Balance: Dynamic Adjustment of Cryptocurrency Deposits

Dominik Harz
Imperial College London
Department of Computing

Lewis Gudgeon
Imperial College London
Department of Computing

Arthur Gervais
Imperial College London
Department of Computing

William J. Knottenbelt
Imperial College London
Department of Computing

ABSTRACT

In cryptoeconomic protocols, financial deposits are fundamental to their security. Protocol designers and their agents face a trade-off when choosing the deposit size. While substantial deposits might increase the protocol security, for example by minimising the impact of adversarial behaviour or risks of currency fluctuations, locked-up capital incurs opportunity costs for agents. Moreover, some protocols require over-collateralization in anticipation of future events and malicious intentions of agents. We present BALANCE, an application-agnostic system that reduces over-collateralization without compromising protocol security. In BALANCE, malicious agents receive no additional utility for cheating once their deposits are reduced. At the same time, honest and rational agents increase their utilities for behaving honestly as their opportunity costs for the locked-up deposits are reduced. BALANCE is a round-based mechanism in which agents need to *continuously* perform desired actions. Rather than treating agents' incentives and behaviour as ancillary, we explicitly model agents' utility, proving the conditions for incentive compatibility. BALANCE improves social welfare given a distribution of honest, rational, and malicious agents. Further, we integrate BALANCE with a cross-chain interoperability protocol, XCLAIM, reducing deposits by 10% while maintaining the same utility for behaving honestly. Our implementation allows any number of agents to be maintained for at most 55,287 gas ($\approx$ USD 0.07) to update the agents' scores, and at a cost of 54,948 gas ($\approx$ USD 0.07) to update the assignment of agents to layers.

1 INTRODUCTION

Capital deposits provide security for cryptoeconomic protocols implemented on decentralised ledgers. Yet, system uncertainties can lead to the need of *over-collateralization*, where more capital than necessary is locked-up to provide a buffer against e.g. currency fluctuations. Moreover, locked capital is expensive: being unable to access funds, agents face an opportunity-cost in the form of the forgone returns that they could have accrued in alternative investments. This is where *cryptoeconomic* protocol design plays an important role in reducing required costs.

Cryptoeconomic protocols (π) combine cryptographic primitives and economic theory. In such protocols, a utility maximising agent chooses which action to perform based on self-interest, ignoring the impact of their action on other agents. However, cheating or other malicious behaviour by self-interested agents may harm the protocol as a whole. Therefore, known protocols typically involve two types of incentives, e.g. [5, 12, 20, 22, 23, 26, 37, 38, 42, 46, 47]: *(i)* payouts motivate agents to act in the best interest of the protocol,

and *(ii)* deposits prevent malicious actions by punish misbehaving agents. Choosing the appropriate deposit, however, is challenging due to the following two sources of uncertainties:

Definition 1 (Private information). A valuation $v_A(\sigma)$ encodes the preference of an agent A for a specific outcome in a protocol depending on an action σ[1]. We define this information as private, i.e. only known to the agent[2].

Definition 2 (Event-dependency). Deposits paid in cryptocurrencies can be subject to external events. Event-dependency is caused by events that occur outside of the underlying blockchain and affects the level of required deposits.

An agent A performing an action in a protocol can *(i)* be intrinsically motivated to harm the protocol or *(ii)* receive an external payment to act maliciously due to e.g. bribing [29]. This preference is encoded in v_A and is unknown to the designer of π. Since such a valuation is *private information*, the required deposit for secure operation becomes unclear. Further, a deposit can be used as insurance for tasks performed outside the underlying ledger, e.g. performing verifiable computations [16, 23, 42], cross-chain assets exchanges [45], or exchanging digital goods [12]. For such usage, the protocol is *event-dependent*, e.g. affected by exchange rate fluctuations. Private information and event-dependency require protocol designers to dynamically calculate the required deposit levels. However, this leads to a dilemma: if the required deposit is too high, the agent is likely to refuse to participate in the protocol. A requesting agent fears that the providing agent has an incentive to cheat if the deposit is too low. In both cases, whether the deposit is set "too low" and "too high" is subjective.

Thus, our central research question arises: *how can deposits in over-collateralised protocols be dynamically adjusted when assuming that agents have private information and there exists a dependency on external events?*

We present BALANCE which provides dynamic adjustment of deposits in over-collateralized protocols, improving agents' financial welfare, without compromising security.

Security intuition. BALANCE allows agents to reduce their capital deposits over a sequence of periodic rounds, while preventing the addition of incentives to act maliciously. In fact, malicious agents that aim to misbehave in a protocol obtain more utility if they do early on, before any reductions in their deposit occurs. Conversely, honest and rational agents receive a higher utility by *consistently* acting in the interest of the protocol. BALANCE achieves

---

[1]For brevity, we write $v_A(\sigma)$ as v_A in the rest of the paper.
[2]See also Chapter 9.4.1 on games with incomplete information in [34].

this property for protocols which feature over-collateralization by reducing deposits to a lower bound. Above the lower bound, the *additional* utility gained from reducing the deposit is *less* than the opportunity cost for locking-up the deposit in the first place. Hence, a malicious agent gains no additional utility from cheating in a later round with a reduced deposit. However, honest and rational agents gain additional utility by reducing the opportunity cost of the locked deposit.

Balance is applicable to protocols that *(i)* are implemented on a decentralised ledger, *(ii)* require agents to over-collateralize, *(iii)* verify interactions with agents through objective specifications ϕ, *(iv)* use a deposit to prevent economically rational agents from performing undesired actions, and *(v)* are subject to at least one of the two identified sources of uncertainty. Suitable examples for integrating Balance are FairSwap [12], TrueBit [42], and XCLAIM [45].

Contributions

- **Introduction of Balance, a mechanism for deposit reduction:** The mechanism maps agents to layers according to their behaviour, where each layer is associated with a different level of deposit. The higher the layer, the lower the required deposit, mitigating over-collateralization. To the best of our knowledge, this is the first mechanism allowing for the dynamic adjustment of cryptocurrency deposits while maintaining the same level of security.

- **A formal analysis of agents' incentives:** we explicitly model agents' utilities with and without Balance. For agents committed to a protocol, we characterise the conditions on incentive compatibility for Balance. Further, we show that given a distribution of agents over honest, rational, and malicious agent types in a protocol, introducing Balance leads to an increase in social welfare. We show that with a *decrease* of the provided deposit, we can achieve a *greater* incentive for economically rational agents to behave honestly.

- **An integration with XCLAIM**: we integrate Balance with XCLAIM [45][3], a protocol that allows two cryptocurrencies to be exchanged without trusted intermediaries, using Solidity[4]. The clear exhibition of both event dependency and private information in XCLAIM makes it a highly suitable Balance use case. We show that Balance reduces the deposit by 13% while maintaining the same level of payoff for complying with the specification of the protocol. The implementation has linear complexity for storing the score of agents depending on the number of layers with a maximum cost of updating a score of a single agent of 55,287 gas ($\approx$ USD 0.07). Further, curating the agents to layers during transition to the next round has constant complexity with 54,948 gas ($\approx$ USD 0.07).

Structure

In Section 2 we provide a definition for cryptoeconomic protocols including contracts and utilities. We provide an overview of Balance and its security assumptions in Section 3. Next, we introduce

the design and functionality of Balance in Section 4. Further, we define a model to prove incentive compatibility and social welfare increase in Section 5. Balance's security is evaluated in Section 6. We present an integration of Balance in XCLAIM in Section 7. Finally, we present related work in Section 8 and conclude in Section 9.

2 BACKGROUND

Cryptoeconomic protocols combine cryptographic primitives and economic theory to create applications on a decentralised ledger. We define a cryptoeconomic protocol as follows.

Definition 3 (Cryptoeconomic protocol π). A cryptoeconomic protocol implements agreements—containing publicly known specifications (i.e. the terms of the agreement), verified through cryptographic primitives in a decentralised ledger—between multiple agents. An agreement encodes an incentive mechanism through payments and deposits, seeking to promote honest behaviour.

The states of a protocol are as follows (cf. Appendix A).

(1) *Committed*: an agent A commits to perform an action σ that fulfils a specification ϕ defined by an agreement. The commitment requires payment of a deposit D.

(2) *Executed*: an agent A performs an action σ. This action either complies with the specification or violates it[5].

(3) *Concluded*: if the action σ complies with the specification, the deposit plus any payments is transferred to the agent. Otherwise, the agent's deposit is destroyed or refunded to another agent that suffered from σ[6].

In this article we use notation as summarised in Appendix E.

2.1 Contracts and agreements

A cryptoeconomic protocol π is implemented through smart contracts. A contract includes agreements $\mathcal{A}$ in the form of:

$$\mathcal{A} = \langle \phi, p, \mathcal{D} \rangle \tag{1}$$

This notion of agreements is similar to [40]. ϕ represents a condition that needs to be fulfilled by an action and can evaluate to either true or false, i.e. $\phi : \sigma \rightarrow \{0, 1\}$. For example, the FairSwap protocol implements such a mechanism by arithmetic circuits [12]. p are the payments made from a receiving agent B to a performing agent A as stated in the agreement. Further, agents need to provide deposits to *commit* to a protocol. $\mathcal{D}$ is a set of deposits $\{D_1, .., D_n\}$ paid by agents to participate in an agreement implemented in a smart contract. The smart contract acts as an escrow of the deposit. If an agent complies with the specification in an agreement, i.e. $\phi(\sigma) = 1$, the deposit is refunded. If an agent performs an action $\phi(\sigma) = 0$ (or no action), the smart contract will not refund the deposit to the agent.

2.2 Action choices

By Definition 3, an agent can only choose to perform an action after it is committed to the protocol. The agent has two action choices.

(1) An agent $A \in P$ performs a *desired action* $\sigma_A \in \Sigma_d$, where Σ_d is defined as the set of all actions that evaluate to true with respect to the specification of an agreement, i.e. $\phi(\Sigma_d) = 1$.

(2) An agent $A \in P$ performs an *undesired action* $\sigma_A \in \Sigma_u$, where Σ_u is the set of all actions that evaluate to false for the specification of an agreement, i.e. $\phi(\Sigma_u) = 0$. This includes also no action $\sigma_A = \emptyset$ as this results in $\phi(\emptyset) = 0$.

