| Tomas Vrba

| Global Technology and Development
| Risk Innovation Lab
| tvrba@asu.edu

GTD 593 Applied Project
School For the Future of Innovation in Society
Arizona State University

Decentralized Governance, Reputation and Mechanism Design for a Global Social Impact Funding Platform on a Public Blockchain

Proposal Filtering, Hyper-Transparency and Tokenization in Trustless Multi-Agent Systems

April 2018

I. Introduction

The main focus of this project is to contribute to the design and creation of a "Charity DAO", a Decentralized Autonomous Charitable Organization (DAO) built on top the public Ethereum blockchain. The goal of this organization is to use blockchain technology to bring more transparency and trust to the charitable, humanitarian, non-profit and foreign aid sectors by providing community and evidence driven funding and investment. Specifically, by building a decentralized and permissionless social impact funding platform accessible to anyone globally in a peer-to-peer fashion driven by the values of effective altruism and social justice. The primary contribution of this project is in the area of governance of decentralized systems in the form of designing a proposal filtering bond and a decentralized reputation system framework.

The main focus of this project are the technical aspects of economic incentive-driven governance in decentralized systems using mechanism design—specifically, designing a bond for filtering funding proposals on a decentralized social impact funding platform; a reputation framework for this platform, and blockchain ecosystem integration. A secondary section begins with presenting a proof of concept of Hyper-Transparency on the public Ethereum blockchain. This is a concept developed specifically with social impact and charitable funding platforms in mind. This work is the result of a three-day, four-person collaboration at the ETHWaterloo Hackathon. The two main outcomes here are the introduction of the concept of Hyper-Transparency and its use-cases in certain contexts; and the development of the Title Token Standard.

The first section of this paper involves technical research into decentralized governance and reputation systems, came about in collaboration with a German non-profit foundation aiming to establish a permissionless, transparent, scalable, and decentralized social impact and funding platform on the public Ethereum blockchain. The working name for this foundation and project is "Charity DAO" which stands for Decentralized Autonomous Organization and which I will also use interchangeably with "Social Impact Funding Platform." Its aim is to bring more transparency and trust to the charitable, humanitarian, non-profit and foreign aid sectors by supporting community and evidence driven investment. The research and conclusions from this project are however more widely applicable in blockchain ecosystems when it comes to implementing governance and incentive models in these decentralized systems.

This project builds on existing research by creating a framework for how reputation could be handled on a decentralized social impact funding platform and adds to the volume of research by providing a practical applied model for handling funding requests in an attempt to limit the role of intermediaries such as curators on a permissionless funding platform. The desired objective of the bond model is to filter and disincentivize via economic means non-serious, fraudulent or spam requests for funding which place a burden on voters and raise the cost of capital for everyone, as well as to create incentives for funding requests to always remain in the area of skill and competence of proposers and to not request outsized funding amounts ill suited to their goals. As such, this bond falls into the field of economic mechanism design concerned with designing incentives or mechanisms with a given goal. From the perspective of decentralized governance models, this mechanism aims to manage incentives and penalties for good and bad behavior with cryptoeconomic principles in mind, ie. minimizing the role of social trust assumptions and utilizing clearly defined rules to manage consensus.[1] As expanded upon below, this bond mechanism could be widely applicable in distributed networks to handle proposals without the need for curators.

The final section introduces the concept of Hyper-Transparency and a proof of concept for a Title Token Standard which were the result of a four-person team effort at the ETHWaterloo Ethereum Blockchain Hackathon in October 2017.[2] The last section builds on this concept and further discusses the justification, use cases and applications of this standard as well as other potential socio-economic impacts of blockchain protocols.

II. Decentralized Governance, Reputation and Mechanism Design for Proposal Filtering

A. Why Use Blockchain for Social Impact

1. What is the blockchain

Despite the rise in popularity of cryptocurrencies, there are many misconceptions about what blockchain technology can and cannot do, what it is good for and how it works on a technical level. Before getting into the "why" of using blockchain technology for social impact, a brief overview of the technology is necessary to facilitate a basic understanding of what it is and how it works.

A definition of a blockchain may change depending on what perspective the technology is approached from. Technically, business-wise, legally, or politically, a blockchain may mean different things. To quote William Mougayar "Technically, the blockchain is a back-end database that maintains a distributed ledger that can be inspected openly. Business-wise, the blockchain is an exchange network for moving transactions, value, assets between peers, without the assistance of intermediaries. Legally speaking, the blockchain validates transactions, replacing previously trusted entities."[3] Even bearing this in mind however, one of the core takeaways when it comes to discussing blockchain technology is that "there is no universal definition of a blockchain, and there is widespread disagreement over which qualities are essential in order to call something a blockchain."[4] "Blockchains" can be public, private, permissioned, permissionless, closed-source, open-source, decentralized or centralized (at which point we are really just referring to a shared database). The point is, context and clarity of language matter when discussing blockchain technology, its use cases, and potential impacts. In the context of

---

[1] Buterin, V. (2018). *Governance, part 2: Plutocracy is still bad*. Retrieved from https://vitalik.ca/general/2018/03/28/plutocracy.html

[2] Bauza, L., Lockyer, M., Vrba, T., & Yambao, M. (2017). *The DAC (decentralized autonomous charity).* ETHWaterloo: Retrieved from https://devpost.com/software/the-dac

[3] Mougayar, W. (2016)., *The business blockchain: Promise, practice, and application of the next internet technology* Hoboken, New Jersey : John Wiley & Sons, Inc.

[4] Jeffries, A. (2018). 'Blockchain' is meaningless. Retrieved from https://www.theverge.com/2018/3/7/17091766/blockchain-bitcoin-ethereum-cryptocurrency-meaning

this paper and the decentralized governance topics it covers, I am focusing purely on public, permissionless, open-source blockchain technology and particularly Ethereum.

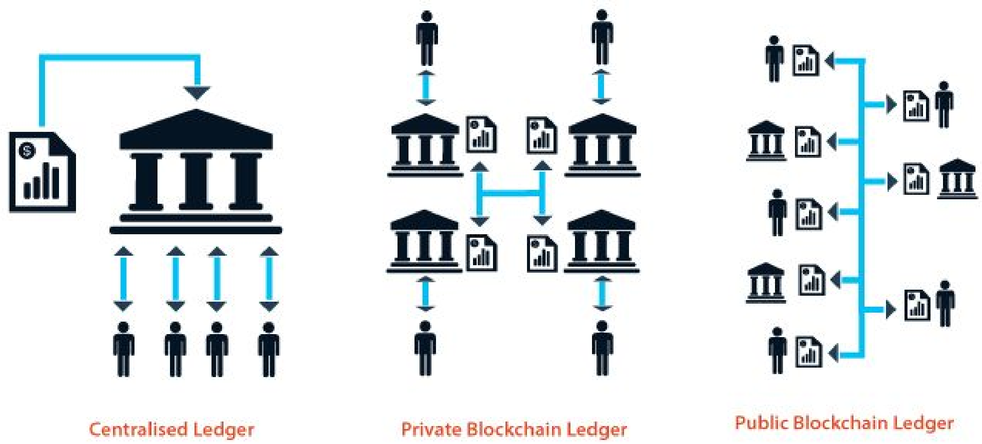

Figure 1. Blockchain, Cryptoeconomics, and the Disintermediation of Trust[5]

While at their core, most blockchains are network databases secured by public key cryptography and whose rules, security and value are driven by individual self-interest, economic incentives, and the principles of game theory (a field of economics that focuses on mathematical representation of conflict and cooperation between rational actors), they are fundamentally a *political* technology driven as much by ideology as by technological and technical innovation. Instead of relying on organizations and institutions such as central banks and governments to provide security features, anti-counterfeiting measures, law enforcement and regulated money supply, cryptocurrencies, as the name suggest, utilize cryptography and public-private key combinations to embed trust, security and rules in the system itself. Thus one of the important aspects of blockchains to keep in mind is that their importance does not necessarily lie in *what* they do, but in *how* they do it. Most anything that can be done with a blockchain, can be done with centralized databases and delegated trust as well. The problem with these systems is however their susceptibility to corruption and manipulation. Thus "Blockchains are far more than a technical solution to solve the double-spend in digital cash; they change the balance of power in networks, markets and even the relationship between the individual and state, all characteristics of Web 3.0"[6]

"Charity DAO" is a non-profit foundation whose aim is to use blockchain technology, specifically the public Ethereum blockchain to bring more transparency and trust to the charitable, humanitarian, non-profit and foreign aid sectors by providing community and evidence driven funding and investment. The goal of the organization is to build a decentralized and permissionless social impact funding platform accessible to anyone globally in a peer-to-peer fashion driven by the values of effective altruism and social justice. Yet in being fully decentralized and permissionless, such a funding platform will have to address many potential attack vectors and perverse incentives which could make it a target for exploitation.

[5] Akmeemana, C. (2016). Blockchain, cryptoeconomics, and the disintermediation of trust. Retrieved from https://www.linkedin.com/pulse/blockchain-cryptoeconomics-disintermediation-trust-akmeemana-14-000-/
[6] Outlier Ventures Research. (2017). *Blockchain-enabled convergence: Understanding the web 3.0 economy.* ().Outlier Ventures Research.

Decentralized autonomous governance models, economic incentives for project evaluators and reputation systems will all be key to effectively protecting such a funding platform from malicious actors seeking to misappropriate or launder funds. This project proposes one proof of concept for such a governance model—a dynamic proposal bond to filter incoming funding proposals without the need for human curators, and one proof of concept of a hyper-transparent system for tracking and monitoring donations and investments.

2. Donor trust, transparency and global deployment

In building a social impact platform whose aim is to facilitate philanthropically driven donations, investments and wealth distribution, one needs to consider where the funding for this platform will be coming from, what are the obstacles standing in the way, and which can be addressed using blockchain technology and how. While macroeconomic factors certainly play a role in the amount of charitable giving, as does the declining membership in religious organizations among young people, these are not necessarily problems that can be tackled with innovative technological solutions. At least not directly. However, another significant reason for why charitable donations and volunteering have been growing at a slower pace than GDP,[7] or even declining by 11% since the early 2000s by other metrics,[8] is lack of transparency and perceived (or real) lack of transparent results and impact evaluations by charitable organizations.[9] The transparency of public blockchains lends itself perfectly to addressing these issues. Using blockchain technology we can build a decentralized autonomous social impact platform that can accept donations in any currency, hold onto its funds in a non-volatile form and deploy them globally within seconds while maintaining complete transparency of individual donations. Such hyper-transparent funding, impact evaluations and auditability combined with disintermediation of financial services have the potential to vastly improve the donor trust in and efficiency of charitable giving and non-profit funding.

