to do with that order flow while still maintaining the appearance of high-fidelity and a good reputation. This could become a new horizon for quantitative trading.

5. UXTO Trading

A major limitation of protocols like TradeLayer and OmniLayer is lack of directly control over raw BTC or LTC UXTOs. The original OmniLayer protocol implementation of a decentralized exchange for Bitcoin to tokens, such as Tether, suffered from a three-step process and serious delivery-failure risks. In TradeLayer we have relaxed the logic used by OmniLayer and assume that buyers will broadcast their order-accept and order-deliver tx simultaneously and have them confirmed within a few blocks. Mempool backlogs make this process nonetheless trepidatious. Since BTC/LTC volumes are the most important volumes in the layer (vs. token-to-token and contracts), due to their liquidity being anchored to a global, pre-existing liquidity complex, this was a necessary problem to solve.

The solution is to package the delivery of the UXTO and the OP_Return documenting the trade and ensuring token delivery, into the same atomic transaction, as illustrated in the below diagram via an example of someone buying the TradeLayer for Litecoin native metacoin, ALL, with 1 LTC:

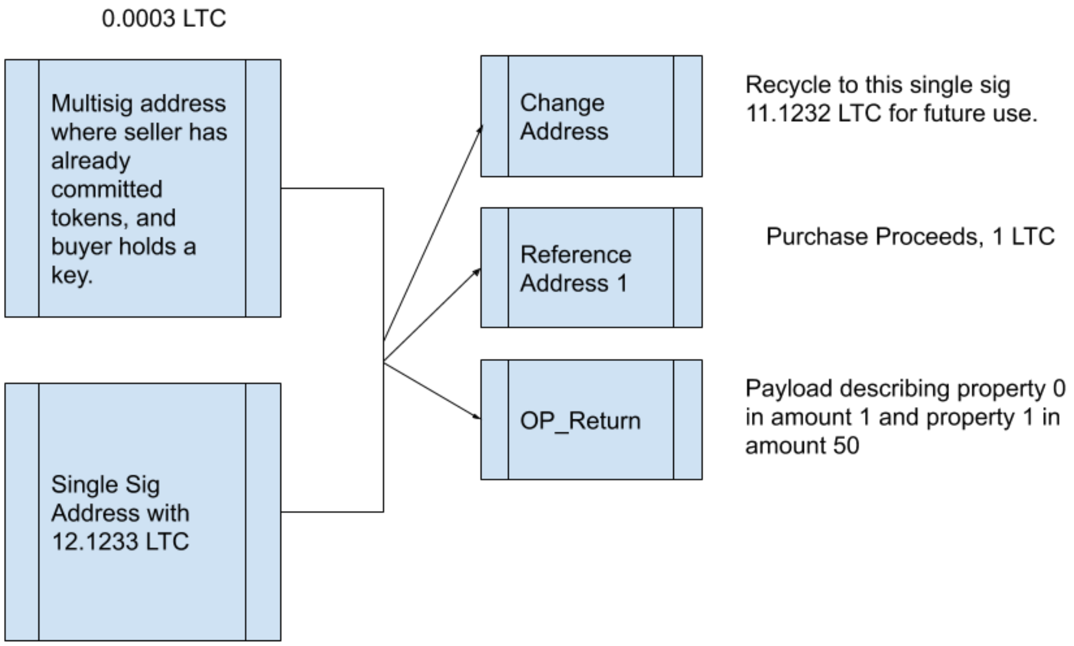
0.0003 LTC
Multisig address where seller has already committed tokens, and buyer holds a key.
Single Sig Address with 12.1233 LTC
Change Address
Reference Address 1
OP_Return
Recycle to this single sig 11.1232 LTC for future use.
Purchase Proceeds, 1 LTC
Payload describing property 0 in amount 1 and property 1 in amount 50

As with other Trade Channel transactions involving tokens or derivatives, the security of the trade hinges on a logic running outside the protocol, such as in the javascript code of a web-wallet, on only co-signing a trade which has all the correct inputs and outputs.