2.3 Utility parameters

Each of the two action choices results in a utility for the agent expressed by $u(\sigma)$. The utility of an action depends on five parameters.

(1) Payment p: determines how much an agent B has to pay for requesting an action in an agreement $\mathcal{A}$ and how much the performing agent A receives by fulfilling the specification in $\mathcal{A}$.

(2) Cost c: captures all costs associated with performing an action. This includes, for example, transaction costs and costs for executing an action σ in $\mathcal{A}$.

(3) Deposit D: summarises the deposits D provided by agents in $\mathcal{A}$.

(4) Expected future return $\mathrm{E}[rD]$: describes the opportunity cost for locking deposit D within an agreement that could be used in another protocol to earn an interest[7].

(5) Valuation v: encodes the *private* preference of an agent for an outcome depending on an action σ.

The private valuation expresses that an agent A prefers performing an action $\sigma \in \Sigma_u$. Assume A is given the task to perform a computation in a protocol like TrueBit [42] and provided a deposit D in the agreement $\mathcal{A}$ to perform a computation that fulfils the specification ϕ. If A has a positive valuation for v_A that does *not* fulfil the specification and the provided deposit D is smaller than v_A, the agent prefers the undesired action. We use the value to encode types of agents (cf. Section 3.2). Note that this private valuation might not be economically rational in the sense that it may reflect non-pecuniary value.

2.4 Performing agent utilities

The following are the utilities an agent $A \in P$ *performing* an action[8].

$$u_A(\sigma_A) = \begin{cases} p - c_A - \mathrm{E}[rD_A], & \text{if } \sigma_A \in \Sigma_d \\ v_A - D_A - c_A - \mathrm{E}[rD_A], & \text{if } \sigma_A \in \Sigma_u \end{cases} \quad (2)$$

The agent tries to maximise its utility by choosing an appropriate action. We note two observations from equation (2). First, performing an undesired action becomes a rational choice if $\sigma_{A,i} \in \Sigma_d, \sigma_{A,j} \in \Sigma_u : u(\sigma_{A,i}) < u(\sigma_{A,j})$. In a protocol π without BALANCE agents can enter or leave the game at any time and their

---

[7]We assume the future return rate is stochastic such that the opportunity cost is not known with certainty. We therefore denote the return rate as the expected return rate, $\mathrm{E}[]$.

[8]We omit the valuation v_A for the desired action $\sigma \in \Sigma_d$ as we assume agents do not have an intrinsic or external motivation to comply with the specifications of agreements in π.

past actions are not considered. Hence, the game consists of single, disconnected rounds, i.e. it is a *single-shot game*. Second, executing a desired action depends on the publicly known payout p as well as the agent specific costs c_A caused by performing the action and the expected loss of interest $\mathrm{E}[rD_A]$. An agent will choose this action if it yields higher utility than the other actions[9].

This specification implies the notion of security that we employ in our construction.

Definition 4 (Security under rational agents). A cryptoeconomic protocol is secure if for an economically rational performing agent A, they would not choose an undesired action because the utility of a desired action is greater, i.e. $u(\Sigma_d) > u(\Sigma_u)$.

2.5 Receiving agent utilities

An agent $B \in P$ *receiving* a certain action has a different utility from the performing agent A. This agent has no influence on the performance of the action of agent A if we assume that the payout cannot be changed after commitments are made. We assume that the agent that receives an action is reimbursed with the deposit that is provided by the performing agent. The utilities depending on A's actions are calculated as in (3).

$$u_B(\sigma_A) = \begin{cases} v_B - p - c_B, & \text{if } \sigma_A \in \Sigma_d \\ D_A - v_B - c_B, & \text{if } \sigma_A \in \Sigma_u \end{cases} \quad (3)$$

2.6 Mechanism utilities

Social welfare describes the sum of all agent utilities in a mechanism. If we combine equations (2) and (3), we can calculate the sum of utilities for performing and receiving agents $\{A, B\} \in P$ by considering the actions of *performing* agents A. Assuming that every A is matched with at least one B, we split P into three groups corresponding to whether the performing agents perform desirable, undesirable or no actions. P is therefore comprised of a group Λ_D, for the pairs of agents containing performing agents who perform desirable actions, Λ_U, for undesirable actions, and Λ_N for no actions. Thus summing for pair of agents $x \in P$, where there are X pairs, such that $2X = |P|$:

$$u(\sigma_A) = \begin{cases} \sum_{x=1}^{X}(v_{x,B} - c_{x,B} - c_{x,A} - \mathrm{E}[rD_{x,A}]), & \text{if } A \in \Lambda_D \\ \sum_{x=1}^{X}(v_{x,A} - c_{x,A} - v_{x,B} - c_{x,B}), & \text{if } A \in \Lambda_U \end{cases} \quad (4)$$

Hence, if we take all actions within a certain amount of time, we can specify the utility for a protocol π. Notably, agents acting individually rationally and truthfully after commitment does not imply that the mechanism as a whole maximises social welfare.

3 SYSTEM OVERVIEW

We introduce a mechanism called BALANCE, a verifiable layered curated registry, that aims to address the research question stated in Section 1. The idea is to use BALANCE as an extension to an existing cryptoeconomic protocol, i.e π_{BALANCE}, to control the amount of deposit an agent has to provide to resolve the dilemma of having "too

---

[9]This is a deterministic model where we assume that if an agents "cheats", it loses its deposit. However, one could extend the model to include a probability of the agent being caught if the specification function ϕ includes non-determinism.
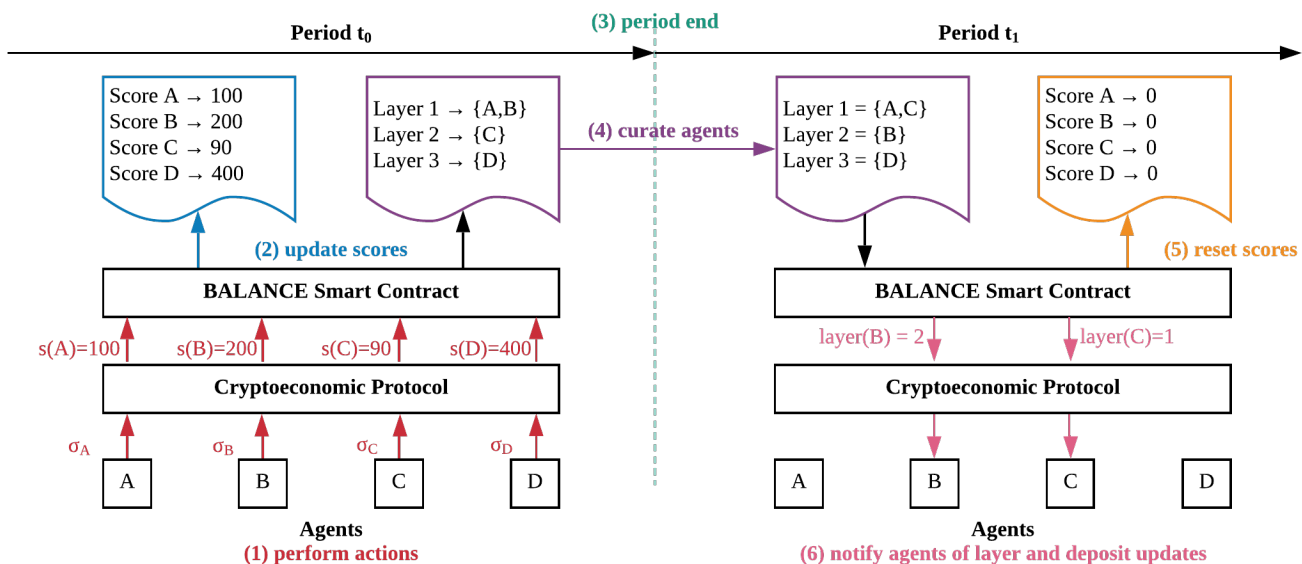
Figure 1: BALANCE is implemented as a smart contract and integrated into an existing cryptoeconomic protocol. Agents A, B, C, **and** D **are assigned to layers in the registry representing their deposit level. In step (1) they are performing actions as part of a cryptoeconomic protocol. By performing these actions, they obtain a score that is forwarded to the BALANCE smart contract. Step (2), is triggered by an agent's actions and updates its scores. Scores reflect a reputation within a round, where the higher the score the better. In step (3), the current round ends. The next interaction of an agent with the cryptoeconomic protocol triggers step (4), the curation of agents to layers. The score of each agent is used to update the mapping of agents to layers, whereby agents can either stay at the same layer, get promoted or demoted by one, or get removed from the registry. Updating the assignment of an agent to a layer is determined by the upper and lower bounds of the layer an agent is currently assigned to. Within that update, step (5) resets the scores of the agent for the next round. Last, in step (6), agents that changed their layer are notified (e.g. via events emitted from a smart contract).**

high" or "too low" deposits. The general intuition of the mechanism is displayed in Figure 1 and is as follows:

- *Commitment*: An agent A has to provide a deposit D_1 to commit to a protocol encoded by a smart contract. For protocols with private information and event-dependency, the deposit is typically set relatively high[10]. An agent is not trusted at any point in time. However, an agent can reduce its deposit by executing actions that are desired by the protocol.
- *Execution*: Agents are tracked by a decentralised registry that stores a ranking of their contributions towards the protocol. Agents are initially assigned to the lowest layer which results in the highest required deposit. When performing desired actions, the agent collects a score within a round.
- *Curation:* If the agent's score is high enough by the end of a round, it moves to the next layer (provided there is a higher layer). The mechanism defines a factor for each layer that is multiplied with the base deposit. This will affect the deposit an agent has to provide. Further, after each round, the score of an agent is reset to retain an incentive to act in

the best interest of the protocol. Thereby, this mechanism introduces a *sequential game*.

We assume that agents will react in a timely manner to resolve any issues with their deposit. Otherwise, agents will move down a layer or be removed from the registry.

3.1 Actors

BALANCE includes three actors with the following roles.