One of the leading beliefs behind decentralized funding platforms such as Charity DAO is that fully transparent charitable organizations with decentralized governance can increase trust and efficiency in how charitable donations, development funding, and foreign aid are raised and distributed. The goal is to provide decentralized fundraising and fund distribution anywhere in the world and quickly deploy raised money for communities in a transparent way. Lack of transparent impact evaluations can make it difficult for potential funders to evaluate social impact projects and separate effective project from ineffective ones, leading to inefficiencies and larger administrative costs for everyone involved. Permissionless blockchains as a technology which facilitates trust and disintermediates transactions and contracts is uniquely suited to address these issues.

---

[7] Blackbaud Institute. (2016). *5th annual charitable giving report: How nonprofit fundraising performed in 2016.* ().Blackbaud Institute. Retrieved from https://institute.blackbaud.com/wp-content/uploads/2017/02/2016-Charitable-Giving-Report.pdf

[8] Generosity for Life, & Indiana University Lilly Family School of Philanthropy. (2018). Charitable profile. Retrieved from http://generosityforlife.org/generosity-data/data-tools/generosity-reports/

[9] Flandez, R. (2012). Donors say they would give more if they saw more results. Retrieved from https://www.philanthropy.com/article/Many-Donors-Would-Give-More-if/156463

A large portion of current blockchain applications fall in the field of financial technology (fintech), aiming to disintermediate and streamline payment and value transfer systems. As such, blockchains are also well suited to the charitable giving, foreign-aid and development sectors which are increasingly utilizing direct payments and cash transfers. Studies are showing[10] these can sometimes be a very effective type of giving as exemplified by organizations such as GiveDirectly[11] as opposed to material aid. The Food and Agriculture Organization of the United Nations (FAO) for example has during recent years gone from 6% of their funding in the form of cash transfers to more than 25%.[12] We are not arguing however that all development funding should be in the form of direct cash transfers, it is not a one size fits all solution.

B. Decentralized Governance and Reputation for Multi-Agent Systems

1. Charity DAO Economic Incentive Model for Proposal Filtering

a) Theoretical Bond Model

In the "ARES Protocol White Paper" Mark and Dodson propose a mechanism for eliminating human curators from the DAO voting process.[13] The need for human creators arises from the fact that in a decentralized, permissionless system, anyone is allowed to submit proposals and there needs to be a way to filter out spam and fraudulent proposals among other kinds of undesirable content. Human curators are trusted individuals selected by the community to moderate this process, yet this introduces elements of centralization, delegation of trust and most of all inefficiency and security (corruption) vulnerabilities. The goal would thus be to create a DAO which can self-regulate the proposals introduced to the voting pool without an intermediation by human curators. Mark's and Dodson's proposed way of achieving this is setting a bond for a proposal prior to allowing voting. "The bond will be forfeit to the Fund if the proposal fails, but returned to the proposer if the proposal succeeds." However, at 100 Ether per proposal submission (on the order of tens of thousands $USD) such bond is not well suited for a social impact funding platform which aims to target populations with limited economic means. Moreover, the size of the bond can only be adjusted by every six months. This is an unnecessarily static and restrictive governance model. These two aspects of the original bond proposal, its size and its fixed nature is where it fails as a functional governance mechanism. Although it may be better suited for venture capital and other such investment projects with large amount of capital, even then, this bond would ideally be determined by market forces based on a variety of factors such as the type of proposal and the requested amount of funding.

[10] From evidence to action: The story of cash transfers and impact evaluation in sub saharan africa (2016). . Oxford: Oxford University Press.10.1093/acprof:oso/9780198769446.001.0001 Retrieved from http://www.oxfordscholarship.com/10.1093/acprof:oso/9780198769446.001.0001/acprof-9780198769446
[11] Haushofer, J. (2016). The short-term impact of unconditional cash transfers to the poor experimental evidence from kenya. The Quarterly Journal of Economics, 131(4), 1973-2042. 10.1093/qje/qjw025

[12] Food and Agriculture Organization of the United Nations. (2018). Cash and voucher programmes. Retrieved from http://www.fao.org/emergencies/fao-in-action/cash-and-vouchers/en/
[13] Mark, D., & Dodson, N. (2017). ARES protocol white paper (Draft)

For these reasons, this paper aims to design such a dynamic bond whose size would be appropriate to its intended use and which will dynamically adjust itself based on information available in the system as opposed to voting.

Designing this kind of bond is a type of coordination game where what matters is the need to keep in mind what kinds of information are available in the system for the filtering of proposals to happen. This filtration process happens before any actual voting or human review of proposals takes place, thus the content or quality of the actual proposal is not taken into direct consideration during this process. Our goal is to incentivize agents who submit proposals to self-select which proposals are worth submitting based on clearly defined economic incentives. An agent would thus be economically disincentivized to submit a proposal unless they believe such a proposal is likely to be approved during objective evaluation. Without taking proposal content directly into consideration when filtering proposals, we need to rely on other information available in the system. There are directly measurable metrics that come with each proposal that can be used to evaluate it and calculate an appropriate bond. The most readily available metrics are the amount of funding the project is asking for, the amount of total funding available in the system (supply), and the amount of total funding requested by all pending proposals (demand). This also gives us, by proxy, the information on the proportion of total available funding a given proposal is requesting. Lastly, we can include a reputation score (described in the following section) of a given agent to further tweak the bond amount based on an agent's past behavior and interactions.

From these considerations, it's possible to construct a function that leads to a variable and context-dependent bond based on the following variables:

B *= Bond required for posting a proposal*

$\boldsymbol{B}\,(F, F_L, R, \Omega)$

F *= Amount of funding requested*

$\boldsymbol{F_L}$ *= Funding Request Limit agreed on by stakeholders below which no bond deposit is required. This is aimed to alleviate monetary and administrative burden for small proposals*

R *= Reputation ranges between 0 and 1 as determined by Trust (T) and Reliability (r) values by the Trust and Reputation Model*

$\boldsymbol{\Omega}$ *= DSR = Demand Supply Ratio*

$DSR(F_S, F_D) = \frac{F_D}{F_S}$

F_S *= Supply of funding (amount held by the DAO)*

F_D *= Demand for funding (outstanding amount requested)*

The above set of variables aims to be an exhaustive set of variables which aims to use all the data directly available in the system about the agent, the proposal and the state of the system when an agent submits a proposal. Through the design principles outlined below, an iterative design process can be used to construct a function where the above variables will interact in a way to produce desirable incentives without giving rise to perverse incentives which would make it possible to game the intending functioning of the system. The general design principles here are:

- The bond size should rise as amount of requested funding rises. This is represented by *F* in the numerator of the bond model.
- The bond should however not rise in such a way to make large projects prohibitively expensive to propose. Thus the size of the bond should rise, but it should rise at a decreasing, rather than increasing or linear rate. This is where the deflationary $F^{1/3}$ in the denominator comes from.
- The bond should also rise if there is too much demand for or not enough supply of funding. That is, if proposals are requesting more funding than is available, the bond for new proposals should automatically rise to discourage new submissions. Vice versa, if supply starts exceeding demand, the bond should decrease to incentivize more proposals.
- Theoretically, a bond that increases at a decreasing rate of change could be viewed as a punishment for requesting smaller amounts whose bonds are disproportionately larger. It could also be viewed as an incentive for proposals to be pooled together to request larger funding amounts, as opposed to submitted separately. Here, the supply limitation should work well enough as a guard against large requests exceeding total supply. When the supply of funds is fixed or limited, the supply limitation should act as a disincentive to submitting funding requests that are too large. In other words, it is not always the case as that larger requests should require a proportionately smaller bond. The bond size behavior is very dependent on the supply of funding as well. The incentives above are handled by the $sqrt(1/\Omega)$ in the power of the denominator
- Lastly, Reputation *(R)* which in our reputation model is on a [0,1] scale should work in conjunction with the other variables to make it affordable for agents with low reputation to still propose projects and build reputation, but to disincentivize them from requesting excessive amount of funding given their limited trustworthiness. This is why proposals below a certain amount may not require any bond at all as represented by the piecewise function. *R* in the denominator produces such an incentive for agents with low reputation to either build reputation first by voting and evaluating other projects, or requesting funds $F < F_L$ for which there is no bond required and build reputation that way.

The equation below attempts to design such a bond with all of the above features and incentives in mind. The following section also provides sample visualizations to illustrate how the size of the bond changes as all of the given variables change.

$$B = \begin{cases} 0 & F \text{ Ò} F_L \\ \dfrac{F}{(F^{1/3} * R)^{\tilde{a}^{-}\frac{1}{q}}} & F > F_L \end{cases}$$

Equation 1. Proposal Filtering Bond

The bond mechanism proposed above aims to address the issue of having a fully (or as fully as possible) decentralized funding platform without the reliance on a pre-determined and limited set of human curators, while limiting the attack vectors from malicious actors by economic disincentives. The bond will function as described originally, ie. forfeit to the DAO for funding other charitable projects if the proposal fails, or returned to the proposer if it succeeds (minus transaction fees). Similarly to the original ARES proposal:

> A proposal can fail in only three ways. First, if a voting quorum is not met. Second, if a voting quorum is met but there are more NO votes than YES votes. Third, if a voting quorum is met, there are more YES votes, but not enough funds left in the Fund contract to pay out the proposal. No partial funding will be paid to any proposals. (Mark & Dodson, 2017)

Detailed description of variables based on their roles in the bond equation

- Amount of funding requested in the proposal (*F*)
    - The size of the bond is determined by the size of requested funding,
    - Additionally, there is a Funding Request Limit (F_L) below which no bond is required. This is represented by the piecewise nature of the function. F_L should be a small amount which would incentivize and allow proposals for local community projects with low capital requirements. This amount could be periodically adjusted but should be limited to never exceed a certain value so that a malicious attack cannot remove the bond entirely.

