MULTICOIN CAPITAL

Token Curated Registries: Features & Tradeoffs

By Tushar Jain, Spencer Applebaum, Myles Snider, & Kyle Samani

Multicoin Capital is a thesis-driven cryptofund that invests in tokens reshaping entire sectors of the global economy. We rigorously research blockchain protocols, teams, and market opportunities to deliver venture capital economics with public market liquidity.

Report Disclosures

As of the publication date of this report, Multicoin Capital Management LLC and its affiliates (collectively "Multicoin"), others that contributed research to this report and others that we have shared our research with (collectively, the "Investors") may have long or short positions in and may own options on the token of the project covered herein and stand to realize gains in the event that the price of the token increases or decreases. Following publication of the report, the Investors may transact in the tokens of the project covered herein. All content in this report represent the opinions of Multicoin. Multicoin has obtained all information herein from sources they believe to be accurate and reliable. However, such information is presented "as is," without warranty of any kind – whether express or implied.

This document is for informational purposes only and is not intended as an official confirmation of any transaction. All market prices, data and other information are not warranted as to completeness or accuracy, are based upon selected public market data, and reflect prevailing conditions and Multicoin's views as of this date, all of which are accordingly subject to change without notice. Multicoin has no obligation to continue offering reports regarding the project. Reports are prepared as of the date(s) indicated and may become unreliable because of subsequent market or economic circumstances.

Any investment involves substantial risks, including, but not limited to, pricing volatility, inadequate liquidity, and the potential complete loss of principal. This report's estimated fundamental value only represents a best efforts estimate of the potential fundamental valuation of a specific token, and is not expressed as, or implied as, assessments of the quality of a token, a summary of past performance, or an actionable investment strategy for an investor.

This document does not in any way constitute an offer or solicitation of an offer to buy or sell any investment or token discussed herein.

The information contained in this document may include, or incorporate by reference, forward-looking statements, which would include any statements that are not statements of historical fact. These forward-looking statements may turn out to be wrong and can be affected by inaccurate assumptions or by known or unknown risks, uncertainties and other factors, most of which are beyond Multicoin's control. Investors should conduct independent due diligence, with assistance from professional financial, legal and tax experts, on all tokens discussed in this document and develop a stand-alone judgment of the relevant markets prior to making any investment decision.

Summary

INTRODUCTION

Token Curated Registries (TCRs) are an idea that gained a fair amount of traction within the cryptocurrency community throughout late 2017 and early 2018. While the idea has yet to be implemented at any meaningful scale, there are a few TCRs live on the Ethereum mainnet and over a dozen projects are considering building TCRs in some form.

In simple terms, a TCR is a list with a native token that is used for curation and governance rights. TCRs aim to create economically incentivized systems for generating high-quality signals. Some crypto enthusiasts have described them as a "new form of crypto-economic primitive." While that's yet to be determined, they can be thought of more concretely as lists that utilize blockchain tokens on the back end to keep track of curation rights.

Humans utilize and rely on lists for all sorts of information gathering. Most lists are simply curated by individuals or entities. These curators create lists for all sorts of different reasons. In the case of websites like Eater or Grubstreet, these publications create high-quality lists of restaurants and food destinations in order to attract more users to their sites, which allows them to charge more from advertisers. In the case of Spotify playlists, curators often create these lists for social recognition, for promotion of affiliated musicians, or simply for altruistic reasons. Other lists like GiveWell.org are created by non-profits simply to offer reliable information about a particular subject.

Proponents of TCRs argue that list curation can be radically improved by giving key participants skin in the game in the form of list-specific tokens that fluctuate in value and can be earned or lost through participation, popularity, and correctness. They believe that an economic system designed around this token can ensure better-quality curation and participation. Further, TCRs aim to be self-sustaining public utilities; rather than being dictated by any central party who can shut it down or change it arbitrarily, the list itself becomes a self-contained system accessible by anyone and controlled by no single entity. Others posit that these systems are at best no improvement upon the status quo and at worst economically broken.

In this paper, we will explore the promises, features, trade-offs, and weaknesses of TCRs. If TCRs are to be useful, they must offer substantial benefits over existing systems for creating and accessing lists, they must be resilient to attacks, and they must actually conform to the economic assumptions outlined above.