6. Incentives

The original Master Protocol/Mastercoin, re-branded to OmniLayer/OMNI, had a flat supply bestowed to the original investors, plus a 10% vesting component that was meant to fund development but ended up predominantly becoming the property of the insider who holds the private key.

Full-validating nodes in the Bitcoin-family protocols do not earn any revenue for their contribution to network resilience.

Centralized exchanges in unregulated jurisdiction have an incentive to wash trade at no true cost to themselves, to pump volume, and by correlation, revenue. The meta-coins ALL and TOTAL for the Litecoin and Bitcoin versions of TradeLayer respectively, exist as cousins to

OMNI, with dynamic emission patterns. The core use case for the meta-coin is to accrue trading fees as a deflationary force, to incentivize liquidity as an inflationary force, and to allow a float of stable currency based on derivatives of the metacoin. Inflation in the metacoin occurs in the following ways:

1) Node Reward: Node operators relay transactions periodically including a cipher about the latest block's resulting TradeLayer consensus hash, at periodic intervals based on the last characters in the new blockheader as it corresponds to the last characters in the user's address, such that a check-in occurs every twelve to twenty four hours, but in an unpredictable pattern. Batching many addresses to farm the Node Reward is possible but the miner fees involved creates some anti-Sybil protection. To protect newcomers with little capital far in the future from this value being crowded by mining oligarchs, there are no race conditions that determine who gets the reward, any valid cipher confirmed that block can earn a pro-rata share, though a dishonest miner could censor their transactions. The inflationary dilution of the Node Reward is very small so greed doesn't ruin what is meant to be a perpetual faucted for marginal newcomers.

2) Liquidity Reward: Market liquidity providers earn the meta-coin as a secondary rebate for having their orders matched by someone else. As there is no distinction between makers and takers in Trade Channels, the CVA/last-look aspect being priced into the quoted bid/ask spread, some Liquidity Reward also accrues to Trade Channel volume. This only applies to native contracts, native token or UXTO trades aren't subsidized so as to not dilute the anti-wash trading mathematics of the fees.

3) A set of addresses holding vesting tokens lock in a reserve balance of Founder Reward that unfreezes as the cumulative volume of the decentralized exchange increases. The volume that makes up the x-axis of this vesting function is, like with Liquidity Reward, denominated in BTC or LTC as these are the only truly native references for liquidity. Token and contract volumes are translated to their BTC or LTC equivalent by way of historical price triangulation. The Founder Reward begins to vest at 1000 LTC and 100 BTC, and maxes out at 10,000,000 LTC and 1,000,000 BTC, based on a proportionate logarithm (e.g. 50% vested at 100,000 LTC).

Trading will initially incur the following fees:

Token-to-Token: Maker, -0.03%; Taker, 0.05%

UXTO Native Coin-to-Token: 0.01%, Token Seller Only

On-chain Native Contracts: Maker, -0.005%; Taker, 0.01%

On-chain Oracle Settled Contracts: Maker, -0.01%; Taker, 0.02%

In Trade Channels:

Token: 0.01% Each Side

UXTO: 0.01%, Token Seller Only

Native Contracts: 0.0025% Each Side

Oracle Contracts: 0.005% Each Side

The taker fees and rebate ratios for the pairs that create settlement data for contracts are more sharp to create a concave risk profile for a would-be market manipulator [5] having to pay taker fees and contend with the liquidity of all the properly incentivized arbitrageurs picking up multiple basis points of rebate income. It's possible for someone quoting an offer to trade, to include in that offer a signal that they expect a rebate, or that they expect the other side to pay fees. Flags in indicators of interest can be used to construct different kinds of Trade Channel funding combinations, which could allow for one party to pay fees and/or rebate the other. This is accomplished with a Change Address belonging to the rebated counterparty, which can receive a BTC or LTC output, pre-funded by depositing the rebate + estimated fees to the Trade Channel.