(1) **Performing agent** A: An agent A is committed to π_{BALANCE} and chooses action $\sigma \in \{\Sigma_d, \Sigma_u\}$ evaluating to true or false with respect to the specification ϕ of an agreement $\mathcal{A}$.

(2) **Receiving agent** B: A receiving agent B requested the performance of an action $\sigma \in \Sigma_d$ as expressed by the specification ϕ in the agreement $\mathcal{A}$. The execution of the action is left with an agent A.

(3) **Registry** $\mathcal{R}$: The registry is a smart contract implemented on a ledger. The registry keeps a mapping of scores of performing agents in a round t, a mapping of agents to layers in a round t, and updates the mapping of agents to scores and layers during transition to the next round $t + 1$. The registry has direct access to the result of a performing agent's action $\phi(\sigma) = \{0, 1\}$ to update the agent's score.

Each actor is identified on the ledger using a private/public key pair. We note that agents can take several identities on the

---

[10]For example, in the Dai stable coin, users have to provide at least a 150% deposit to issue Collateralized Debt Positions (CDPs).

blockchain, i.e. Sybil identities are possible. The adjustment of deposit and consequences of performing actions is bound to a public key. An agent A that performs several desired actions to reduce its deposit can have a Sybil identity A'. However, the identity A' is in the view of the registry $\mathcal{R}$ a separate agent and hence will not benefit from the reduced deposit of the desired action sequence by identity A. Further, we show in Section 6.2, that agents in BALANCE cannot profit from Sybil attacks.

3.2 Agent types and adversaries

The performing agent $A \in P$ influences the utility for both performing and receiving agents. Most importantly, the security of a cryptoeconomic protocol depends on a core assumption: adversaries are economically rational and computationally bounded. Considering (2), an agent A chooses to perform a desired action if the utility of the desired action is higher than the utility of the undesired action, i.e. $p - c_A - E[rD_A] > v_A - D_A - c_A - E[rD_A]$. Inspired by the BAR model [1], this notation allows us to express three types of agents, from which we can deduct the adversaries.

Definition 5 (Performing agent types). We categorise performing agents into the following three types.

(1) Type $\mathcal{T}_d$: This type will always perform a desired action as its utility resulting from a desired action is larger than the valuation of an undesired action, i.e. $v_A \ll D_A + p$.

(2) Type $\mathcal{T}_u$: This type will always perform an undesired action as $v_A \gg D_A + p$. By definition, the deposit is never large enough to prevent an undesired action by A.

(3) Type $\mathcal{T}_r$: This type is undecided in the single-shot game setting which decision to take as $v \approx D_A + p$.

By this definition, a cryptoeconomic protocol is only able to ensure that economically rational adversaries, i.e. type $\mathcal{T}_r$, and honest participants perform desired actions. Malicious adversaries as expressed by the type $\mathcal{T}_u$ cannot be economically motivated to perform a desired action. These agent types will perform the undesired action as their valuation for the undesired outcome (from the perspective of the protocol designer) is much greater than the economic damage due to the loss of their deposit.

3.3 System properties

BALANCE aims to resolve the dilemma of balancing deposits in cryptoeconomic protocols with event dependency and private information. We achieve the following properties.

(1) **Auditability**: Performing and receiving agents A and B have public insight into the amount of deposit $\mathcal{D}$ and layer assignment of performing agents $A \rightarrow \mathcal{L}$ (Section 3.4).

(2) **Transparency**: Assuming performing agents A_1 and A_2 provided the same initial deposit D, but are now assigned to different layers such that $D_{A_1} \neq D_{A_2}$, B does not need to choose between A_1 and A_2 as BALANCE ensures that each deposit provides the same utility for A_1 and A_2 to act in the interest of B (Section 5.5).

(3) **Sybil resistance**: A performing agent A cannot gain additional utility from BALANCE by creating Sybil identities A' (Section 6.2).

(4) **Reduction of opportunity costs**: BALANCE reduces the opportunity costs of locked-up deposits $E[rD_A]$ for performing agent A by reducing deposit depending on the layer assignment $A \rightarrow \mathcal{L}$ (Section 5.1).

(5) **Strategy-proofness**: A performing agent A behaves truthfully with respect to its valuation of an outcome v_A independent of its type $\mathcal{T}$ (Section 5.2, 5.4).

(6) **Social welfare increasing**: The joint welfare of performing agent A and receiving agent B is increased with BALANCE (Section 5.6).

3.4 Blockchain model

BALANCE operates as part of π implemented on a blockchain. This blockchain provides a ledger functionality as for example defined in [3, 4, 9, 35]. We assume that the ledger has *finality* as in consensus protocols like [7, 31, 41]. Additionally, we assume given the number of participants n in a consensus protocol, the number of Byzantine faults is $f < n/3$. In cases where a consensus protocol does not provide finality, like Nakamoto consensus [6, 33], we assume that a security parameter k is set such that with high probability the transaction is securely included [15]. Apart from that, we assume that less than 1/3 of miners are malicious [13].

Aside from the ledger, we assume there is an execution environment and a scripting language available that supports the creation of smart contracts to implement agreements $\mathcal{A}$ and the registry $\mathcal{R}$, for example the Ethereum Virtual Machine [43]. We assume that the agreement $\mathcal{A}$ including its specification ϕ, payments p, and deposits $\mathcal{D}$ is implemented by a smart contract SC such that SC $\Leftrightarrow \mathcal{A}$. Further, we require the possibility of creating generic data structures to store the information about agents' scores and layers as part the registry $\mathcal{R}$ and that such information is readable by performing and receiving agents A and B. Moreover, we assume that the cryptographic primitives in the ledger are securely implemented.

> **Property 1 Auditability**
>
> *The deposit $\mathcal{D}$ and mapping of performing agents to layers $A \rightarrow \mathcal{L}$ is publicly available within $\mathcal{R}$.*

We make no further assumptions regarding the blockchain. Miners and adversaries are able to re-order transactions, censor transactions, and read the information from transactions before and after they are included in the blockchain.

4 BALANCE

BALANCE extends an existing protocol π. This extension includes two properties. First, we include a score s into $\mathcal{A}$ such that:

$$\mathcal{A} = \langle \phi, p, s, \mathcal{D} \rangle \tag{5}$$

s is added to the current agent's score if the specification ϕ of the agreement evaluates to true. All details of the agreement are public knowledge. Second, we introduce the registry $\mathcal{R}$ consisting of layers as well as functions to update the score of agents and to curate agents into layers.

4.1 Integrating layers

The layer that an agent is assigned to determines how much deposit they need to provide. We use BALANCE to assign agents to deposit layers, making this a verifiable layered curated registry as illustrated in Figure 6. Intuitively, the higher the score of an agent, the higher the layer an agent is assigned to. For our purposes, layer x is *higher* than layer y if the deposit that an agent needs to provide in layer x is *lower* than the deposit needed in layer y.

We formally define a finite order of layers $\{L_1 \prec \ldots \prec L_\omega\}$, where each layer $L_m \in \mathcal{L}$ has a lower bound l, an upper bound u and a deposit factor $f \in \mathcal{R}^+$. The factors determine the ordering of layers. Each layer maps to a deposit level in order, $\mathcal{L} \to \mathcal{D}$. For instance, L_2 requires deposit D_2.

$$L_i = \langle l, u, f \rangle \tag{6}$$

The set of agents are mapped to the layers, i.e. $P \to \mathcal{L}$. Each agent is only mapped to a single layer at any point in time. The *layer* function returns the layer an agent is currently assigned to. We define the layer(A) as the current layer of an agent A, layer(A)+1 as the higher layer with a lower deposit factor, and layer(A) − 1 as the lower layer with a higher deposit factor.

The layers are used to calculate the deposit an agent needs to provide. We define a base deposit D_{base} (for the lowest layer) and a factor f that are used to calculate D_m, where $m \in \mathcal{L}$. Further, the deposit D_A of an agent A is determined by the base deposit and the factor of the layer the agent is currently assigned to.

$$D_m = D_{\text{base}} f m \tag{7}$$

$$D_A = D_{\text{base}} f_{\text{layer}(A)} \tag{8}$$

Factors are ordered similarly to the layers. Formally, factors are a finite order $\{f_\omega < \ldots < f_1\}$ where the factor corresponding to the highest layer L_ω is the *smallest*.

> **Property 4 Reduction of opportunity costs**
> BALANCE reduces deposit lock-up.

4.2 Updating scores

An agent that just committed to an agreement starts at the lowest layer with the highest deposit factor. An agent can move up or down the layers in the registry depending on its score s_A.

Agents can increase their score in any round by performing actions. The score for an action is determined from the agreement, as defined in (5). The *update* function takes an action by an agent and updates the agent's score.

$$\text{update} : \sigma_A \to s_A \tag{9}$$

4.3 Curating agents

We define a function *curate*. This function takes as input all agents in an agreement $\mathcal{A}$ and results in a new assignment of agents to layers depending on their scores. If an agent's score is higher than the upper bound of its currently assigned layer, the agent progresses to the higher layer. If an agent's score is lower than the lower bound of its currently assigned layer, the agent falls back to the lower layer.

$$\text{curate}(A) = \begin{cases} \text{layer}(A), & \text{if } l_{\text{layer}(A)} \leq s_A \leq u_{\text{layer}(A)} \\ \text{layer}(A) + 1, & \text{if } u_{\text{layer}(A)} < s_A \\ \text{layer}(A) - 1, & \text{if } s_A < l_{\text{layer}(A)} \end{cases} \tag{10}$$

As layers are a finite order, an agent cannot progress above the highest layer or below the lowest layer. An agent can leave the registry by de-registering. This detail is left open to the actual underlying cryptoeconomic protocol. We note that most protocols require the agent to perform the action they initially committed to. If an agent leaves prematurely, i.e. before the concluded phase, the agent's deposit is usually destroyed. However, an agent can certainly participate in multiple agreements in the same contract. The scores of these multiple agreements are added to the overall score of the agent. An agent already in the highest layer can in the best case remain there. An agent that is currently in the lowest layer and misses its lower bound is excluded from the registry. Further, any agent that performs an undesired action in any agreement part of a smart contract is immediately excluded from the registry and the deposit is destroyed or refunded to another agent.