BACKGROUND

Mike Goldin of ConsenSys proposed the first TCR in the TCR 1.0 white paper (and later revised in the TCR 1.1 blog post), as a part of AdChain in early 2017. Goldin's goal with TCRs was manifold: to decentralize the curation process for any type of list, to economically incentivize high-quality curation, and to cultivate better signals by forcing participants to have skin in the game.

Goldin's paper served as an early introduction to the idea of TCRs, but he acknowledged limitations and potential attack vectors in his design and encouraged others to build on the ideas he proposed. Since that time, many different projects have done just that. Goldin himself has explored problems and potential solutions, and many other teams have come up with unique approaches to building TCRs. Goldin also built the first live TCR in AdChain.

AdChain is the first production TCR, and it will be important to watch how well it performs in the real world. Many other prominent projects, including Messari and District0x, are currently in the process of building TCRs. The recently announced DIRT Protocol is creating a generic framework for TCRs. Some of these projects utilize the same basic model proposed by Goldin, while others have changed the model in important ways. We'll examine the different approaches in this paper.

TCR Design

In its most basic form, a TCR is a list whose items are decided upon by a set of token holders who have staked deposits in a one-token-one-vote system. In many TCR designs, a user's tokens can be slashed—aka programmatically and automatically confiscated—under some set of (presumably bad) circumstances. Users who propose that an item be included in, excluded from, or removed from a list must stake a deposit, or in other words, put their money where their mouth is. After staking tokens, anyone else can challenge the initial claim by staking an equal (or in some cases greater) number of tokens. Token holders (of the same TCR) ultimately vote on the proposal, and the people who proposed the winning proposal are rewarded and the people who proposed the losing proposal are slashed.

As Goldin describes, most TCRs have three types of participants: consumers, candidates, and token holders.

- **Consumers** are those who utilize the lists for information discovery
- **Candidates** are those who wish to be included on the list
- **Token holders** are those who participate in list curation

In addition to the three major categories of participants, each TCR has a **native token**. The native token of the TCR represents *ownership in the curation rights of the list being generated*. Token

holders can vote on many things: what is included in the registry, who can be added and removed and when, who wins in the case of disputes, what the system parameters are, and more.

The TCR framework is designed as follows:

Consumers interact with the TCR as they'd interact with any list: They use it for the purpose of information discovery. Just like restaurant-goers use Yelp to find a list of "the best Italian restaurants in Manhattan," consumers of other types can use TCR-generated lists in the same way. Consumers have no preference about whether or not the curation process is centralized or decentralized. They simply want good information and will stop using sources that don't provide that. Well-designed TCRs should offer consumers an interface that's comparable to a list provided by a centralized service.

Candidates are those that want to be included on the list. If the list were a ranking of the top venture capital firms, then those firms would be incentivized to be on the list as it may help them attract more investors (who, in this case, would be the consumers). Candidates must pay a deposit (paid in a certain amount of the TCR's native token, in most cases) in order to be considered for a listing. If the candidate is accepted by token holders, she keeps her deposit as a slashable stake. If she is rejected, the deposit is confiscated, and possibly redistributed to tokenholders. This is meant to disincentivize bad applicants, but it may ultimately disincentivize even honest applicants, as we'll explore later.

Token holders are the most important stakeholder group in a TCR. Token holders are those who own the native token and thus participate in voting and curation.

In a simple implementation, candidates apply to be included in the list by staking a deposit. Most of the time, the deposit is made in the native token of the TCR, though some models require deposits made in other assets such as ETH. During the review period that follows, token holders vote on whether the candidate merits inclusion in the list. Any user who does not think that the candidate should be included must issue a challenge by staking a matching deposit.

If a challenge is issued, then token holders must vote on whether the candidate should be included or not. Most TCRs use a commit-and-reveal voting scheme to prevent voters from simply siding with the existing majority. Once a certain threshold of votes is reached and the voting period ends, a decision is finalized, and the losing party is slashed and the winning party is rewarded. In most cases, half of the losing party's tokens go to the winning party, and the other half go to the voters who participated and voted for the majority decision. Some TCR implementations also implement a slash on the token deposits of the *voters* who voted for the minority decision, to prevent token holders from splitting their voting power evenly between sides in order to get a guaranteed payout.

If no one presents a challenge and the candidate is accepted, her deposit is kept in a smart contract as a slashable stake. However, even after a candidate is included in the registry, any user can

present a challenge and ask token holders vote on it. This is intended to keep candidates honest even after being accepted.