- Reputation of the proposer (*R*)
  - Individuals (or organizations) with low reputation will face a higher proportional bond, while individuals (or organizations) with higher reputation will face a lower bond for submission of a proposal
- Demand for funding and supply of funding (Ω)
  - Proportionally to the total supply of funding (sum of donations held by the Charity DAO), the larger the demand for funding (ie. the more competition for funding), the higher the bond to submit a proposal
  - The smaller the demand, proportionally to the supply, the smaller the bond

The economic incentives in this bond model ensure that individuals with low reputation face higher bonds for submitting proposals, and this bond is also proportional to the size of the funding request as well as the current supply and demand for funding. Thus agents are incentivized to build reputation by starting with smaller projects, for example in their local communities, or to submit verified information and impact evaluations of previous projects to increase their reputation. Furthermore, no single variable (reputation, funding request size, supply/demand) is designed to be more important than any other in determining the size of the bond. A balance of all the inputs is desirable as not to make proposal submission prohibitively expensive. For example, having a high reputation is no guarantee that an agent will face a relatively small bond. If there is a much larger demand for funding than the DAO can currently provide, even an agent with high reputation will be automatically disincentivized to make a large funding request. Given limited supply, if the agent goes ahead and submits a request anyway, they face a high likelihood of losing their bond, which in turn increases supply of funding and reduces bond size for everyone else.

Inversely, having a low reputation does not mean an agent will always face a prohibitive cost of entry into the system. Such an agent can do multiple things: submit small proposals to reduce bond size and gradually build reputation; wait for supply of funding to increase; or have a trusted third party post the bond for the agent in a system of guarantee. The "Visualizing Bond Size" section below aims to build intuitive understanding of this mechanism, and explores in much more detail how each variable affects the bond size.

b) Visualizing Bond Size

For a more intuitive understanding of how the size of the bond varies as the variables (reputation, supply, demand, and funding requested) change, this section offers sample visualizations of the the proposal bond as described in the previous section. The series of visualizations aims to offer a better understanding of how the size of proposal bond changes with each variable in the model. The three main variables in the model are Amount of funding requested (*F*); Reputation ($0 < R < 1$); and Demand Supply Ratio (γ) of funds requested from the DAO and held by the DAO. Additionally, there is a Funding Request Limit (F_L) below which no bond is required. The following three sets of visualizations will each hold two of the three variables fixed to show the impact of the variable in question.

For the corresponding visualizations where relevant variables are fixed, the values of these fixed variables will be:

$F = \$1{,}000{,}000$ $\quad\quad$ $\gamma = 1$

$R = 0.5$ $\quad\quad$ $F_L = \$1{,}000$ $\quad\quad$ *F_L will always be held constant*

Figure 2. Effect of varying Reputation, holding *3* and γ constant

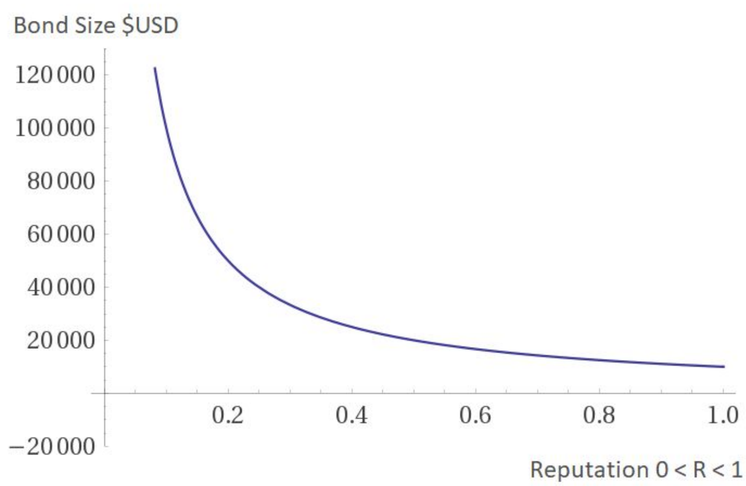

Figure 2 shows the non-linear relationship between bond size and reputation. With very low reputation (R<0.2) the posted bond is very expensive relative to the total amount of funding requested. Therefore, agents who request funds are incentivized to either build reputation first by voting and evaluating other projects, or requesting funds F < FL for which there is no bond required and build reputation that way.

Figure 3. Effects of varying Supply and Demand against Reputation, holding *3* constant

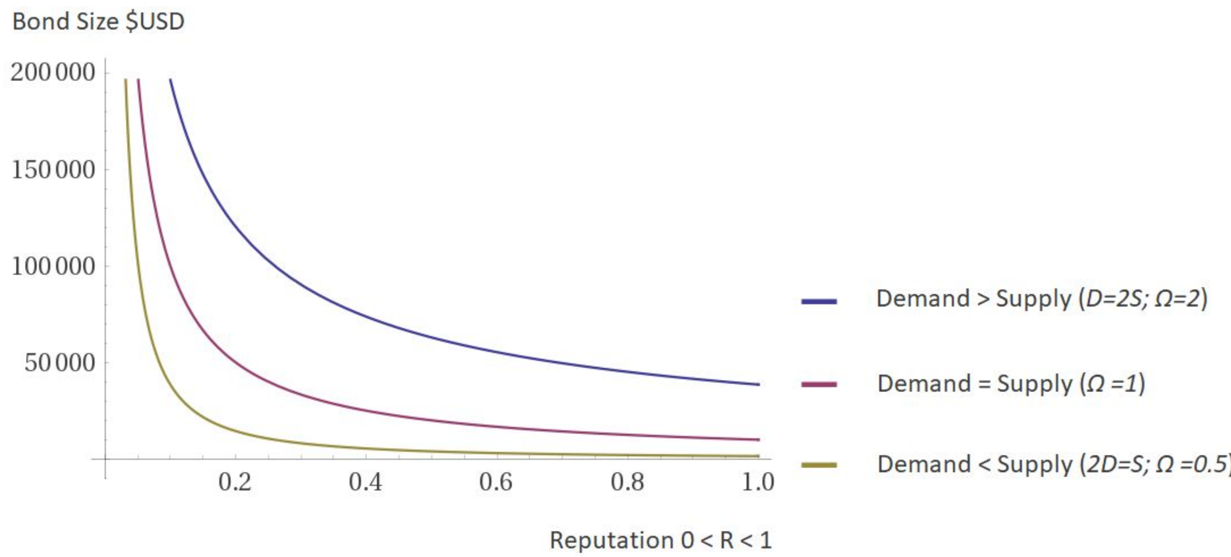

Figure 3 shows the effects of varying supply (funds held by the Charity DAO fund) and demand (outstanding funding requests). Specifically, we see that as demand outpaces supply (by a factor of two in the example above) the required bond sizes for submitting a proposal rise, disproportionately affecting those with lower reputation, thus discouraging new entries into the system and spamming of proposals. The opposite is true for when supply is larger than demand. The required bond becomes much cheaper for everyone and only those with really low reputation (R<0.1) still face a significant barrier to entry.

Figure 4a. and 4b. Effects of varying Requested Funding, holding R and ɣ constant

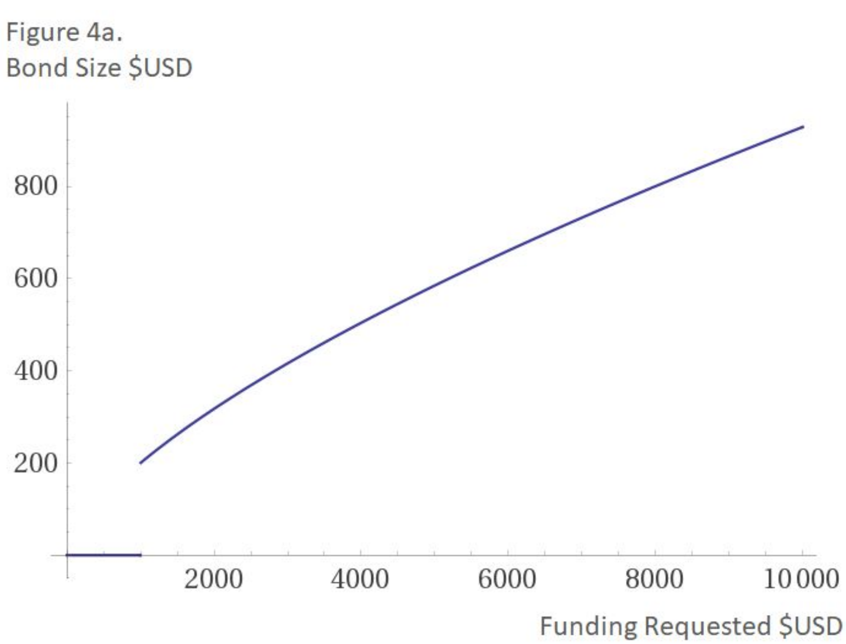

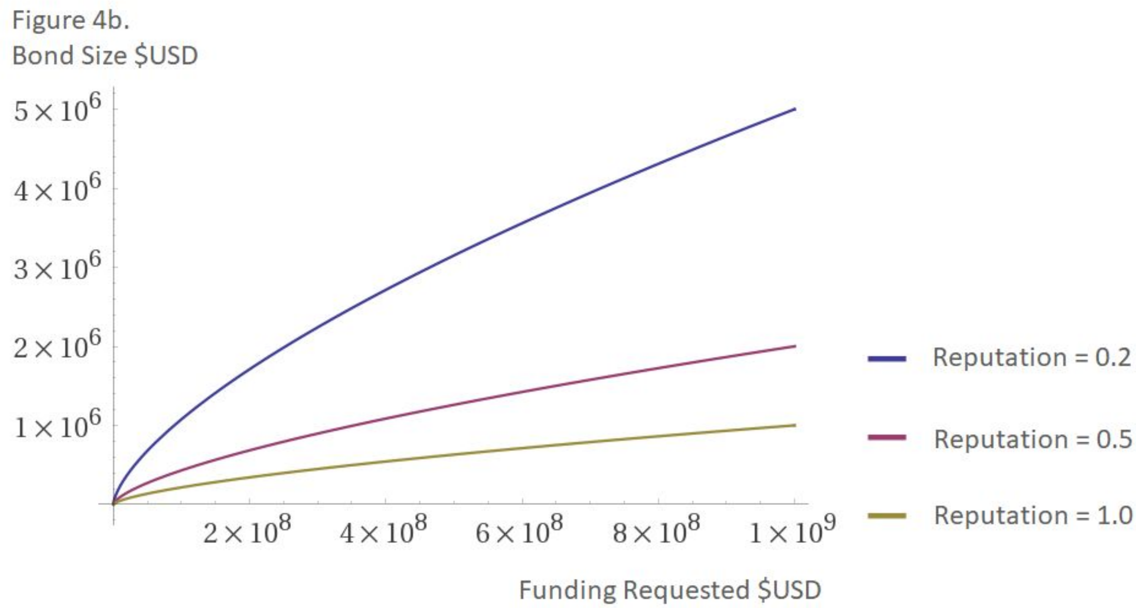

Figures 4a. and 4b. show how bond size varies with different amounts of funding requested at two different scales. Figure 4a. shows the range $0 - $10,000 with reputation held constant at 0.5 to visualize how with FL set to $1,000 the required bond is $0.