4.4 Determining a time period

According to our blockchain assumptions, BALANCE requires a minimum number of blocks to consider transactions confirmed depending on the security parameter k [15]. Further, BALANCE requires a common period t over all its agreements $\mathcal{A}$. We can determine a *minimum* time period by considering the agreement that requires the longest time even in the case of multiple chains (e.g. XCLAIM as in Section 7). However, the time between blocks Δ is not constant on all chains (e.g. [11]). In the case of a single chain it is sufficient to express the time as a number of blocks, i.e. t depends solely on the security parameter k and the number of transaction n_{tx} to execute an agreement such that $t = kn_{\text{tx}}$.

If BALANCE is used in the multi-chain case, we need to include a buffer b_Δ to account for deviations in Δ. The time period in seconds for *one* chain is expressed by $t_{\text{chain}} = b_\Delta(n_{\text{tx}}k\Delta)$. We can then calculate the total time by summing the time periods for all involved chains within an agreement. Last, the time period can be expressed as a number of blocks by dividing the total time by the average Δ of the chain on which BALANCE is implemented.

However, we note that the time period also determines how often agents need to perform desired actions to remain or progress in the layers. Hence, protocol designers need to consider how often agents execute actions within the cryptoeconomic protocol.

5 INCENTIVE COMPATIBILITY

In our incentive proofs, we make the following assumptions. First, agents are always able to perform the action that they choose to perform. Second, for clarity of exposition, we make the simplification that if an agent performs a single desired action at a certain time, they transition to the higher layer. This is a simplification because in the full specification, agents would need to perform a number of actions in order to change layer. Third, we assume that BALANCE immediately detects an undesired action. In this case, the agent takes on a new identity in the protocol and is able to continue to interact with the protocol starting again from the lowest layer.

Forth, agents can reduce their deposit upon entering the next layer. Fifth, we assume that in the long run, the market for registries such as BALANCE would be perfectly competitive, in the sense that the price (or payments) made in the system will be equal to the marginal cost of each transaction. Sixth, we assume that performing an undesired action is free of direct costs[11].

5.1 Action choice

The performing agent commits to the protocol by providing a deposit D_m. The deposits required by each layer are such that $D_1 > D_2 > ... > D_\omega$: BALANCE rewards the performance of desired actions by letting agents move to the next layer, thereby decreasing the amount of deposit they need to lock up. The utility that the agent gets depends on their choice of action from $\{\sigma_A \in \Sigma_d, \sigma_A \in \Sigma_u, \sigma_A \in \Sigma_\emptyset\}$. Given that the performing agent is committed to the protocol, a utility maximizing agent will never choose to do nothing ($\sigma_A \in \Sigma_\emptyset$) as they would receive a positive utility from committing either a desired or an undesired action (depending on their valuation, v_A).

Figure 2 presents the payoffs that the receiving agent would receive by deciding to perform either a desired or undesired action. The performing agent receives a payoff after each move. If we consider a single-shot game at each round, the agent has to decide between the resulting utility of two actions, i.e. the desired and undesired action.

We can express the condition for choosing a desired action for an agent A at layer $m \in [1, \omega]$ with the following equations:

$$v - c_A - \mathrm{E}[rD_m] - D_m < p - c_A - \mathrm{E}[rD_m] \quad (11)$$

This amounts to requiring that

$$v < D_m + p \quad (12)$$

Theorem 1. Considering a single-shot game, if the decision to perform a desired action holds true for the highest layer L_ω then it will also hold true for all previous layers.

PROOF. If the agent in the highest layer decides to perform a desired action, equation (12) holds true. As $D_\omega < ... < D_1$, the utility of the agent at layer L_ω is higher than at the previous layers. Hence, a rational agent deciding to perform a desired action at layer L_ω makes the same decision in the previous layers. This also allows it to reach the highest layer. □

5.2 Incentive compatibility for types $\mathcal{T}_d$ and $\mathcal{T}_u$

We now extend the analysis to consider the intertemporal payoffs resulting from the agents' decisions. If an agent always does a desired action, i.e. an agent of type $\mathcal{T}_d$, they get a payoff of $p - c_A - \mathrm{E}[rD_1]$ in the first period. In the second period, invoking a discount factor $0 < \delta < 1$ and a real rate of interest r, agents receive a payoff of $(\frac{\delta}{1+r})(p - c_A - \mathrm{E}[rD_2])$ in today's terms. The payoff is discounted reflecting the fact that future payoffs are less valuable than present ones, because a payoff today would receive a rate of return r. This continues until the agent is in the highest layer, L_ω. In the highest layer, agents would receive a payoff of $(\frac{\delta}{1+r})^{\omega-1}(p - c_A - \mathrm{E}[rD_\omega])$ in today's terms. Thus summing to infinity yields the following

---
[11]For example, an agent might go offline and not perform any action in $\mathcal{A}$.

payoff $u(\sigma_d)$ for agents who always performs a desired action each round:

$$u(\sigma_d) = \sum_{t=0}^{\omega-1} (\frac{\delta}{1+r})^t (p_A - c_A - \mathrm{E}[rD_{t+1}])$$
$$+ \sum_{t=\omega}^{\infty} (\frac{\delta}{1+r})^t (p_A - c_A - \mathrm{E}[rD_\omega]) \quad (13)$$

Where convergence of the latter term requires that $|(\frac{\delta}{1+r})| < 1$.

If we assume that an agent is of type $\mathcal{T}_u$, the agent will indefinitely perform an undesired action: at the lowest layer L_1, the utility of the undesired action is higher than the utility of the desired action. Hence, this agent's utility is as follows.

$$u(\sigma_d) = \sum_{t=0}^{\infty} (\frac{\delta}{1+r})^t (v - c_A - \mathrm{E}[rD_1] - D_1) \quad (14)$$

Assuming that $|(\frac{\delta}{1+r})| < 1$, such that the series of payoff converges, this can be expressed as:

$$u(\sigma_d) = \frac{(1+r)(v - c_A - \mathrm{E}[rD_1]) - D_1}{1+r-\delta} \quad (15)$$

Theorem 2. Agents of types $\mathcal{T}_d$ and $\mathcal{T}_u$ in BALANCE act individually rational and truthful to their valuation.

PROOF. By Definition 5, agents that are committed to the protocol cannot increase their utility by hiding their true valuations: the protocol is strategy-proof. The utility of each agent in an infinite time horizon is given by equations (13) and (14), respectively. Thus, BALANCE is incentive compatible for types $\mathcal{T}_d$ and $\mathcal{T}_u$. □

Property 5 Strategy-proofness
BALANCE offers strategy-proofness for types $\mathcal{T}_d$ and $\mathcal{T}_u$.

5.3 Decision boundary for type $\mathcal{T}_r$

The rational agent of type $\mathcal{T}_r$ is indifferent between performing an undesired or desired action in the highest layer L. Hence, consider the following strategy s_{cycling} for an agent:

Definition 6 (Layer-Cycling). The agent performs a desired action in the first $\omega - 1$ layers but always commits an undesired action once in layer ω.

We index these cycles with γ, such $0 \leq \gamma \leq \infty$. For instance, the first iteration of this strategy is denoted by $\gamma = 0$; the next iteration by $\gamma = 1$. If we consider a one-shot game, whether the agent optimally commits a desired or undesired action is determined by (12). However, this does not consider the fact that once the agent performs an undesired action they need to work back up through the layers, with greater opportunity costs for their locked-up deposit, by performing $\omega - 1$ desired actions to reach the highest layer again. Thus, we provide a framework which captures the intertemporal trade-off that agents face in BALANCE in committing an undesired action once in the highest layer.

For each cycle complete γ starting at $t = 0$, the payoff of such a strategy can be expressed as

$$u_A = v - c_A - D_\omega - \mathrm{E}[rD_\omega] + \sum_{t=1}^{\omega-1} (\frac{\delta}{1+r})^t \{p - c_A - \mathrm{E}[rD_t]\} \quad (16)$$
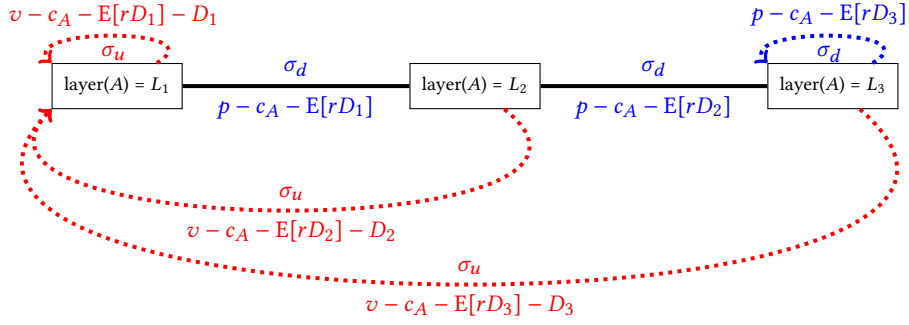
Figure 2: Layer transitions for agent A in a simplified 3 layer system. The red dashed lines correspond to an undesirable action(σ_u); the black solid branches correspond to a desirable action (σ_d). The agent starts in L_1. If the agent chooses σ_d they are moved into layer L_2, receiving payoff $p - c_A - \mathrm{E}[rD_1]$, thus reducing the agent's deposit to D_2. In contrast, if the agent chooses σ_u in L_1 they receive payoff $v - c_A - \mathrm{E}[rD_1] - D_1$ and are removed from the registry. Assuming agents can rejoin the protocol, this is equivalent to being returned to L_1. This process is repeated analogously in L_2 and L_3.

where D_ω denotes the deposit made in the highest layer L_ω. The corresponding expression for infinite cycles is provided in equation (27) in the appendix.