The slashing mechanism and reward split described above serve three purposes:

1) Incentivize token holders to actually participate in voting
2) Incentivize users to challenge bad candidates while also discouraging malicious challenges
3) Disincentivize low-quality candidates from applying for inclusion in the registry

The parameters of individual TCRs can be optimized in different ways. Each TCR must specify the number of votes required to enact a decision as well as a time frame for decision making.

Advantages of TCRs

TCRs offer a few notable advantages over existing curation models:

DECENTRALIZED CURATION

The primary proposed benefit of TCRs is that the curation process becomes decentralized. Instead of having a single company, individual, or organization that controls the list, TCRs empower a decentralized group of actors to do so. This limits the power of any single entity, eliminates points of failure, and can in some instances offer a more reliable information signal.

CENSORSHIP RESISTANCE

TCRs, by their very nature, offer a form of censorship resistance that centralized services cannot. Because no single entity controls the list, TCR token holders would require majority collusion in order to block any candidate from being listed or delisted. Further, any participant with access to the market can buy tokens and contribute to curation. Centralized alternatives have the ability to censor arbitrarily at will. The decentralized nature of TCRs also means that they cannot be shut down.

The TCR model also has the potential to lower the barriers to entry for participation in the curation process. Anyone with a wallet and access to the token markets can participate in the TCR, without requiring the permission of any central party.

TRANSPARENCY

TCRs record all actions and decisions on a blockchain, so there is a clear audit trail of every decision that was made. The entire history of the list, with its various additions and removals, is recorded. All

of the governance decisions made are in the public domain, which adds a level of transparency that almost no centralized entity offers.

UP-TO-DATE INFORMATION

TCRs offer incentives to keep curators engaged such that the information in the list is curated on an ongoing basis. Rather than creating a single list that is then published in its final form, TCRs are ongoing lists that can change as new information is incorporated.

REWARDING MODERATORS

In other models (Yelp, for example), any user can participate, but these users aren't rewarded for their work. This creates a few major problems. Only a small set of users actually participate, and their inputs are often subjective and unreliable. Because participation doesn't have a cost, the system can easily be gamed. And because users are not incentivized to contribute, very few actually do.

ECONOMIC INCENTIVES AND PENALTIES

As mentioned above, existing curation models don't always give participants economic incentives. In some cases, list providers are motivated to increase pageviews in order to earn more from advertisers. These incentives don't always directly result in high-quality curation. In other instances, neither the curators nor the candidates have any value on the line. This can lead to gamification of the system, and it can also lead to poor signals.

The traits described above examples are *theoretical* advantages of TCRs, as almost none have yet been observed in practice. It's not yet clear that they will play out in the real world, or whether these advantages outweigh some of the additional problems TCRs present.

Disadvantages of TCRs

The first question to ask in examining TCRs is whether centralized list curation is actually a problem at all. Centralized list providers certainly have the ability to do harm. They can censor participants, offer poor curation services, accept bribes, and shut down. In some cases they fail to reward participants, while in others the rewards are based on metrics (like SEO-driven page views) unrelated to the quality of the actual list. But TCRs are not immune to all of these issues, and and as a result of being decentralized, suffer from a new class of problems.

The foundational assumption of TCRs is that token holders will be incentivized to *honestly* curate a high-quality list because doing so will cause the value of their tokens to increase.

Whether or not most TCRs will work in practice depends a lot on the validity of this assumption. In theory, a good list will attract the best applicants, who will then be willing to pay more to be included. This will ultimately drive demand for the token, which will increase its value.

However, in many instances this fundamental assumption fails in the real world. The dynamics of TCRs do not *necessarily* incentivize high-quality curation. Rather, they incentivize token holders to act in a way that maximizes the value of their holdings and earnings. In some cases this may be aligned with high-quality curation, but not in all. Given certain attack vectors that leverage flawed risk/reward mechanics, users may generate more profits by attacking or gaming the system than they will by acting honestly. Further, there are many scenarios in which TCRs are objectively worse than centralized list-making services.