The scale of figure 4b. goes all the way to $1,000,000,000 ($1 Billion) to show the potential impact of growth of the funding platform and large funding requests on the bond size. We see that at this scale, even though bond size is designed to increase at a decreasing rate, reputation plays a highly important role. Looking at a funding request of $1B, an agent with baseline reputation of 0.2 would have to deposit a $5,000,000 bond, while an agent with 1.0 reputation deposits a bond of only $500,000. A relatively small amount given the amount of funding requested.

Figure 5.
Effects of varying Requested Funding, with a limited fixed supply of funds, holding *K* constant

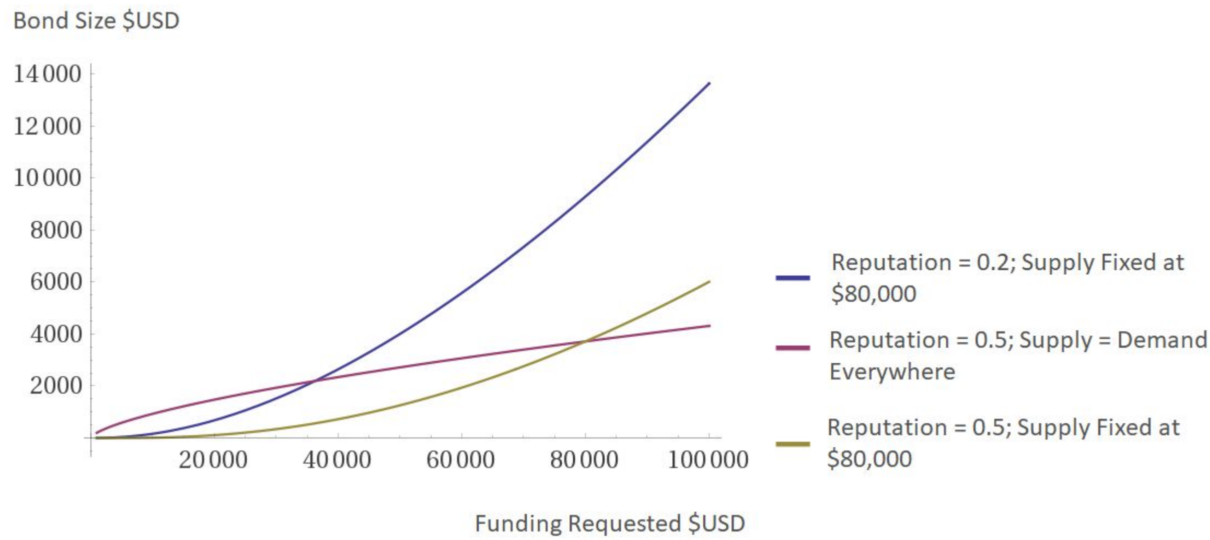

One potential drawback of having the size of the bond increase at a decreasing rate of change as we can see in Figures 4a. and 4b. is that it could be viewed as a punishment for requesting smaller amounts whose bonds are disproportionately larger. It could also be viewed as an incentive for proposals to be pooled together to request larger funding amounts, as opposed to submitted separately. One possible way to handle this issue is to introduce more complexity into the function by taking for example average (mean) reputation and making the bond size linear (or exponential) if one is below average reputation and keeping it at diminishing bond size otherwise. Thus agents with relatively low reputation would face a faster increasing bond size, while agent with relatively high reputation would face a more slowly increasing bond size. However, with the dynamic aspect of the model given to us by the Demand Supply Ratio, this should not be necessary.

The supply limitation should work well enough as a guard against large requests exceeding total supply as seen in Figure 5. above. The figure shows what happens when the supply of funds is fixed or limited as opposed to maintaining a dynamic $\boldsymbol{\gamma}$ ratio which was the assumption in previous models. We are looking at an example of funding requests between $0 and $100,000 and we can see that by fixing the supply of funds at $80,000 we change the bond size function to increase exponentially. The supply limitation thus acts as a disincentive to submitting funding requests that are too large. In other words, it is not always the case as Figure 4. might show that larger requests always require a proportionately smaller bond. The bond size behavior is very dependent on the supply of funding as well.

Additionally, having a smaller bond size for larger requests might function as an incentive for encouraging cooperation among projects with similar goals or in same geographic areas by encouraging them to pool requests together. This might eliminate some overlap between funding requests and make evaluation more efficient for voters as well.

2. Charity DAO Reputation System Framework

One of the key promising features under development in the blockchain ecosystem is the concept of self-sovereign identity and potentially reputation associated with this identity. Instead of relying on third parties (ever present "Log In" with Facebook or Google buttons), users would be able to own their data and manage their identities independent of any platform, making individuals free to switch services at will and be compensated for the value their user data generates. This concept lends itself to the idea of a global blockchain reputation system and score, making it seemingly possible to outsource reputation management to other specialized protocols as described in more detail in the "Charity DAO Integration in the Blockchain Ecosystem" section below. Thus, a "reputation protocol computes and maintains reputation scores of verified entities, globally on the Ethereum blockchain. Any existing and upcoming (...) dApps can fetch the reputation of any entity, with a short breakdown on how the score was calculated".[14] While true, this approach over-simplifies decentralized deputation management. While it should be possible to eventually aggregate individual reputation data into a global value, the data need to come from individual dApps and protocols in the first place. Thus it is necessary for any applications or protocols whose functionality relies on local reputation (within the context of a given application or protocol) to develop and manage their own reputation systems based on data native to the context of the given protocol or application.

For these reasons, even if there eventually exists a global reputation for each user in a blockchain ecosystem, individual protocols and applications such as Charity DAO still need to know what a user is likely to behave like within this particular system. Such specialized use cases have a need for a more granular reputation information about specific sets of behaviors. This section thus describes the desired reputation system features for the Charity DAO platform and outlines a framework for how the individual components of the trust, reliability and reputation values interact. An important consideration to keep in mind with regards to this section, as mentioned before, is that reputation system design, especially when it comes to public, decentralized and permissionless systems is in its infancy. Design of such systems requires deep expertise and knowledge in a variety of fields to ensure it is robust, trustworthy and impossible or at least prohibitively expensive to attack and game. At the time of writing, to the author's knowledge, there is no real, fully decentralized, trustless and public reputation system in live production in existence. It is in the scope of this limited project to only give a theoretical overview of what a potential reputation system for a social impact funding platform could look like, given that "Rating system design is difficult, and [any] rating/reputation system will need to

[14] The Bee Token. (2018). The bee token, the future of the decentralized sharing economy Retrieved from https://s3-us-west-2.amazonaws.com/beenest-public/whitepaper/bee_whitepaper_v3.pdf

evolve along with the network (...) but for starters a simpler system will serve adequately as the initial condition for ongoing refinement".[15]

The proposed Charity DAO Reputation system is based on the ReGreT[16] and FIRE[17] trust and reputation models. Final reputation values is broken down into a Direct Trust value and multiple aspects of reputation: Witness, Neighborhood and System Reputation. The number of impressions, reliability and credibility of all of the above are then combined into a reliability measure which, combined with the trust value serves as a normalized measure of reputation. The Appendix also includes a visual guide/framework for navigating the reputation system. In this model, "**Trust**" refers to a state achieved via authentication, which is the process by which we achieve confidence in a party's identity, plus trust from all direct interactions where parties rate each other. "**Reputation**" refers to historical data which inform donors' and proposers' subjective predictions about an (identified) party's future behavior. Our **Final Reputation Score** (R) then consists of a **Trust** (T) value and a **Reputation Reliability** (r) value, each designed to be between 0 and 1 in this model. Reputation reliability then consists of **Rating Variability** and **Number of Impressions** giving us a score based not only on an agent's activity and the perception of that activity by other members of the system, but also of the agent's reliability based on rating deviation of individual impressions. We use the ReGreT specified functions below to compute the reliability value using a normalized number of impressions and subjective reputation deviations, both between 0 and 1; giving us a reliability value between 0 and 1.

$$Dt(IDB_p^a) = 1\ \bar{\imath} \underset{li \in IDB_p^a}{\times} Ɉ(t,\ t_i)\ \bar{A}|W_i\ \bar{\imath}\ R^t(IDB_p^a)|$$

Equation 2. Reputation Deviation. Reproduced from Sabater, 2002.

"This value goes from 0 to 1. A deviation value near 0 indicates a high variability in the rating values (this is, a low credibility of the reputation value from the subjective reputation deviation point of view) while a value close to 1 indicates a low variability." (Sabater, 2002)

[15] SingularityNET. (2017). SingularityNET: A decentralized, open market and inter-network for AIs Retrieved from https://public.singularitynet.io/whitepaper.pdf
[16] Sabater i Mir, J., & Sierra García, Carles. (n.d.). Trust and Reputation for Agent Societies.
[17] Huynh, T. D., Jennings, N. R., & Shadbolt, N. (2004). FIRE: An integrated trust and reputation model for open multi-agent systems. Frontiers in Artificial Intelligence and Applications, 110, 23-2

$$
Ni(IDB_p^a) = \begin{cases} sin(\frac{\pi}{2 \cdot itm}|IDB_p^a|) & |IDB_p^a| \in [0,\ itm] \\ 1 & \text{otherwise} \end{cases}
$$

Equation 3. Number of Impressions to Calculate Reputation. Reproduced from Sabater 2012.