Now consider the stream of payoffs to an agent who always commits a desired action once in the highest layer. We can group the stream of payoffs into sections of size ω, to correspond to the number of rounds required for a player playing **Layer-Cycling** to complete one full cycle. Thus the equivalent utility for performing desired actions for a single cycle γ is as follows:

$$u_A = \sum_{t=0}^{\omega-1} (\frac{\delta}{1+r})^t \{p - c_A - \mathrm{E}[rD_\omega]\} \qquad (17)$$

Similarly, the corresponding expression for infinite cycles is provided in equation (25) in the appendix. Lemma 2 in the appendix shows says that comparing the payoffs of a equations (17) and (16) over infinite cycle repetitions $\gamma \to \infty$ is the just the same as comparing those equations for a single cycle. Thus, equating (27) and (25) provides us with an expression for the value of v at the boundary.

$$v = p + D_\omega + \sum_{t=1}^{\omega-1} (\frac{\delta}{1+r})^t (\mathrm{E}[rD_t] - \mathrm{E}[rD_\omega]) \qquad (18)$$

Equation (18) captures several salient aspects of BALANCE. Firstly, the boundary valuation depends on the levels of p and D_ω. An increase in the payment that the agent receives or an increase in the deposit in the final layer serve to increase the value of v required such that the agent is indifferent between committing a desired or undesired action. The minimum valuation v also depends on the different in opportunity costs (forgone returns) on deposits, with the opportunity cost of funds in lower layers contributing more: $(\frac{\delta}{1+r})^t$ declines as t increases.

As detailed in Appendix C.1, explicitly invoking deposit ratios between the layers, such that $D_\omega = f_\omega D_{\text{base}}$ and $D_t = f_t D_{\text{base}}$, and setting $p = 0$, enables (18) to be rewritten as

$$v = f_\omega D_{\text{base}} + \sum_{t=1}^{\omega-1} \left(\frac{\delta}{1+r}\right)^t \left(\mathrm{E}[rD_{\text{base}}(f_t - f_\omega)]\right) \qquad (19)$$

Note that on the assumption that an agent can advance a single layer m per time period t, to make this relationship between layers

and time explicit we swap m for t. Rearranging this expression for f_ω, and assuming a linear relationship between the smallest factor f_ω and the other factors f_m (see Appendix C.1) enables the indifference boundary to be plotted - see Figure 3.

5.4 Incentive compatibility region for type $\mathcal{T}_r$

We can characterise the incentive compatibility region for the rational agent as follows.

Theorem 3. Provided $f_\omega > f_{\bar\omega}$, performing σ_d at the highest layer (and not **Layer-Cycling**) is incentive compatible for type $\mathcal{T}_r$.

PROOF. We set D_{base} as a relative value where $D_{\text{base}} = 1$. Assuming a linear relationship between the smallest factor f_ω and the other factors f_t, as detailed in Appendix C.2 (35), enables us to express the boundary condition for incentive compatibility for agents of type $\mathcal{T}_r$ $f_{\bar\omega}$ as follows:

$$f_{\bar\omega} = \frac{v\omega - v - \sum_{t=1}^{\omega-1} \left(\frac{\delta}{1+r}\right)^t \mathrm{E}[r(f_1\omega - f_1 t)]}{\omega - 1 - r\sum_{t=1}^{\omega-1} \left(\frac{\delta}{1+r}\right)^t \mathrm{E}[r(\omega - t)]} \qquad (20)$$

Thus provided $f_\omega > f_{\bar\omega}$, incentive compatibility holds for $\mathcal{T}_r$. □

> **Property 5 Strategy-proofness**
> *BALANCE offers strategy-proofness for type $\mathcal{T}_r$.*

5.5 Security by incentives

BALANCE seeks to be transparent to receiving agents B. B does not have to choose between performing agents A based on their layer assignment, but selects A as without BALANCE.

Theorem 4. Given an initial deposit D_1 for a performing agent A of type $\mathcal{T}_r$ and BALANCE with ω layers, A cannot decrease its utility for performing a desired action and does not have a higher utility for an undesired action assuming $f_\omega > f_{\bar\omega}$.

PROOF. If $A \to L_m$ at a point $t = 0$, we consider two cases: First, A can consider an *infinite* horizon. By Theorem 3, we know that if $f_\omega > f_{\bar\omega}$, the utility for performing a desired action is greater

Figure 3: Boundary of f_ω for five initial factors of f_1 = $\{1.0, 1.25, 1.5, 1.75, 2.0\}$. The five functions describe the decision boundary for an economically rational agent to decide between a desired and undesired action given a number of time steps t. If f_ω is set above the decision bound for a given f_1, the rational agent performs a desired action. Assuming $r = 0.05$, $\delta = 0.9$, $\omega = 30$.

than the undesired action for A starting at L_1. By the infinite time horizon the starting layer L_m becomes irrelevant. Hence, A has the same utility for performing a desired action at layer L_m as at L_1.

Second, by Theorem 1, if agent A performs a desired action at the highest layer, A also performs desired action at previous layers. Further, the utility at the lowest layer L_1 is smaller than the utility at the subsequent layers. If A only considers a single time step, and A has performed a desired action at layer L_1, A can only increase its utility by performing a desired action in the next layer. □

Property 2 Transparency

BALANCE is transparent to a receiving agent B.

5.6 Comparison and social welfare

Next, we show that adding BALANCE increases social welfare for agents who perform desired actions.

Theorem 5. For a distribution of agent types $p(x)$, where $x \in \{\mathcal{T}_d, \mathcal{T}_u, \mathcal{T}_r\}$, adding BALANCE increases social welfare.

PROOF. As in (4), if performing agents perform desired actions, the utility is equal to $\sum_{x=1}^{X}(v_{x,B} - c_{x,B} - c_{x,A} - E[rD_{x,A}])$, where x is a pair of agents in the set of pairs X. without BALANCE, each pair is required to have deposit D_1 each period. Over a cycle length $t = \omega$, the total utility for all pairs of agents over the cycle is given by

$$\sum_{t=0}^{\omega-1}\sum_{x=1}^{X}(\delta/(1+r))^t((v_{x,B} - c_{x,B} - c_{x,A} - E[rD_{x,1}])) \quad (21)$$

In contrast, with BALANCE, total welfare is given by

$$\sum_{t=0}^{\omega-1}\sum_{x=1}^{X}(\delta/(1+r))^t((v_{x,B} - c_{x,B} - c_{x,A} - E[rD_{x,t+1}])) \quad (22)$$

For agents performing desired actions, BALANCE improves welfare when (22) is greater than (21). Yet, since in all but $t = 0$, the opportunity cost term ($[rD_{x,t+1}]$) is smaller in the latter equation, with BALANCE permitting agents to post smaller deposits at higher

layers, (22) must be the larger. Therefore, BALANCE improves social welfare for agents performing desired actions. For agents who always perform undesired actions, social welfare is the same with BALANCE as without it. □

Property 6 Social welfare increasing

BALANCE is social welfare increasing.

6 SECURITY ARGUMENTS

We discuss a range of attack strategies against BALANCE and how to mitigate them.

6.1 Single-shot attack

In single-shot attack, $s_{single-shot}$, the only objective of agent A is to attack B through committing an undesired action. A would then consider how to execute such an attack with least cost. In particular, A may be able to reduce the cost of attack by playing a modification of $s_{cycling}$ where A progresses through the layers and performs the undesired action at the highest layer. We show that if f_ω is set above a bound, A would not gain additional utility by waiting until it is in the highest layer before performing the attack.

Lemma 1. If A plays a strategy $s_{single-shot}$ it cannot gain additional utility given that the f_ω is set above $f_{\bar\omega}$.

PROOF. This strategy has two implications. First, A performs $\omega - 1$ desired actions that contribute to the social welfare of the protocol. Second, A increases its utility at $t = \omega$ for the undesired action with being punished with the smaller D_ω instead of D_1. However, for in each of $\omega - 1$ rounds, A has incurred an opportunity cost. As shown in Lemma 2 in Appendix C.1, $s_{single-shot}$ results in the same utility expression and boundary for f_ω as $s_{cycling}$. Hence, if f_ω is set to be greater than $f_{\bar\omega}$, A does not gain additional utility from playing $s_{single-shot}$. A should commit the undesired action at D_1 if A intends to perform a undesired action. Otherwise, A should perform the desired action. BALANCE provides in the single-shot attack the same security as protocols without BALANCE, however, offers desired agent types to reduce their deposit over time. □

6.2 Reputation boosting

An agent A could create Sybil identities to request performance of actions and fulfil them itself. This results in improving the agent's score and reducing its deposit. We denote this $s_{boosting}$, a common attack vector in reputation like systems [19, 21].

Theorem 6. Reputation boosting $s_{boosting}$ is not rational in BALANCE if the cost of this strategy, $c(s_{boosting})$, is higher than the expected saving from the reduction of deposit.

PROOF. We consider a scenario where A can take the role of a receiving agent B. Further, we assume that the cost for requesting a service within an agreement $\mathcal{A}$ by $c(s_{boosting}) = c_B$. Without loss of generality, we assume that an instance of BALANCE includes two layers, L_1 and L_2. A performs a desired action by fulfilling its own request (per role B). In the next step, A can either execute a desired or an undesired action. In either case, the increase of utility is determined by the reduction in opportunity cost at each layer,

i.e. $(\frac{\delta}{1+r})E[rD_2] - E[rD_1]$. A sufficient condition for s_{boosting} to not be a rational strategy is:

$$c_B > E[rD_1] - \left(\frac{\delta}{1+r}\right)E[rD_2] \tag{23}$$

Therefore, we prevent s_{boosting} by limiting the maximum deposit that can be provided by a single identity of A to D_{max}. If A wants to exceed D_{max}, it needs to commit to an agreement with an additional identity A'. □

The *commit* stage of cryptoeconomic protocols typically requires an on-chain transaction which incurs a cost, e.g. [37, 38]. If requesting is *free* in a protocol, boosting would be rational. Thus, to prevent *griefing* by B, cryptoeconomic protocols typically require a small amount of deposit [12] as friction. Since B also has an expected loss of interest on this deposit, this incurs a cost expressed by $E[rD_B]$.

> **Property 3 Sybil resistance**
> BALANCE *is resistant to the creation of Sybil identities.*

6.3 Action delay

In strategy s_{delay}, A, who has already performed sufficient desired actions to move up a layer, delays a desired action from t until $t + 1$. In the new round, with a reset score, the delayed action would count towards the new score. If there are significantly more requesting agents B than performing agents A, a backlog of requests would occur. We propose two remedies. First, A could be required to perform the actions within a time-limit enforced by the cryptoeconomic protocol (e.g. HTLCs). Second, BALANCE could be modified to use *dynamic time rounds*. Dynamic time could be based on the number of *active* performing agents A and the number of requests by receiving agents B. For instance, the next round could be started after a certain fraction (or all) of the performing agents A could have at least fulfilled one request by B.