OPEN MARKET COMPETITION

One can reasonably argue that decentralized list curation offers almost no tangible benefit over centralized list curation given open market competition among centralized list providers. As an example, let's look at the restaurant review industry. Yelp, Eater, and The Infatuation all offer competing services to provide users with lists of the best restaurants in a given area. Each of these services is centralized in the sense that it is controlled by a single company, and each has a business model (advertising) that provides a financial return for providing list curation services, among others. The services are free to all users, and users are also free to switch between services for any reason or no reason at all. The administrators of these lists are attempting to maximize the value they can get from advertisers, just as TCR token holders are attempting to maximize the amount of value they can get from their token holdings. The former is a proven business model, while the latter is entirely experimental. And while there are issues with each of these centralized businesses, competition amongst them, plus zero switching costs for users, generally negates many of these issues. These models are satisfactory for the vast majority of users. We have yet to see a case where decentralized list curation offers objectively better products and incentives than centralized alternatives.

ADDRESSING SUBJECTIVITY AND UNOBSERVABLE TRUTHS

Another major issue with TCRs is that they can only function when the information being voted on is publicly observable, objective, and easy to verify (Aleksandr Bulkin of CoinFund explores this topic thoroughly in this article).

The reason that TCRs do not work well for subjective information is that the curation process quickly becomes a Keynesian Beauty Contest. For token holders, the rational strategy is not to vote on the candidates that they personally find to be of the highest quality, but rather the candidates that they think *other voters* will perceive to be of the highest quality. Subjective criteria offer no tangible coordination signal for voters. Further, users may be disincentivized to challenge listings that they disagree with. Why? Because they are taking on higher risk in doing so. If users are voting

on a publicly verifiable and objective truth, then they know that there is a strong coordination signal for all voters. If they are voting on a subjective truth (best song, best restaurant, best doctor, etc), then a user who issues a challenge is not taking a risk based upon something they know to be true, but rather something they subjectively observe and hope that other token holders also agree with, regardless of the underlying truth, if such an underlying truth can even be determined.

Imagine I am a token holder in a TCR that ranks the best doctors in a given area. If I have a bad experience with a particular doctor, then I can issue a challenge to the doctor's inclusion. But this is a risky strategy, since my bad experience with this doctor could be a subjective one that may not be shared by other voters. I run the risk of losing my deposit if voters disagree with me, and there is no reliable way to predict voters' subjective interpretations.

This issue has recently been raised with the AdChain registry. AdChain's goal is to provide advertisers with a list of high-quality online publishers who do not try to game page views (gaming page views is a way for publishers to effectively steal money from advertisers). Interestingly, Facebook was rejected from the registry because certain token holders don't trust Facebook (as evidenced by this Reddit thread). The goal of the AdChain TCR isn't to protect users' best interests, but rather to show advertisers the best place to display their ads, which at times might conflict with user interests on a given site. Although there are many reasons to dislike Facebook, there is no material evidence that suggests that Facebook is intentionally gaming its viewership metrics (though there have been bugs in the past that Facebook has disclosed).

This is also a problem if the information being voted on is not publicly or easily observable and verifiable. If I have privileged information about a candidate, that may change my decision about whether or not she merits inclusion in the TCR. I can issue a challenge based on that information, but voters who do not have that information cannot act upon it. This can be a problem even if the information itself is objective. I can go public with the information, but voters who cannot *verify* that information for themselves once again won't be likely to act upon it. This issue can break the entire system, since information that is highly relevant to the curation of the list simply won't be taken into account, and the quality of the list will suffer.

Centralized list providers, on the other hand, do not suffer from this issue. They can change a list based on any information they have, whether it is objective, subjective, public, or private. This moves the layer of competition from the list itself to competition among lists.

TOKEN PRICING AND CUSTOMER ACQUISITION

One major issue with TCRs is that the price of entry must be high enough to discourage bad candidates, while also being low enough to not price out honest entrants entirely. If token holders are driven by an incentive to maximize the value of the tokens, then the token may become expensive enough that it simply doesn't make sense for candidates to spend the money to apply. If the token price is too low, then the system becomes much easier to attack.

TCRs that offer cheap entry will be flooded with low-quality candidates. Token holders may be unable to screen each one individually, and some poor-quality candidates are likely to slip through and be included, defeating the purpose of the TCR entirely. Further, if candidates don't have to bear any real cost for inclusion in the registry, they will have very little incentive to abide by the terms of curation. This results in a low-quality list that consumers will have no use for.

On the other hand, TCRs with expensive entry simply price out high-quality candidates. Candidates for many types of lists are not used to having to pay for inclusion and in some cases may be completely unwilling to do so, especially for new and unproven TCRs with few consumers. Further, if candidates are required to stake a large, expensive deposit, then challenges become prohibitively expensive too, since challengers must match the candidate's deposit. Striking and maintaining a balance between entry price, challenge price, and gameability will be a challenge for all TCRs.