Sabater’s model takes into account many impressions to form a “judgement” about an agent’s characteristics. Taking into account a large number of impressions as opposed isolated experiences allows the model to increase the degree of reliability of its judgement over time. When this reliability reaches a maximum value, Sabater says “we call [this the] intimate level of interactions (*itm* from now on). From a social point of view, this stage is what we know as a close relation. More experiences will not increase the reliability of our opinion from then on.” (ibid)

As per ReGreT reputation model specifications,

> “All the modules work together to offer a complete trust model based on direct knowledge and reputation. However, the modular approach in the design of the system allows the agent to decide which parts it wants to use. For instance, the agent can decide not to use neighbourhood reputation to calculate a reputation value or rely only on direct trust to calculate the trust on an agent without using the reputation module. Another advantage of this modular approach is the adaptability that the system has to different degrees of knowledge. The system is operative even when the agent is a newcomer and it has an important lack of information. As long as the agent increases its knowledge about the other members of the community and the social relations between them, the system starts using other modules to improve the accuracy of the trust and reputation values.” (Sabater)

Within the modular framework of trust, system, neighborhood and witness reputation, we have four types of agents: donors; donors + voters; donors + delegates; and proposers/receivers, ie. the agents or projects funded by Charity DAO. At the basic level, after each interaction all parties involved rate each other, on a [0, 1] scale. “In this simple version, an Agent’s rating is the distribution of past rating decisions” (SingularityNET, 2017). On top of these baseline ratings, which fall under witness reputation, system reputation consists of counts and variability of voting participation, success rate of proposals and projects voted for, funds donated, percentage of Charity DAO funds requested, and funds donated through forfeited bonds. Neighborhood reputation consists of: In the case one is a delegate, who are the voters voting for the delegate, and what is their reputation, reputation of projects, delegates and other voters in the voter’s network, agents associated with a proposed project, agents associated with source of funds. Finally, witness reputation additionally consists of on the ground evidence from projects funded, proposed or voted for, expertise approval, votes received for a proposal, and success rate of proposed projects. These initial categories are subject to review and change based on system testing once the network is implemented on a small scale.

The second aspect of the Final Reputation Score which is the Trust value reflects a process by which we achieve confidence and trust in an agent's identity and thus consists of verified third party references and identity verified through an identity management or KYC (Know Your Customer) system. Initially, this could be an external centralized KYC provider, however, the preference is to use a decentralized KYC or identity management system as described below in blockchain ecosystem integration.

Additional Features and Considerations of the Charity DAO reputation system:

When it comes to calculating Reputation reliability, voting participation could be weighed on a curve of total participation. This would mean that if there is a large number of proposals, voters cannot be expected to vote on the majority of them, and should not be punished for not doing so. Quality of review should be incentivized over quantity. Agents could additionally be clustered into voting rings based on expertise and assigned to particular kinds of proposal in a specialized division of labor based on knowledge. In that regard, precautions would have to be taken to prevent "regulatory capture" of subcategories of proposals by special interests and elite groups of experts. To prevent this from happening, we could imagine limited incentives for voting outside one's expertise to encourage interdisciplinary review and collaboration, but not to the extent where misinformed agents vote on proposals without adequate knowledge

When delegating votes to delegates, considerations need to be made for how the delegates' reputation is affected by the neighborhood reputation of those voting for them and how the reputation of the proposer is affected. For example, preventing (sybil) attacks of agents with low reputation sabotaging reputation of delegates on purpose. Such risk might even disincentivize a delegate to create excellent proposals, because the better a proposal and the more voters vote for it, the more the delegate's reputation will regress to the mean of all the voters, as the sample of voters grows, especially if the proposal is controversial. A mechanism to address this risk could function in the following way: voters can cast two types of votes, **for** and **against** a given proposal. Reputation of both groups of voters is tracked and evaluated. We should expect a normal distribution of reputations in both the for and against groups with the averages in both groups being close to the system average. If a discrepancy is detected, a review is triggered. For example, if voters voting against a proposal have a significantly lower reputation than those voting for a proposal, we could assume the proposal is being sabotaged by a malicious attack. Inversely, if the reputation of those voting for the proposal is significantly lower, we could assume a coordinated attack is happening trying to misappropriate funds through a sub-par proposal.

Many reputation models incorporate some time decay so more recent ratings are weighted more heavily than ones in the distant past. While this is an excellent approach for many contexts, I would argue that when it comes to a social impact funding platform and funding charitable and development projects, the rating of past projects and interactions should not be discounted given the sensitivity of these kinds of projects.

Additionally, one agent may not necessarily be limited to one vote, but "consecutive ratings from the same source are heavily discounted, whereas ratings from a wide variety of sources are considered more reliable" (SingularityNET, 2017). Votes can also weighed differently based on reputation and an agent may change its vote at a cost which is proportional to an agent's reputation. Moreover, an option could be implemented to vote for more than one delegate or rank delegates in order of preference, this would in turn generate more neighborhood and clustering data based on voter behavior. There should be a minimum reputation value required for one to become a voter. Upon reaching this value a "stake is deposited by the Agent, to be forfeited should its rating (in some dimension) fall below [this] threshold" (ibid). Prediction markets, once in production could be implemented so that "curators are rewarded for predicting another Agent's future rating, via [the market]. If an Agent obtains an artificially inflated rating, then the network should reward the first curators that see through the ruse and penalize it. The same mechanism allows for up-and-coming Agents to more rapidly acquire a good reputation if it's well deserved" [3]

The following figure (next page) gives a simplified visual representation of this reputation framework. The FIRE and ReGreT reputation models, are attached in the Appendix.