6.4 Competition and cooperation

Performing agents compete to fulfil agreements, to receive payments and reduce their deposit requirements. However, it is possible that the number of requests is insufficient for A to maintain its assignment to the high layers of BALANCE. In such cases, two strategies for A are as follows.

First, A can be *non-cooperative* and play a strategy s_{compete} in which A (i) tries to fulfil requests as fast as possible by e.g. detecting request transaction early through memory pool sniffing and (ii) actively prevent other performing agents from fulfilling requests by e.g. executing eclipse attacks [17].

Second, A can be *cooperative*, executing a strategy $s_{\text{cooperate}}$ to collaborate with other agents. A could form groups with shared interests through e.g. multi-signature wallets with other agents. A set of agents can act as a single entity to reduce their pooled deposit and fulfil agreements as a whole. This concept is similarly applied to layer-one protocols as so called mining pools. Lewenberg et al. formulate an analysis for cooperative games in the Bitcoin mining network [24]. A cooperative game for cryptoeconomic protocols would consider the number of performing agents A, their overall deposit $\mathcal{D}$, and a likelihood of having open requests by B.

In both cases, BALANCE does not change the possible strategies in the cryptoeconomic protocol. Agents can play either strategies with and without BALANCE being present.

7 APPLICATION: XCLAIM

XCLAIM is a generalised protocol that allows exchange of cryptocurrencies across different blockchains without a trusted intermediary [45]. The protocol employs third parties that provide a deposit to be eligible to participate in the protocol. Further, XCLAIM reduces the cost of atomic cross-currency swaps in comparison to atomic-swaps realised with hashed time-lock contracts (HTLC). We choose XCLAIM as an example of how to apply BALANCE as it is subject to exchange rate risk (event-dependency) and heterogeneous valuations of individual cryptocurrencies (private information).

7.1 XCLAIM protocol

XCLAIM is divided into three sub-protocols. First, the *issue protocol* enables creation of cryptocurrency-based assets (CbAs) from a backing chain onto an issuing chain. Second, in the *swap protocol*, senders and receivers on the issuing chain can exchange the issued CbAs. Finally, in the *redeem protocol*, a redeemer exchanges his CbA for the original coin on the *backing* chain. For example, Alice could issue Bitcoin-backed tokens on the Ethereum chain. Alice could then send her Bitcoin-backed tokens to Bob on Ethereum in exchange for Ether. Bob can take his Bitcoin-backed tokens to receive the equivalent amount of Bitcoins on Bitcoin. XCLAIM includes six roles.

(1) *CbA Requester*: Creates a backed token $i(b)$ on an issuing chain from locking a coin b on the backing chain.
(2) *CbA Sender*: Sends a backed token $i(b)$ on the issuing chain.
(3) *CbA Receiver*: Receives a backed token $i(b)$ on the issuing chain.
(4) *CbA Redeemer*: Destroys the backed token $i(b)$ on the issuing chain to redeem the coin b on the backing chain.
(5) *CbA Backing Vault (vault)*: a non-trusted third party enabling the issue protocol and executing the redeem protocol. The CbA Requester locks her coins b with the vault while the CbA Redeemer upon destruction of $i(b)$ redeems b from the vault.
(6) *Issuing Smart Contract* (`iSC`): Implements the sub-protocols on the issuing chain.

The vault only becomes active during the redeem sub-protocol as any actions in the issuing sub-protocol are executed by the CbA Requester and the `iSC`. To ensure correct behaviour, the vault has to provide a deposit on the issuing chain which allows the issuing and redeeming of CbAs. The `iSC` serves as an escrow for the deposit of the vault. The deposit is a promise by the vault to release the backing tokens b when a valid redeem request is made by destroying $i(b)$. Under the assumption of an economically rational vault, the vault releases b as otherwise its deposit is refunded to the CbA Redeemer. Further, the `iSC` verifies correct execution of the issue and redeem process by verifying transactions occurring in the backing chain on the issuing chain via a chain relay.

Vault agreements

Roles: A mapping from XCLAIM to BALANCE roles.

(1) **Performing agent** A: Vault
(2) **Receiving agent** B: CbA Requester and CbA Redeemer
(3) **Registry** $\mathcal{R}$: Integrated in the iSC

Events: An update to the exchange rate at time t changes the required ideal deposit $\lceil D_A \rceil$ and $\lfloor D_A \rfloor$. The current deposit D_A might fall below $\lceil D_A \rceil$ or $\lfloor D_A \rfloor$.

Private information: A might value a currency higher such that even if $D_A > \lfloor D_A \rfloor$, A behaves economically rational when performing an undesired action.

$\mathcal{A}_1$: **Issue commitment**

ϕ **True**, if A provides a deposit $D_A \geq \lceil D_A \rceil$.
    False, if $D_A < \lceil D_A \rceil$.
p If ϕ of $\mathcal{A}_{1,4}$ is true, A receives p_A paid by B in its issue request.
$\mathcal{D}$ A's deposit D_A is refunded to A if ϕ of $\mathcal{A}_1$ is false.

$\mathcal{A}_2$: **Redeem request** (*Precondition:* ϕ of $\mathcal{A}_1$ is true.)

ϕ **True**, if A fulfils a *valid* redeem request by B within a time w.
    False, if A fails to provide a proof of executing B's request within w.
p If ϕ of $\mathcal{A}_2$ is true, A receives payment $p_A{}^{a}$.
$\mathcal{D}$ A's deposit D_A is transferred to B if ϕ of $\mathcal{A}_2$ is false.

$\mathcal{A}_3$: **Liquidation** (*Precondition:* ϕ of $\mathcal{A}_1$ is true.)
Trigger: Exchange rate update provided by oracle O.

ϕ **True**, if D_A is above $\lfloor D_A \rfloor$.
    False, if D_A is below $\lfloor D_A \rfloor$.
p There are no payments for performing this action.
$\mathcal{D}$ A's deposit D_A is used to perform a redeem request as stated in $\mathcal{A}_1$.

$\mathcal{A}_4$: **Buffered deposit** (*Precondition:* ϕ of $\mathcal{A}_1$ is true.)
Trigger: Exchange rate update provided by oracle O.

ϕ **True**, if D_A is above $\lceil D_A \rceil$.
    False, if D_A is below $\lceil D_A \rceil$.
p There are no payments for performing this action. A cannot participate in the issue protocol while ϕ is false.
$\mathcal{D}$ There is no impact on the deposit.

---

[a] XCLAIM does not define how payments are made during redeem (cf. Section VII-F [45]). We assume that a fee is already paid during issue.

Parameters of BALANCE

Assumptions:

r 0.05. A earns 5% of its deposit by participating in other protocols.
δ 0.9. A discounts future income by 10%.

Time constraints:

t 528 blocks (on Ethereum)

Actions: Actions depend on the agreements and have an associated score.

$\mathcal{A}_1$ **Desired** σ_d: A provides $D_A \geq \lceil D_A \rceil$.
   Undesired σ_u: A provides $D_A < \lceil D_A \rceil$.
   Score s: 0. A should not receive a reduction of deposit for adding deposit in small quantities to reduce its overall deposit. By setting $s = 0$ for $\mathcal{A}_1$ we prevent such behaviour.

$\mathcal{A}_2$ **Desired** σ_d: A fulfils B's redeem request within w.
   Undesired σ_u: A fails to execute B's redeem request within w.
   Score s: > 0. A receives an additional incentive to execute redeem requests and can in return reduce its deposit.

$\mathcal{A}_3$ **Desired** σ_d: If $D_A < \lfloor D_A \rfloor$, A adds additional deposit D'_A such that $D_A + D'_A \geq \lfloor D_A \rfloor$.
   Undesired σ_u: A does not add deposit new deposit or an insufficient amount such that $D_A + D'_A < \lfloor D_A \rfloor$.
   Score s: 0. A should not receive a reward for being temporary below the liquidation bound.

$\mathcal{A}_4$ **Desired** σ_d: If $D_A < \lceil D_A \rceil$, A adds additional deposit D'_A such that $D_A + D'_A \geq \lceil D_A \rceil$.
   Undesired σ_u: A does not add deposit new deposit or an insufficient amount such that $D_A + D'_A < \lceil D_A \rceil$.
   Score s: > 0 should receive an incentive for staying above the ideal deposit, however, this might result in purposely falling under the ideal deposit.

Layers: The layer factors express the boundary for $\lceil D_A \rceil$. We leave $\lfloor D_A \rfloor$ as suggested by XCLAIM at 1.05.

ω 12. Based on considering a time window of 24 hours and the minimum time of 528 blocks.
f_1 2.06. Based on our analysis of exchange rate and order books.
f_ω 1.85. Based on applying f_1 to (36) and considering the exchange rate and order book analysis.

Figure 4: An overview of the agreements of a vault in XCLAIM and parameters for an integration of BALANCE into XCLAIM.

7.2 Deposit dilemma for XCLAIM vaults

Vaults in XCLAIM provide a deposit in the currency of the issuing chain to insure against a risk on the backing chain. Exchange rate fluctuations are external events that affect the security assumption of a vault fulfilling a redeem request. Sudden drops in the exchange rate results in insecure protocol states where the provided deposit is less than the value of the coin b on the backing chain. Hence, rational vaults have an incentive to refuse redeem requests even under the threat of having their deposit taken away. Further, vaults have a private valuation for the outcome of the redeem protocol. If a vault values the backing coin b higher than the required deposit, the vault has an incentive to not fulfil the redeem request. XCLAIM mitigates exchange rate fluctuations and, implicitly, the potential detrimental effects of private information through staged deposits.

(1) In *secure operation*, a CbA Requester can issue new CbAs as D_A is greater than an the ideal deposit $\lceil D_A \rceil$.
(2) In the *buffered deposit* stage, $D_A < \lceil D_A \rceil$ but D_A is still above a lower bound deposit $\lfloor D_A \rfloor$. In this stage, a CbA Requester cannot issue new CbAs based on this vaults deposit. The vault can provide additional deposit to increase its deposit rate back to the ideal rate.