As more traders hold the meta-coins, the effectiveness of the "house" potentially wash trading for its own aggrandizement becomes infeasible. The community becomes The House. The cartelization we've seen traditionally with large exchanges, and again with cryptocurrency exchanges, can be diluted by the increasingly broader base of holding in the beneficiary asset. However Wash Trading (WT) remains a pattern attack vector, essentially a Sybil Attack, to earn these rewards from less-than-bona-fide activity, the only to properly price a liquidity rebate is to make Wash Trading a way of buying the coin, and thus the inflation rate per unit of volume must converge on or below the net-cost of WT. Regulators can detect WT in on-chain environments, but only use it for prosecution if it is within their jurisdiction and the participants are identifiable, net-costs making WT just an inefficient way of going long ALL or TOTAL is a mathematical alternative.

7. Reclaiming UTXO Space

Bitcoin-protocol design practitioners have a civil duty to design for data-efficiency where possible to no incur a tragedy of the commons. We can't go to 0 dust thresholds because the potential spam of mostly friction-less OP_Return strings could fill up the backlog for practically no cost. We could go to lower dust thresholds and more importantly larger OP_Returns, such that multi-state referencing transactions such as Send-To-Many are doable with fewer 100 byte signatures. Dust outputs however, all require their own 100 byte signatures to clean up and recycle them into future transactions, crowding the UXTO set and heightening RAM requirements to run a client. Prioritizing that clean-up by inserting the dusts into unsuspecting wallet users tx's is bad user experience design. Independent Segregated Witness transactions, or under an alternative mechanism that helps reduce the cost of multi-input transactions, should probably be the industrial-best-practices way of periodically cleaning up large numbers of dusts with minimal cost. Soft-forks to increase OP_Return payload limits can create efficiencies vs. increasing the number of reference outputs in a transaction.

8. Wrapping UXTOs as Tokens

We can conceive of four kinds of collateral that would be useful on TradeLayer:

1) Metacoin and its dCurrency derivatives
2) Issued tokens backed by legally prescribed value (banked dollar coins, securities)
3) Wrapped BTC or LTC held custodially on an auditable set of multisignature addresses
4) BTC or LTC Wrapped non-custodially using new OP_Codes in the base protocol

Work in Bitcoin OP_Code design by Jeremy Rubin with OP_CheckTemplateVerify [6] poses an example of ways to add new building blocks to these transactions. A BTC UXTO could be pledged alongside a counterparty to an address bound by OP_CheckTemplateVerify to create a binary pay-out situation. Similar compounded binary option transactions not dependent on OP_CTV have been prototyped, resulting in a smoothed out PNL similar to a vertical options spread, there is potential in that research. But wouldn't the best thing be a new OP_Code that directly does the UTXO wrapping? That is why the TradeLayer team supports a proposed OP_Repo that would take any Layer-esque protocol's tx marker as a parameter, cache UTXOs natively, create repository tokens for use in a layer protocol, and only pay out to a redeemer of those tokens.

In legacy finance, the repo market for high-quality collateral is an often misunderstood but fundamental part of the system, its carotid artery. rBTC tokens that are nativel redeemable for UTXOs could become similarly important to the future evolution of TradeLayer and Bitcoin as a full-stack financial system.

9. Custodial Security

There are two main security techniques needed to make decentralized currency secure for most people, who are generally exposed to the risk of physical duress over bearer assets. The first technique is segregating assets over time, a Savings Address feature allows rate-limits, time-locks and cancelable cheque payments that ameliorate the risk of non-physical compromise (i.e. 'hacking').

The second technique involves segregating control over the assets with multisig back-up key holders; traditionally such a strategy only works for people hiring custodial trustees or hiding hand-written keys in miscellaneous dig sites and bank deposit boxes. To make safe custody truly scalable, the cost of customer service for modularly factoring out multisig custody must approach a combination commission-based finder's fee for lost keys' assets, and a low fixed annual fee. State Channels involving bonded collateral may be a semi-decentralized way to hold such low-cost trustees accountable.

In a physical duress situation a pure Savings Address defense makes violence the high probability strategy for the assailant, to prevent the holder from canceling the transaction. Whereas, a 3rd party service needing secondary verification within twenty minutes to not cancel the payment, only operating during daytime business hours, changes the game theory to favor faster, less violent shakedowns for loose single-signature pockets of funds. It's the combination of the two that solves custodial security.