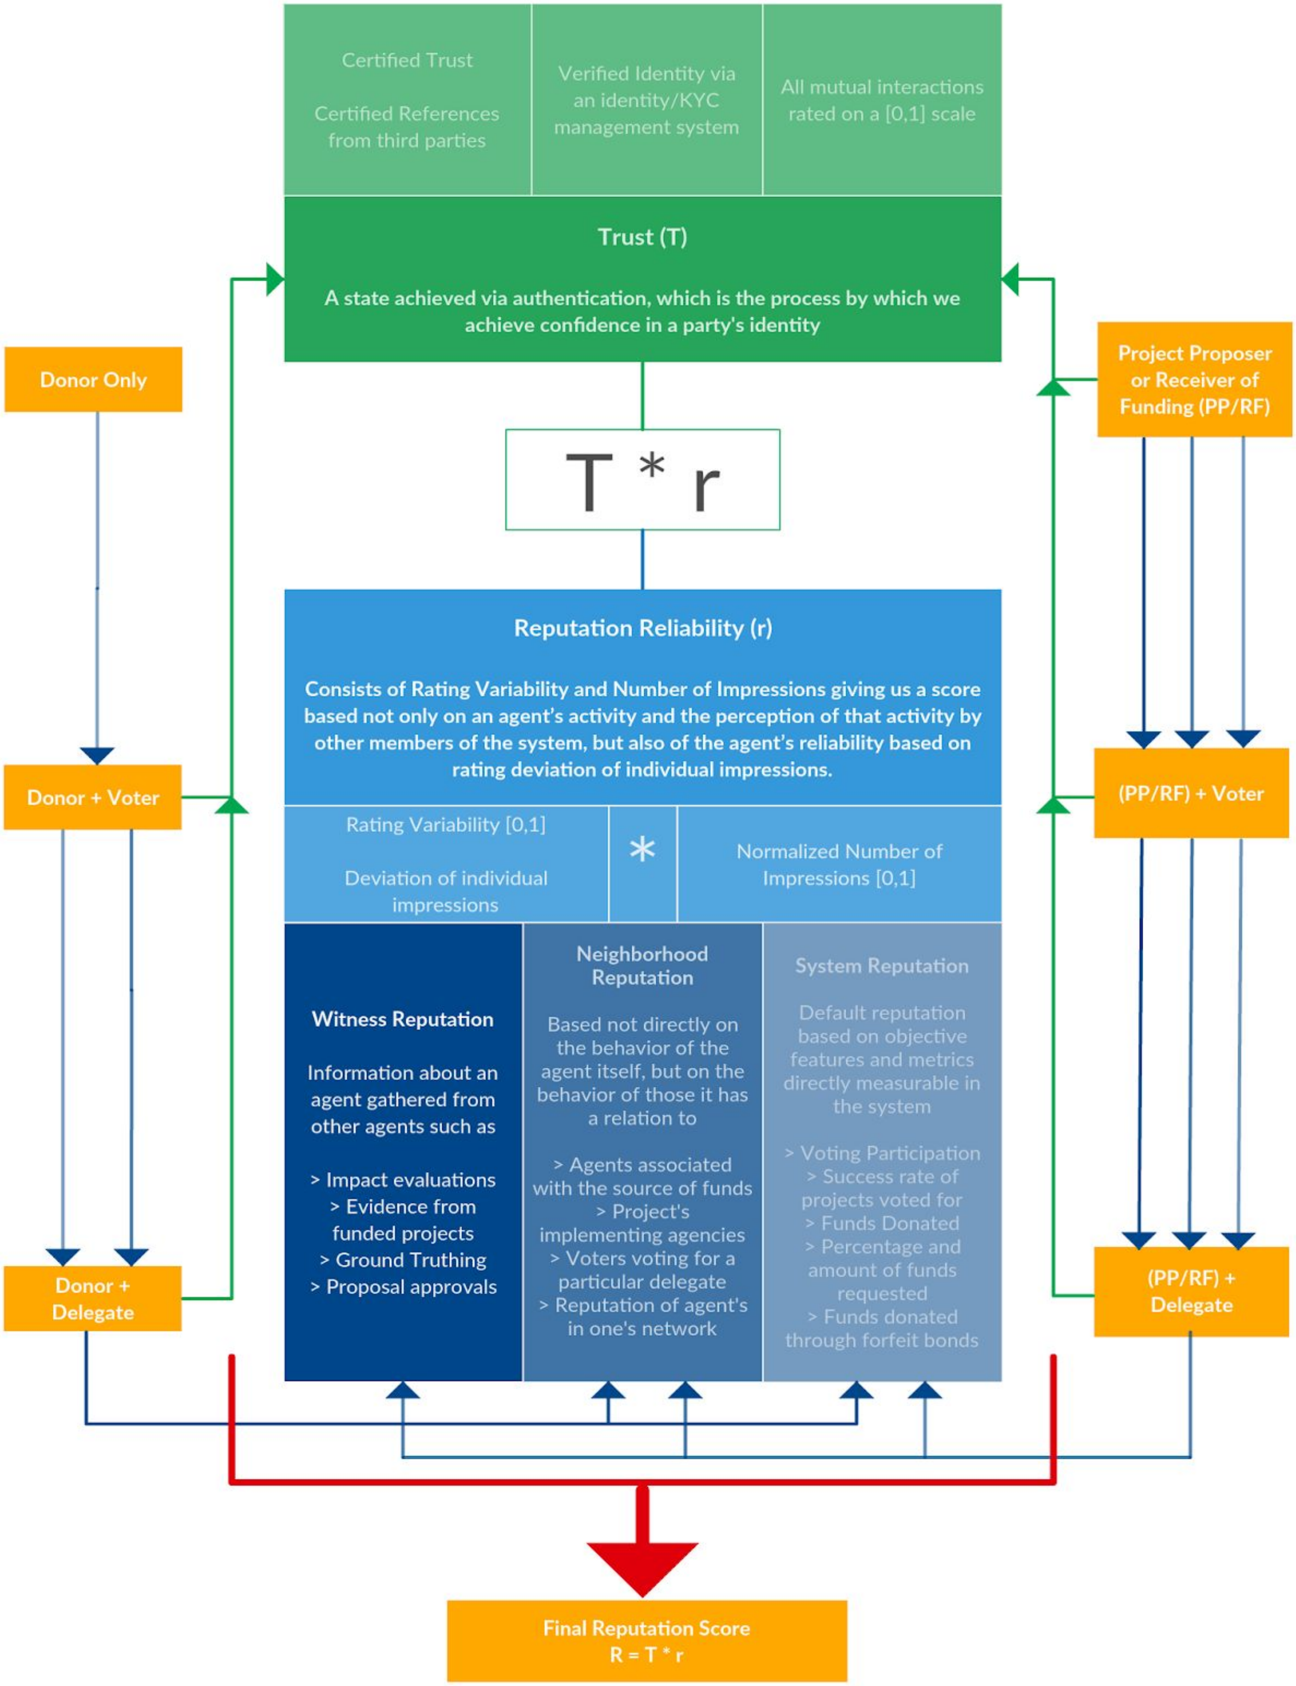

Figure 8. Charity DAO Reputation System Framework

3. Charity DAO Integration in the Blockchain Ecosystem

This section describes how a social-impact funding platform could be integrated with other services within the blockchain and decentralized application ecosystem. However, before outlining this integration, we need to understand a little of how this ecosystem, on a public permissionless blockchain, works when it comes to protocols and applications as compared to the current web 2.0. Currently, value produced by protocols such as TCP/IP, HTTP or proprietary corporate protocols (Google File System) has been largely captured in the form of data ownership in applications on top of these protocols (such as corporate ownership, control and monetization of user data). “This relationship between protocols and applications is reversed in the blockchain application stack. Value concentrates at the shared protocol layer and only a fraction of that value is distributed along at the applications layer”[18]. Instead we have a global network database that any user or application can plug into; a Shared Data Layer[19] “by replicating and storing user data across an open and decentralized network rather than individual applications controlling access to disparate silos of information, we reduce the barriers to entry for new players and create a more vibrant and competitive ecosystem of products and services on top” (ibid). This shared data and application layer with native application and protocol tokens allows us to build modular and decentralized applications which can utilize this shared layer without permission, greatly reducing barriers to entry and giving us the possibility to develop and monetize open-source protocols in a competitive way by “creating internal assets and selling them for development.”[20] How does this affect a social impact funding platform such as the Charity DAO?

In light of the above, we can imagine how a social-impact funding platform could be integrated with other services within the blockchain ecosystem. This means we can plug into existing protocols and decentralized applications which offer service valuable to us without having to negotiate permissions and licensing agreements with the owners or developers of a given service. Simply by using the service’s native token we provide value to the network while allowing us to focus on the specialized task of developing a social impact funding platform and outsourcing as many other tasks as possible to other specialized solutions in the decentralized ecosystem. Beyond Charity DAO, this is a fundamental economic shift that may make possible a global transformation in productivity and cooperation.

---

[18] Monegro, J. (2016). Fat protocols Union Square Ventures. Retrieved from https://www.usv.com/blog/fat-protocols

[19] Monegro, J. (2014,). The shared data layer of the blockchain application stack. Retrieved from http://joel.mn/post/104755282493/the-shared-data-layer-of-the-blockchain

[20] Buterin, V. (2014). *Decentralized protocol monetization and forks* Retrieved from https://blog.ethereum.org/2014/04/30/decentralized-protocol-monetization-and-forks/

A platform such as Charity DAO will have many needs apart from raising money, voting on how these funds get used and distributing them. The platform will need to process these payments, manage the identities of donors and recipients, evaluate and manage projects, hold onto non-volatile funds that don't fluctuate in value or decide and enforce a system of governance. These however are not specialized tasks that every organization should focus funds and energy into developing. These are tasks that are generalizable and common to many organizations. As such we can make use of shared and specialized protocols developed for these purposes by others who have more expertise in these areas. Fundamentally, shared data layers and decentralized permissionless applications allow us to become more specialized and focus on our comparative advantages. Below is a list of blockchain ecosystem projects that outlines how the tasks listed above could be handled by other specialized protocols and applications currently being developed. These are by no means final, rather just a highlight of existing decentralized solutions that an organization such as Charity DAO could utilize.

a) **Identity Management:** uPort
   Self-sovereign key, identity, credential and reputation management

b) **Payment Processing and Fiat Pegging:** Request Network, 0x Protocol, Maker DAO
   Request Network allows for on-chain payment requests in any currency; 0x Protocol can be utilized to convert payment on a decentralized trading platform into a stable asset or coin such as the Dai developed by Maker DAO (See section C below for example implementation of the Title Token using these protocols)

c) **Donation Processing, Impact Investment, Impact Evaluation, and Social Impact Project Management:** Giveth, Alice, ixo Foundation
   All of the above projects are building decentralized social impact platform with some overlap with Charity DAO. Together, these could be used to manage a variety of aspects of decentralized giving, impact evaluation, project monitoring and implementation

d) **Governance and Proposal Management:** Boardroom, Aragon, district0x

   These projects are building solutions for managing decentralized organizational structure, communities, coordination, proposals and voting

e) **Proposal Review, Rating and Management:** SinguarityNET, Colony

   As a platform for open organizations, Colony can be used to manage the incentives and rewards for reviewers and voters, and recruiting individuals with relevant expertise all around the world to review proposals. Protocol such as SingularityNET, an open marketplace for Artificial Intelligence could be used to outsource portions of the review process that can be automated to competing AI algorithms

4. Implementation challenges of decentralized and distributed governance systems

There are additional governance aspects and challenges that need to be considered when implementing the proposed bond mechanism as well as other governance aspects of Charity DAO, and distributed systems in general. One such consideration is what percentage of voters need to approve a proposal for it to pass. Given the nature of Charity DAO as a social impact funding platform, proposals should be aligned to be as beneficial and impactful as possible while generating little controversy, thus I would propose that for a proposal to pass, it needs to have an approval of a 60% supermajority of eligible voters.

Additionally, as the system is ideally based on consensus or supermajority approval, another challenge lies in handling urgent proposals. If a voting period is on the order of days, weeks or even months, say one week for proposal review, one week for voting and one week of grace period for dispute, how is a catastrophic situation handled in the case of a sudden large scale disaster and relief funding? Does the DAO offer a fast-track for proposals? What are the costs and who can submit these? What vulnerabilities, potential for exploitation and attack vectors does this introduce into the system? Given the relatively long time it can take to reach consensus in distributed systems and in many cases the non-finality of these decisions, a decentralized funding platform such as the proposed Charity DAO may not be an ideal system for handling urgent proposals and funding requests, as rapid response requires elements of centralization and executive authority in a decision making process.

Finally, a significant challenge faced by distributed systems reliant on network effects to function properly lies in the initial bootstrapping of the network. This is especially true for a social impact funding platform such as Charity DAO where a coordinated malicious attack during the bootstrapping phase could take over the governance of the network and misappropriate funds for purposes not intended by the platform. I would thus argue that during the conceptual, development, implementation and bootstrapping phases, development of distributed systems and protocols requires a degree of central coordination, executive authority and vetting of initial participants in the system. In the case of Charity DAO, this might mean a limited pilot project, only open to certain entities with a proven record of proposing and funding successful charitable causes. Reputation and rating system design is difficult, especially in distributed systems. All potential pitfalls and attack vectors cannot be foreseen in theoretical research and will only be discovered during implementation and live production. Any such system will need to evolve with the network. Given these considerations, it is important to make clear that different forms of failure are very likely in the bootstrapping phases of such systems and so this phase needs to be approached with utmost care and diligence. For politically sensitive platforms such as social impact funding, as well as for distributed system design, this highlights the importance of starting in small environments where potential damage from failure is limited. From a regulatory perspective, this highlights the importance of regulatory sandboxes and openness to innovation which lets such projects innovate and fail without fear of inappropriate punishment.

C. Hyper-Transparency and Socio-economic Impacts of Blockchain Technology

1. Hyper-Transparency and Title-Token Proof of Concept

Chapter one focused on the technical aspects of economic incentive-driven governance in decentralized systems using mechanism design—specifically, designing a bond for filtering funding proposals on a decentralized social impact funding platform; a reputation framework for this platform, and blockchain ecosystem integration. Chapter two begins with presenting a proof of concept of Hyper-Transparency on the public Ethereum blockchain, a concept developed specifically with social impact and charitable funding platforms in mind. This work is the result of a three-day, four-person collaboration at the ETHWaterloo Hackathon. The two main outcomes here are the introduction of the concept of Hyper-Transparency and its use-cases in certain contexts; and the development of the Title Token Standard. The standard establishes a way to trace individual tokens on the Ethereum blockchain. Each title token minted keeps track of its history on the blockchain by making use of Ethereum's event logs, a title token always keeps track of its originator (the donor). The donor can log onto an app and see everywhere her donation has gone. Use cases go beyond charitable giving and not for profit foundations. Using this concept we can trace the path of every single donation which is publicly visible from the moment it is given, to the moment it is spent, no matter how many intermediary organizations it moved through to get there. We can build a decentralized autonomous organization that can accept payments in any currency, hold onto its funds in a non-volatile form and deploy them globally within seconds while maintaining complete transparency.