(3) In the *liquidation* stage, a vault's deposit is automatically liquidated if it falls below $\lfloor D_A \rfloor$. The deposit rate is *dynamically* adjusted using an oracle O importing the exchange rate.

7.3 Application of BALANCE

We integrate BALANCE with XCLAIM by defining the agreements of vaults. These include their specification, payments, and deposits as well as the score. Further, we define the length of the time period t in which an agent can perform actions. Last, we set the initial deposit factor f_1 and lower bound f_ω. Our definitions are described in Figure 4. We leave the quantification of scores for actions and determining lower and upper bounds of layers as future work. We argue that BALANCE does not affect the security arguments of XCLAIM in Appendix D.3.

Determining a minimum time. XCLAIM assumes a minimum number of blocks to consider transactions confirmed depending on the security parameter k [15]. As outlined in Section 4.4, we determine the period by considering the agreement with the most transactions. Agreements $\mathcal{A}_{1,3,4}$ depend on one issuing chain transactions while $\mathcal{A}_2$ depends on one backing chain and two issuing chain transactions. We use a buffer of 2 such that the time $t = 2(k_b \Delta_b + 2k_i \Delta_i)$. For the case where Bitcoin is the backing and Ethereum the issuing chain, we assume $k_{\text{ETH}} = 12$, $\Delta_{\text{ETH}} = 15$ s,

$k_{\text{BTC}} = 6$, $\Delta_{\text{BTC}} = 10$ min. This results in $t \approx 132$ min or $t \approx 528$ blocks on Ethereum.

Determining scores. Within XCLAIM, the issue commitment $\mathcal{A}_1$ is equivalent to the commit stage. We choose $s = 0$ for $\mathcal{A}_1$ to prevent agents from splitting up their deposit into smaller parts. Further, we exclude a reward for $\mathcal{A}_3$ resolving $D_A < \lfloor D_A \rfloor$ as ideally D_A should always be greater than $\lceil D_A \rceil$. We define a positive score for the desired action in $\mathcal{A}_2$. Further, we argue that having a positive score in $\mathcal{A}_4$ provides an additional incentive for A to re-balance its deposit in case of exchange rate fluctuations. In turn, this improves the chance of CbA Requesters being able to participate in the issue sub-protocol.

Determining a number of layers. We determine the number of layers ω by considering which time period a performing agent A is accounting for in its decision to decide between Σ_d and Σ_u. We assume that A considers a 24 hour time period. Next, we calculate ω by dividing 24 hours by the time representation of t. Based on $t = 528$ blocks, we set $\omega = 12$.

Determining factors. XCLAIM suggests deposit stages of 2.0 for $\lceil D_A \rceil$ and 1.05 for $\lfloor D_A \rfloor$. We leave $\lfloor D_A \rfloor$ as suggested. We focus on reducing $\lceil D_A \rceil$ for a concrete implementation of XCLAIM for Bitcoin and Ethereum. To verify the suggested deposit $\lceil D_A \rceil$, we perform an analysis of the daily exchange rate fluctuations and private valuations by parsing order books of exchanges (cf. Appendix D) for Bitcoin and Ethereum. Based on our analysis, and assuming $r = 0.05$ and $\delta = 0.9$ as well as a linear relation between layer factors, we set $f_1 = 2.06$ and $f_\omega = 1.85$. Thereby, we achieve a reduction of 10% of the deposit.

7.4 Implementation

We implemented BALANCE for Ethereum in Solidity v0.5.0[12]. The implementation consists of around 170 lines of code. The core functions of BALANCE are *update* and *curate* defined by (9) and (10). The cost experiments are conducted under the assumption that all agents in the current registry change their layer (i.e. maximum updates in the smart contract)[13]. We conducted experiments with up to ten layers. The `update` function amounts to 55,287 gas costing around USD 0.07. The execution of the `curate` function costs 54,948 gas which is equivalent to around USD 0.07. We reduce the linear complexity of the `curate` function (10) to constant by executing the assignment of agents for the next round in the `update` function. The `curate` function updates the round counter and activates the mapping for the next round.

8 RELATED WORK

To the best of our knowledge, BALANCE is the first reputation-based system for the dynamic adjustment of cryptocurrency deposits. There are three discernible strands of related literature.

The first strand relates to Token Curated Registries (TCRs) [30], inspiring the layered aspects of BALANCE [27]. A TCR formally represents a set $\mathcal{R}$ in which elements n can be included in a set through a token-based voting mechanism. A variety of different TCR types

have been proposed [10, 14, 27, 30, 39]. Notably, ranked TCRs $\mathcal{R}_O$ enable agents to vote on the rank (i.e. position) of an element in a set of elements such that $n_i \prec n_{i+1}$. Further, layered TCRs are a set $\mathcal{R}_\mathcal{L}$, consisting of distinct subsets where $\cup_{i=1}^{n} R_i \subseteq \mathcal{R}_\mathcal{L} \wedge \cap_{i=1}^{n} R_i = \emptyset$. However, while BALANCE takes inspiration from a ranked and layered TCR construction, TCRs require voting by individuals with tokens, adding significant and potentially unwarranted complexity [2].

The second strand concerns *reputation* aspects of BALANCE. [44] uses a notion of reputation to define a miner's power in terms of its work performed over the lifetime of a blockchain, as opposed to instantaneous mining power, in order to mitigate the vulnerability of blockchains to 51% attacks, where an adversary momentarily possesses more than 50% of the total mining power. Another system for reputation management in decentralised systems, but this time for users of the system as opposed to miners, is [25]. The mechanism uses cryptocurrencies to measure trust: using deposits between different agents, the authors construct a web-of-trust like architecture. BALANCE is different, creating only direct trust relationships between agents and actions are directly evaluated through the agreements in a smart contract. Taxonomies of reputation based systems [18, 19, 36] indicate comparators for BALANCE and other reputation based systems. BALANCE is a *quantitative* as opposed to qualitative trust system, expressing a reputation in terms of a deposit factor. Reputation itself is accrued by agents through direct experience, i.e. direct interaction with a smart contract. In particular, there is no transitive reputation between peers: the trust that *others* place in an agent does not confer trust onto the agent [14]. Some systems such as Pisa [28] do allow for transitive trust, allowing a reputation to be indirectly established.

A third strand of literature is pursued by [8]. Seeking to guard against losses of users' funds by centralized exchanges [32], the work focuses on providing a privacy-preserving system for proving the solvency of an exchange, i.e. that an exchange controls sufficient reserves to settle the account of each user.

9 CONCLUSION

BALANCE is an application-agnostic system, intended as an extension to existing cryptoeconomic protocols, that allows for the reduction of cryptocurrency deposits without compromising their security. By explicitly modelling agents' utilities, we show that it features an incentive-compatible mechanism that rewards agents for the performance of desired actions by reducing their required deposits, and therefore the opportunity costs of those deposits. Moreover, we show that the addition of BALANCE increases social welfare. We also implement BALANCE, integrating it with XCLAIM.

The primary motivating force for agents in our construction is the expected reduction in opportunity costs for the agent resulting from forgone returns on deposits. If we modify the assumption of perfect competition in Section 5, such that agents receive a positive payment from performing a desired action, payments p could also constitute a motivating force. For protocols where a reduction in the deposit is not practical or not desired by the protocol designers, the factor in each layer could be used to calculate the payment. In

---

[12]Implementation available at https://github.com/nud3l/layered-tcr.
[13]All USD conversions are made with an Ethereum exchange rate of USD 230 and a gas price of 6 GWEI.

[14]One caveat is that in one sense trust is transitive across agreements, since an agent may participate in multiple agreements within the same smart contract, retaining the same deposit factor.

this case, the factor would increase with every layer so that agents in the lowest layer receive a payment of pf_1, where e.g. $f_1 = 0.6$.

To the best of our knowledge, this paper is the first to use a curated registry to enable dynamic deposit requirements. One question that arises is that given v is private information, at what level should deposits and factors be set such that the proportion of agents who find it incentive compatible to perform desired actions is optimal for the protocol (or society) as a whole? In addition, we plan to explore different parameter configurations such as overlapping boundaries and restrict the number of agents per layers, as well as extend the model to cover probabilistic formulations of the specification.

ACKNOWLEDGMENTS

REFERENCES

[1] Amitanand S Aiyer, Lorenzo Alvisi, Allen Clement, Mike Dahlin, Jean-Philippe Martin, and Carl Porth. 2005. BAR fault tolerance for cooperative services. *ACM SIGOPS Operating Systems Review* 39, 5 (2005), 45. DOI:http://dx.doi.org/10.1145/1095809.1095816

[2] Aditya Asgaonkar and Bhaskar Krishnamachari. 2018. Token Curated Registries - A Game Theoretic Approach. (2018). http://arxiv.org/abs/1809.01756

[3] Christian Badertscher, Peter Gaži, Aggelos Kiayias, Alexander Russell, and Vassilis Zikas. 2018. Ouroboros Genesis: Composable Proof-of-Stake Blockchains with Dynamic Availability. *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security - CCS '18* 780477 (2018), 913–930. DOI: http://dx.doi.org/10.1145/3243734.3243848

[4] Christian Badertscher, Ueli Maurer, Daniel Tschudi, and Vassilis Zikas. 2017. *Bitcoin as a transaction ledger: A composable treatment*. Vol. 10401 LNCS. 324–356 pages. DOI:http://dx.doi.org/10.1007/978-3-319-63688-7{_}11

[5] Lorenz Breidenbach, Phil Daian, Floriantra Er, and Ari Juels. 2018. Enter the Hydra: Towards Principled Bug Bounties and Exploit-Resistant Smart Contracts. In *27th USENIX Security Symposium (USENIX Security 18)*. 1335–1352. https://thehydra.io/paper.pdf

[6] Vitalik Buterin. 2015. On Public and Private Blockchains - Ethereum Blog. (2015). https://blog.ethereum.org/2015/08/07/on-public-and-private-blockchains/