10. Calculations

Max Negative Yield based on Borrow Rates

0 = (Fi(x) - (BTC Borrow Rate + Metacoin Borrow Rate + Fees))-Risk-free Rate

Fi is a function that takes the average discount or premium of a perpetual swap contract, x, and returns an annualized rate of return based on the resulting swap payments. This income calculation function may be complex pending further research, and is truncated here for simplicity. If the swap longs can borrow and sell-short at a certain rate, they must receive an income rate that will pay a net-return at least as large as the prevailing risk free rates available in the derivatives markets.

Counterparty Value Adjustment:

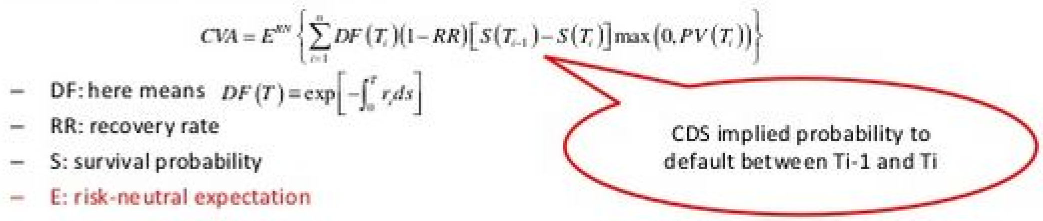

$$CVA = E^{RN}\left\{\sum_{i=1}^{n} DF(T_i)(1-RR)\left[S(T_{i-1})-S(T_i)\right]\max\left(0, PV(T_i)\right)\right\}$$

- DF: here means $DF(T)=\exp\left[-\int_0^T r_s ds\right]$
- RR: recovery rate
- S: survival probability
- E: risk-neutral expectation

CDS implied probability to default between Ti-1 and Ti

11. Conclusion

While on one hand many of the problems Bitcoin sought to solve have yet to be solved, it's encouraging that the size, intellectual heft and temerity of the open source scientific movement it launched has given us enough shoulders of giants to climb over in presenting the solutions described in this paper.

Bitcoin succeeded because it combined four novel R&D threads from various predecessors, and leveraged the popular seigniorage of the block reward to encourage more network value than diluted through inflation. The same is going to be true for decentralizing trading liquidity.

The power of public-private key cryptography enables fast trading with strong assurances in a decentralized environment native to the Bitcoin protocol. There are unique challenges to trading in this environment but also quantitative methods for high frequency algorithms to model and adjust. With enough liquidity, we can facilitate a lot of volume, reducing the high fees and custodial risk in the industry, and perhaps one day having enough UXTO->Token volume that native contracts become manipulation resistant and totally independent currency backed by this Bitcoin derivatives trading becomes possible. The reason why Bitcoin is interesting monetarily is that its pre-programmed inflation rates and resulting scarcity create a benchmark for the time value of money that the centrally manipulated interest rates of fiat currencies could wither against. By pricing derivatives premium in this decentralized environment, we can create the first truly free market to price-discover the time value of money in human history.

Peer-to-peer electronic cash is a basic human right. Human beings should not have to sign an unfavorable legal contract in order to have the ability to transact in the stable unit of account most popular in their home region. We can bring decentralized currency to billions of human individuals who most profoundly deserve access to it.

References

[1] Bitcoin Whitepaper
https://bitcoin.org/bitcoin.pdf

[2] OmniLayer Reference Spec
https://github.com/OmniLayer/spec

[3] Dr. Christian Decker, OP_SIGNOINPUT
https://en.bitcoin.it/wiki/BIP_0118

[4] Antoine Savine, A Brief History of Discounting:
https://www.slideshare.net/AntoineSavine/a-brief-history-of-discounting

[5] Futures Manipulation with “Cash Settlement”
https://sci-hub.tw/10.1111/j.1540-6261.1992.tb04666.x

[6] Jeremy Rubin, OP_CHECKTEMPLATEVERIFY:

https://github.com/bitcoin/bips/blob/master/bip-0119.mediawiki