One of the features of blockchains is that they are public ledgers, meaning that all transactions can be seen and tracked by anyone. And while this feature allows us to see all of the transactions ever made, it doesn't allow us to track individual "coins" or "tokens." The tokens are entirely fungible and unlike cash or bars of gold which carry serial numbers and can technically (albeit with difficulty) be individually traced, we cannot currently do this on the Ethereum blockchain. When a large number of transactions enter a pool of funds (such as a charity) and then exit as the funds are distributed to individual projects, we lose track of whose money went where. Being able to see transactions is thus not quite good enough, we want to track a donation (the exact donation) down a chain of transactions connecting the donor to the exact good the donation is used for. At first sight, this might appear as an unnecessary complexity in the system or even a gimmick. After all, we can already trace all the transactions and see where funds end up, at least proportionately. But we think creating this level of transparency would make donors feel more comfortable about giving.

That is why our ETHWaterloo team proposes developing the Serializable Title Token standard as an addition to the existing ERC20 and especially the ERC721 non-fungible token standard. We by no means want to see this Title Token standard used by everyone for everything. Privacy has its place on the blockchain and there are valid reasons for why we would want certain communications and transactions to be anonymous and untraceable. Our Title Token Standard fulfills a niche where an extreme level of transparency is preferable. Such as charitable giving or government spending.

Our proof of concept is a Decentralized Autonomous Charity (DAC) that offers "TitleTokens" in return for donations. The DAC is able to split and track the individual donation title tokens into child projects. Donations are accepted in ETH and then converted using the 0x.js API[21] into Sai (a more stable coin, the value of which is maintained at around 1 USD) and held in a pool for a given funding period. Recently, MakerDAO launched the live version of the Dai stable coin, replacing Sai which was running on an Ethereum test network. In a live version of the Charity DAO, Dai would be utilized. The reason for the conversion is to maintain consistency in the currency value the donor chooses to participate and make local currency transactions more seamless. Title Tokens are minted upon receipt of a donation. Each Token is able to be tracked down the supply chain, i.e. the DAC gives all or a portion of Title Tokens to a project, the project gives Title Tokens to a hospital, the hospital then trades the Title Tokens back to the DAC for their representational amount of Dai that they exchange for local money to purchase whatever they need. What is special about TitleTokens is that you can track the entire path of individual donations from start to finish, enabling donors to know exactly where their money went at each step of the process. For the technical implementation of this proof of concept, see https://github.com/ETHDAC and for a demo see http://y2u.be/55r4Xsi6RtM

Use Cases Beyond Charity DAO and Social-Impact

The above concepts of Hyper Transparency and Title Tokens are not strictly limited to charitable use cases. Here is a brief discussion of the wider implications and applied use cases of the above research with regards to decentralized governance and blockchain ecosystems. The applications and use cases of a hyper-transparent token go far beyond charitable giving. Foundations, governments, universities and other non-profit and research organizations that provide grant funding and stipulate conditions could easily see exactly what *their* money is being used for and that it satisfies their conditions. As the transactions would be recorded publicly (or

[21] The 0x protocol is an open, permissionless protocol allowing for ERC20 tokens to be traded on the Ethereum blockchain without the need for a centralized exchange. The 0x API (application programming interface) is a standard relayer API, a clearly defined method of communication between a program and the 0x protocol allowing developers and users easy access and trading across multiple sources of liquidity using the same standard protocol.

in a semi-public ledger or state channel visible to the donor/funder and the recipient), the donor could check in real-time how their money is being spent and whether conditions are satisfied. This could to an extent alleviate the administrative burdens of annual reports, reviews, and audits. While these may still be necessary, they may require significantly less time and workforce.

Tax payments could also be tokenized via our Title Token standard, bringing a never before seen level of transparency to government spending. One could see exactly where *their* money ends up. Such a transparency may not only be satisfactory to the individual taxpayer, but may also enable better analysis and scrutiny of spending giving more insights to policy makers as well. One could go as far as imagining a mechanism for being able to have a limited form of input on how exactly one’s tax payment is allocated.

Title Tokens could be issued for individual items in a supply chain, keeping track of the entire record from manufacture to consumption as millions of items enter and exit factories and warehouses. With individual items in a supply chain tokenized on a blockchain, it would become easier to track down the origins of defective items or disease outbreaks. It would also become even easier to identify waste, inefficiencies and bottlenecks and address them. Such systems are already being prototyped in some cases and the addition of a Title Token standard to these systems could add value.

XXXX Q{zmw¢ t{z!

In conclusion, the main focus of this project is to contribute to the design and creation of a “Charity DAO”, a Decentralized Autonomous Charitable Organization built on top the public Ethereum blockchain. The primary contribution of this project is in the area of governance of decentralized systems in the form of designing a proposal filtering bond and a decentralized reputation system framework. While the goal of this organization is to use blockchain technology to bring more transparency to the development sector by providing community and evidence driven funding and investment via a decentralized and permissionless social impact funding platform, the particular governance models and incentives developed as part of this project are more widely applicable in the area of decentralized governance and blockchain ecosystems. Building on existing research by creating a framework for how reputation could be handled on a decentralized social impact funding platform and adding to the volume of research by providing a practical applied model for handling funding requests in an attempt to limit the role of intermediaries such as curators on a permissionless funding platform are contributions that many projects, protocols and platforms in the decentralized ecosystem may find valuable.

Apart from the economic incentive-driven governance model for decentralized systems—specifically, designing a bond for filtering funding proposals on a decentralized social impact funding platform; a reputation framework for this platform, and blockchain ecosystem integration—the other main output of this project is a proof of concept of Hyper-Transparency on the public Ethereum blockchain, a concept developed specifically with social impact and charitable funding platforms in mind. While not directly connected to the Charity DAO, it could be implemented within its protocol to trace donations and investments. As discussed above, this proof of concept if developed fully could also be applied to many other areas which require transparency, such as public sector funding, taxation and other government projects.

Lastly, an important takeaway is that all of this research is still highly theoretical and when implemented in practice, these concepts and models will not be perfect and will require iterations of improvement. Much more research also remains to be done in the area of decentralized reputation system design. This project only touched the surface of complexity when it comes to designing such a system. Many of the other decentralized protocols and applications an organization such as Charity DAO would rely on, such as identity management, payment processing, voting mechanisms etc., are as well in very early stages of development. And while development of these systems is progressing rapidly, there still remain many technical obstacles to overcome when it come to making these systems usable and efficient for everyday operation. It may still take years before many of the visions outlined in this project come to be realized.

XeX O| |oznt¥!

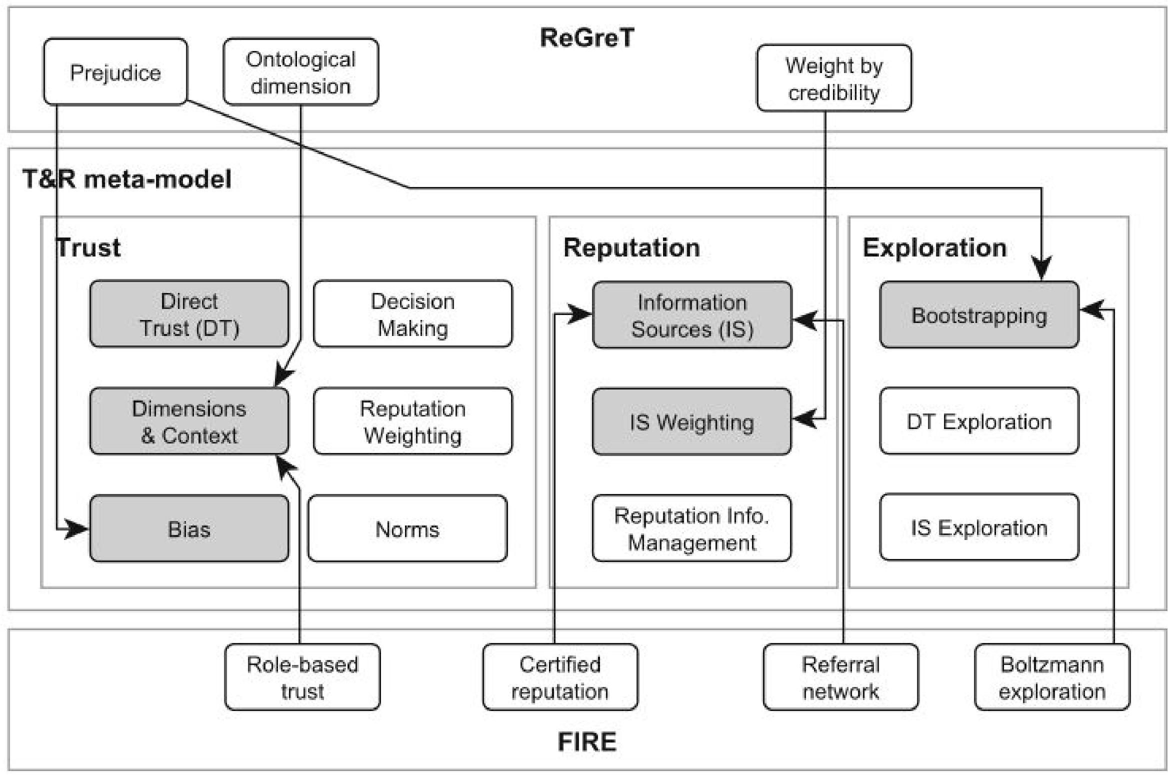

Fig. 11 Mapping of ReGreT and FIRE to the trust and reputation meta-model

Figure 9. Mapping of ReGreT and FIRE to the trust and reputation meta-model (Hoelz & Ralha, 2015)