One way to avoid some problems here is to have deposits and challenges paid in ETH or stablecoins, rather than being paid in the native token. This eliminates the friction points of fluctuations in the price of the native token, and also changes the basic model of the token. The token price becomes a function of the cash flow users can earn, rather than the demand from candidates purchasing it. However, even with this model, token holders will have to coordinate to come up with the best price, in ETH or dollar terms, that strikes a balance between pricing and customer acquisition.

For centralized list providers, consumers are their customers. They can cater their services to offer the best experience for consumers. In TCRs, both consumers *and* candidates are stakeholders, and they have divergent interests because earnings and token price appreciation are mostly a function of candidates applying. This creates a misalignment of incentives in TCRs; voters should want to offer the best experience and information to the consumers, but they also must cater to the wants, needs, and preferences of candidates, or the TCR will not not work. That which is good for consumers isn't always good for candidates, and vice versa. Managing divergent stakeholder interests algorithmically is extremely challenging, and few have done it successfully.

UTILITY OF PUBLIC REPUTATIONS

When consumers seek lists, one of the primary attributes they assess is the public reputation of the entity providing the list. This is one of the main reasons why centralized list providers are able to attract consumers—they build a public reputation for providing high-quality signals over an extended period of time. This is an extremely important and powerful part of the reason that this model works.

With TCRs, the notion of verifiable, public identity is lost in almost every implementation. Token holders can be entirely anonymous and have no reputation at stake. While this could be an argument for centralized list curation over decentralized list curation, it could theoretically be

possible for TCRs to implement some sort of reputation system for participants. This could involve a separate protocol that allows people to build up reputation points, or it could just require users to tie their actual identity to their tokens using something like Civic.

Use Cases for TCRs

While we remain generally skeptical of most TCR implementations, there are some narrow use cases in which the unique features offered by TCRs outweigh some of the disadvantages of the model.

TCRs could be useful in cases where the following are all true:

- There must only be one canonical list
- The list must be censorship resistant (related to above)
- Applying to be part of the list must be permissionless and global
- The information being presented is objective, public, and relatively easy to verify
- Centralized alternatives *don't exist* or *function poorly because of improper incentives*

Further, we generally believe that the TCRs that succeed will not follow the basic TCR model, but rather will structure their token as a work token, in a fashion similar to the method employed by Messari. This model offers more value to token holders and avoids some of the pricing issues for new candidates.

It's likely that 2019 is the year we see TCRs become tested in battle. Given the number of implementations currently in development, it seems clear that soon the market will test many of the assumptions explored in this document. We have not yet seen a TCR that satisfies all the criteria listed above and is also close to going live. We don't yet have strong conviction that these systems will work, even if these criteria are met. But we look forward to seeing TCRs of this type deployed.

Conclusion

TCRs are nascent and experimental. The few existing implementations are in their earliest stages, and leave critical questions unanswered. With a dozen or more TCR implementations slated to debut within the next year, we'll soon see market stress test core TCR mechanics.

It's entirely possible that hobbyist TCRs emerge for things like playlist curation, movie rankings, and other subjective criteria. They could become a fun way for groups of enthusiasts worldwide to collaborate and share information. But the value of the native tokens for these TCRs will likely be in the form of social capital rather than financial capital—digital "brownie points" or "credibility" with near zero direct financial value at stake. As soon as any of these systems become economically

valuable enough, users will be incentivized to game and attack them. Given that all existing TCRs are proof-of-stake systems, their security stems, at least in part, from the network's value. Low-value systems are inherently less secure.

We expect that most TCRs will be unable to compete with centralized alternatives. We also expect to see some TCRs successfully attacked by malicious actors. However, we believe it is possible that we will see some successful niche TCR implementations that function as intended and provide robust information signals.

It's possible that TCRs, in their current form, are simply the first stages of what might become a useful tool for digital information sharing. Information markets that reward participants are certainly something we could foresee in the future. However, most TCRs in their current form are not complete or fully optimized systems. Improvements to the model must be made, and critical flaws must be addressed before these systems can become widely useful at scale. We remain skeptical of most current implementations, but look forward to seeing more advancements in this unique crypto-economic design space.