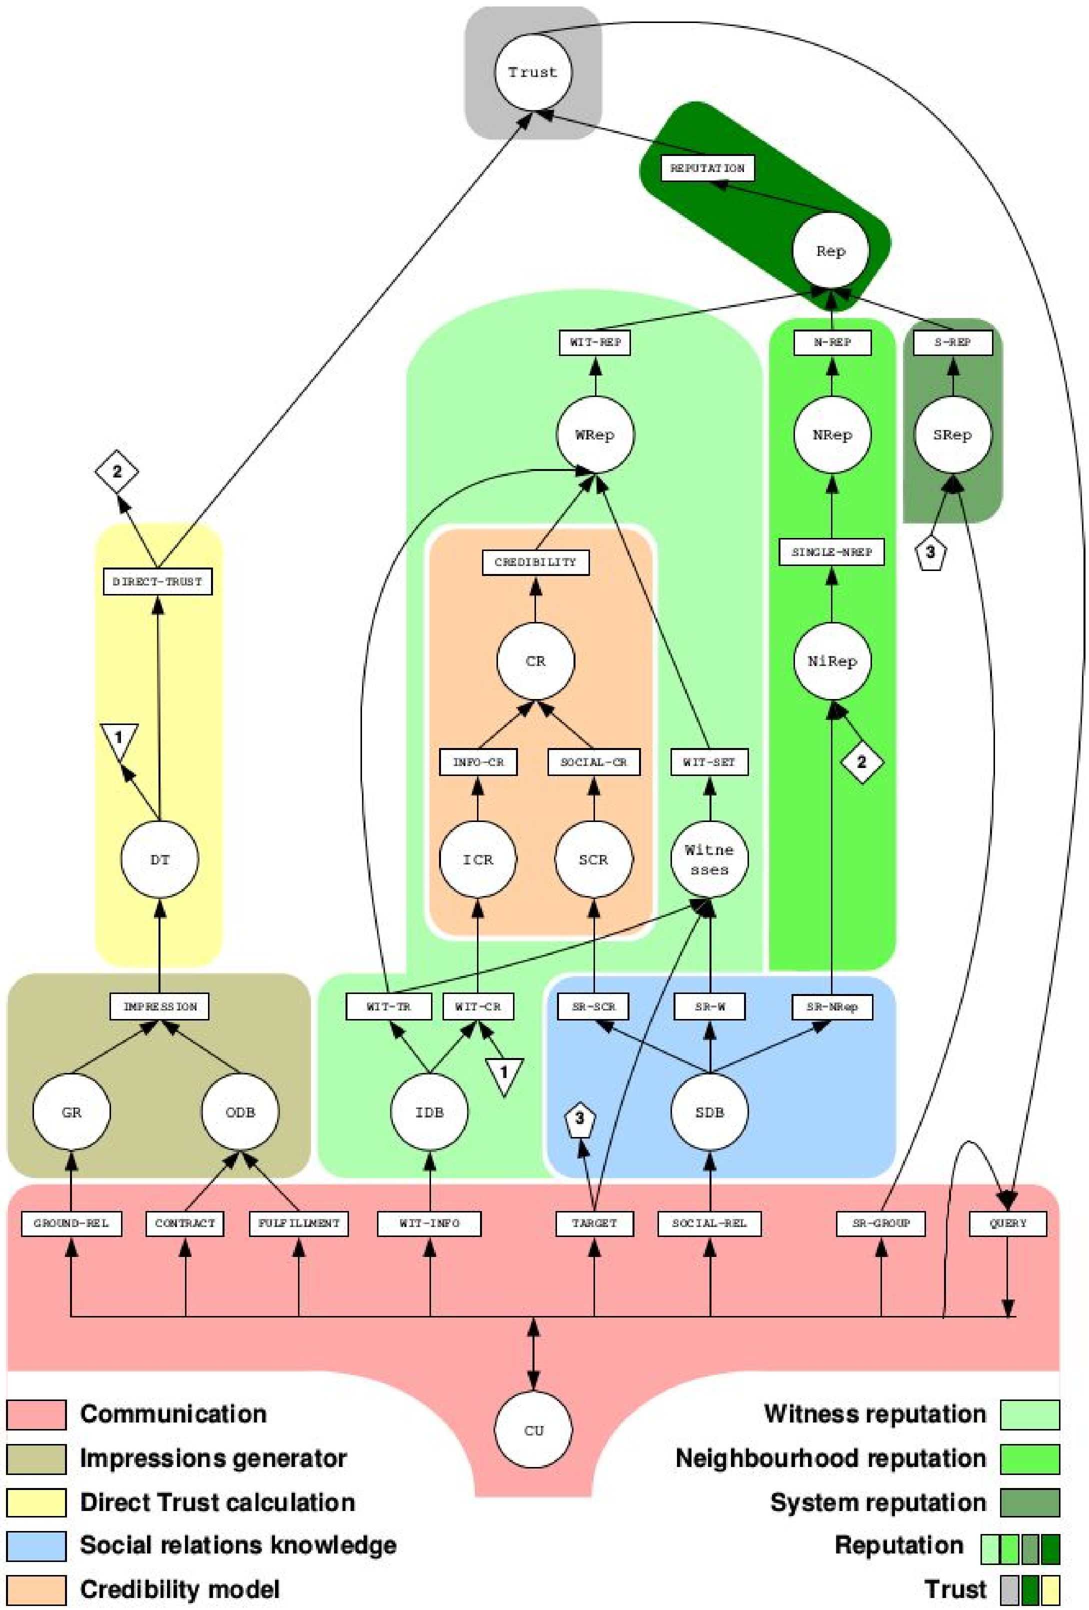

Figure 10. Multi-context specification of the ReGreT system (Sabater, 2002)

e☓ Ptlwt{r~j|s¦ jzn aopo~ozmo !

Akmeemana, C. (2016). Blockchain, cryptoeconomics, and the disintermediation of trust. Retrieved from https://www.linkedin.com/pulse/blockchain-cryptoeconomics-disintermediation-trust-akmeemana-14-000-/

Bauza, L., Lockyer, M., Vrba, T., & Yambao, M. (2017). *The DAC (decentralized autonomous charity)*. ETHWaterloo: Retrieved from https://devpost.com/software/the-dac

Blackbaud Institute. (2016). *5th annual charitable giving report: How nonprofit fundraising performed in 2016.* ().Blackbaud Institute. Retrieved from https://institute.blackbaud.com/wp-content/uploads/2017/02/2016-Charitable-Giving-Report.pdf

Buterin, V. (2014). *Decentralized protocol monetization and forks* Retrieved from https://blog.ethereum.org/2014/04/30/decentralized-protocol-monetization-and-forks/

Buterin, V. (2018). Governance, part 2: Plutocracy is still bad. Retrieved from https://vitalik.ca/general/2018/03/28/plutocracy.html

Flandez, R. (2012). Donors say they would give more if they saw more results. Retrieved from https://www.philanthropy.com/article/Many-Donors-Would-Give-More-if/156463

Food and Agriculture Organization of the United Nations. (2018). Cash and voucher programmes. Retrieved from http://www.fao.org/emergencies/fao-in-action/cash-and-vouchers/en/

From evidence to action: The story of cash transfers and impact evaluation in sub saharan africa (2016). . Oxford: Oxford University Press.10.1093/acprof:oso/9780198769446.001.0001 Retrieved from http://www.oxfordscholarship.com/10.1093/acprof:oso/9780198769446.001.0001/acprof-9780198769446

Generosity for Life, & Indiana University Lilly Family School of Philanthropy. (2018). Charitable profile. Retrieved from http://generosityforlife.org/generosity-data/data-tools/generosity-reports/

Haushofer, J. (2016). The short-term impact of unconditional cash transfers to the poor experimental evidence from kenya. *The Quarterly Journal of Economics, 131*(4), 1973-2042. 10.1093/qje/qjw025

Hoelz, B., & Ralha, W. (2015). Towards a cognitive meta-model for adaptive trust and reputation in open multi-agent systems. Autonomous Agents and Multi-Agent Systems, 29(6), 1125-1156.

Huynh, T. D., Jennings, N. R., & Shadbolt, N. (2004). FIRE: An integrated trust and reputation model for open multi-agent systems. Frontiers in Artificial Intelligence and Applications, 110, 23-27.

Huynh, Trung Dong & R. Jennings, Nicholas & R. Shadbolt, Nigel. (2006). Certified reputation: how an agent can trust a stranger.. 1217-1224. 10.1145/1160633.1160854.

Jeffries, A. (2018). 'Blockchain' is meaningless. Retrieved from https://www.theverge.com/2018/3/7/17091766/blockchain-bitcoin-ethereum-cryptocurrency-meaning

Mark, D., & Dodson, N. (2017). *ARES protocol white paper* (Draft ed.)

Monegro, J. (2014,). The shared data layer of the blockchain application stack. Retrieved from http://joel.mn/post/104755282493/the-shared-data-layer-of-the-blockchain

Monegro, J. (2016). *Fat protocols* Union Square Ventures. Retrieved from https://www.usv.com/blog/fat-protocols

Mougayar, W. (2016). In Buterin V., ebrary I. and MyiLibrary (Eds.), *The business blockchain : Promise, practice, and application of the next internet technology* Hoboken, New Jersey : John Wiley & Sons, Inc.

Outlier Ventures Research. (2017). *Blockchain-enabled convergence: Understanding the web 3.0 economy.* ().Outlier Ventures Research.

Pinyol, I., & Sabater-Mir, J. (2013). Computational trust and reputation models for open multi-agent systems: A review. Artificial Intelligence Review, 40(1), 1-25.

Sabater i Mir, J., & Sierra García, Carles. (n.d.). Trust and Reputation for Agent Societies.

Sabater, Jordi. (2002). REGRET: A reputation model for gregarious societies.

SingularityNET. (2017). *SingularityNET: A decentralized, open market and inter-network for AIs* Retrieved from https://public.singularitynet.io/whitepaper.pdf

The Bee Token. (2018). *The bee token, the future of the decentralized sharing economy* Retrieved from https://s3-us-west-2.amazonaws.com/beenest-public/whitepaper/bee_whitepaper_v3.pdf