[7] Miguel Castro and Barbara Liskov. 1999. Practical Byzantine Fault Tolerance. *Proceedings ofthe Third Symposium on Operating Systems Design and Implementation* February (1999), 173–186. DOI:http://dx.doi.org/10.1.1.17.7523

[8] Gaby G Dagher, Benedikt Bünz, Joseph Bonneau, Jeremy Clark, and Dan Boneh. 2015. Provisions: Privacy-preserving proofs of solvency for Bitcoin exchanges. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*. ACM, 720–731. http://users.encs.concordia.ca/~clark/papers/2015_ccs.pdf

[9] Bernardo David, Peter Gaži, Aggelos Kiayias, and Alexander Russell. 2018. Ouroboros Praos: An Adaptively-Secure, Semi-synchronous Proof-of-Stake Blockchain. In *Advances in Cryptology – EUROCRYPT 2018*, Jesper Buus Nielsen and Vincent Rijmen (Eds.), Vol. 46. Springer International Publishing, Cham, 66–98. DOI:http://dx.doi.org/10.1007/978-3-319-78375-8{_}3

[10] Simon de la Rouviere. 2017. Introducing Curation Markets: Trade Popularity of Memes & Information (with code)! (2017). https://medium.com/@simondlr/introducing-curation-markets-trade-popularity-of-memes-information-with-code-70bf6fed9881

[11] Christian Decker and Roger Wattenhofer. 2013. Information propagation in the bitcoin network. In *Peer-to-Peer Computing (P2P), 2013 IEEE Thirteenth International Conference on*. IEEE, 1–10.

[12] Stefan Dziembowski, Lisa Eckey, and Sebastian Faust. 2018. FairSwap: How To Fairly Exchange Digital Goods. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security - CCS '18*. ACM Press, New York, New York, USA, 967–984. DOI:http://dx.doi.org/10.1145/3243734.3243857

[13] Ittay Eyal and Emin Gn Sirer. 2014. Majority Is Not Enough: Bitcoin Mining Is Vulnerable. In *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*. Vol. 8437. Berlin, Heidelberg, 436–454. DOI:http://dx.doi.org/10.1007/978-3-662-45472-5{_}28

[14] Sebastian Gajek. 2018. Graded Token-Curated Decisions with Up-/Downvoting. (2018). https://medium.com/coinmonks/graded-token-curated-decisions-with-up-downvoting-designing-cryptoeconomic-ranking-and-2ce7c000bb51

[15] Juan Garay, Aggelos Kiayias, and Nikos Leonardos. 2015. The bitcoin backbone protocol: Analysis and applications. In *Advances in Cryptology-EUROCRYPT 2015*. Springer, 281–310. http://courses.cs.washington.edu/courses/cse454/15wi/papers/bitcoin-765.pdf

[16] Dominik Harz and Magnus Boman. 2019. The Scalability of Trustless Trust. In *2nd Workshop on Trusted Smart Contracts (Financial Cryptography and Data Security)*. 279–293. DOI:http://dx.doi.org/10.1007/978-3-662-58820-8{_}19

[17] Ethan Heilman, Alison Kendler, Aviv Zohar, and Sharon Goldberg. 2015. Eclipse Attacks on Bitcoin's Peer-to-Peer Network. In *24th USENIX Security Symposium (USENIX Security 15)*. 129–144. https://www.usenix.org/system/files/conference/usenixsecurity15/sec15-paper-heilman.pdf

[18] Ferry Hendrikx, Kris Bubendorfer, and Ryan Chard. 2015. Reputation systems: A survey and taxonomy. *J. Parallel and Distrib. Comput.* 75 (1 2015), 184–197. DOI:http://dx.doi.org/10.1016/j.jpdc.2014.08.004

[19] Audun Jøsang, Roslan Ismail, and Colin Boyd. 2007. A survey of trust and reputation systems for online service provision. *Decision Support Systems* 43, 2 (3 2007), 618–644. DOI:http://dx.doi.org/10.1016/j.dss.2005.05.019

[20] Harry Kalodner, Steven Goldfeder, Xiaoqi Chen, S Matthew Weinberg, and Edward W Felten. 2018. Arbitrum: Scalable, private smart contracts. In *Proceedings of the 27th USENIX Conference on Security Symposium*. USENIX Association, 1353–1370.

[21] Sepandar D Kamvar, Mario T Schlosser, and Hector Garcia-Molina. 2003. The Eigentrust algorithm for reputation management in P2P networks. In *Proceedings of the twelfth international conference on World Wide Web - WWW '03*. ACM Press, New York, New York, USA, 640. DOI:http://dx.doi.org/10.1145/775152.775242

[22] Ahmed Kosba, Andrew Miller, Elaine Shi, Zikai Wen, and Charalampos Papamanthou. 2016. Hawk: The blockchain model of cryptography and privacy-preserving smart contracts. In *Symposium on Security & Privacy*. IEEE. http://eprint.iacr.org/2015/675.pdf

[23] Ranjit Kumaresan and Iddo Bentov. 2014. How to use bitcoin to incentivize correct computations. In *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security*. ACM, 30–41. http://www.cs.technion.ac.il/~idddo/incentivesBitcoin.pdf

[24] Yoad Lewenberg, Yoram Bachrach, Yonatan Sompolinsky, Aviv Zohar, and Jeffrey S Rosenschein. 2015. Bitcoin Mining Pools: A Cooperative Game Theoretic Analysis Categories and Subject Descriptors. *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2015)* (2015), 919–927.

[25] Orfeas Stefanos Thyfronitis Litos and Dionysis Zindros. 2017. Trust Is Risk: A Decentralized Financial Trust Platform. (2017).

[26] Loi Luu, Jason Teutsch, Raghav Kulkarni, and Prateek Saxena. 2015. Demystifying incentives in the consensus computer. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*. ACM, 706–719. http://www.comp.nus.edu.sg/~prateeks/papers/VeriEther.pdf

[27] Trent McConaghy. 2018. The Layered TCR. (2018). https://blog.oceanprotocol.com/the-layered-tcr-56cc5b4cdc45

[28] P McCorry, S Bakshi, I Bentov, Andrew Miller, and … 2018. Pisa: Arbitration Outsourcing for State Channels. (2018). https://www.cs.cornell.edu/~iddo/pisa.pdf

[29] Patrick McCorry, Alexander Hicks, and Sarah Meiklejohn. 2018. Smart Contracts for Bribing Miners. In *Financial Cryptography and Data Security. FC 2018.*, Vol. 10958. Springer Berlin Heidelberg, 3–18. DOI:http://dx.doi.org/10.1007/978-3-662-58820-8{_}1

[30] Mike Goldin. 2017. Token-Curated Registries 1.0. (2017). https://medium.com/@ilovebagels/token-curated-registries-1-0-61a232f8dac7

[31] Andrew Miller, Yu Xia, Kyle Croman, Elaine Shi, and Dawn Song. 2016. The honey badger of BFT protocols. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*. ACM, 31–42.

[32] Moore, Tyler and Christin, Nicolas. 2013. Beware the middleman: Empirical analysis of Bitcoin-exchange risk. In *International Conference on Financial Cryptography and Data Security*. Springer, 25–33.

[33] Satoshi Nakamoto. 2008. Bitcoin: A Peer-to-Peer Electronic Cash System. https://bitcoin.org/bitcoin.pdf. (Dec 2008). https://bitcoin.org/bitcoin.pdf Accessed: 2015-07-01.

[34] Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V. Vazirani. 2007. *Algorithmic Game Theory*. Vol. 1. Cambridge University Press, Cambridge. 1–754 pages. DOI:http://dx.doi.org/10.1017/CBO9780511800481

[35] Rafael Pass, Lior Seeman, and Abhi Shelat. 2017. Analysis of the Blockchain Protocol in Asynchronous Networks. Vol. 10210. 643–673. DOI:http://dx.doi.org/10.1007/978-3-319-56614-6{_}22

[36] Isaac Pinyol and Jordi Sabater-Mir. 2013. Computational trust and reputation models for open multi-agent systems: A review. *Artificial Intelligence Review* 40, 1 (2013), 1–25. DOI:http://dx.doi.org/10.1007/s10462-011-9277-z

[37] Joseph Poon and Vitalik Buterin. 2017. Plasma: Scalable Autonomous Smart Contracts. https://www.plasma.io/. (Aug 2017). https://plasma.io/plasma.pdf

Accessed: 2017-08-10.

[38] Joseph Poon and Thaddeus Dryja. 2016. The bitcoin lightning network. https://lightning.network/lightning-network-paper.pdf. (2016). https://lightning.network/lightning-network-paper.pdf Accessed: 2016-07-07.

[39] Achill Rudolph. 2018. Ranking Token Curated Registries. (2018). https://medium.com/coinmonks/ranking-token-curated-registries-e9a92dc85b31

[40] Tuomas W. Sandholm and Victor R. Lesser. 2001. Leveled Commitment Contracts and Strategic Breach. *Games and Economic Behavior* 35, 1-2 (2001), 212–270. DOI:http://dx.doi.org/10.1006/game.2000.0831

[41] Tendermint. 2016. Introduction to Tendermint - Tendermint. (2016). https://tendermint.com/intro

[42] Jason Teutsch and Christian Reitwießner. 2017. A scalable verification solution for blockchains. (2017).

[43] Gavin Wood. 2014. Ethereum: a secure decentralised generalised transaction ledger. (2014).

[44] Jiangshan Yu, David Kozhaya, Jeremie Decouchant, and Paulo Verissimo. 2019. RepuCoin: Your Reputation is Your Power. *IEEE Trans. Comput.* (2019), 1–1. DOI:http://dx.doi.org/10.1109/TC.2019.2900648

[45] Alexei Zamyatin, Dominik Harz, Joshua Lind, Panayiotis Panayiotou, Arthur Gervais, and William J. Knottenbelt. 2019. XCLAIM: Trustless, Interoperable, Cryptocurrency-Backed Assets. In *Proceedings of the IEEE Symposium on Security & Privacy, May 2019.* 1254–1271. DOI:http://dx.doi.org/10.1109/SP.2019.00085

[46] Fan Zhang, Ethan Cecchetti, Kyle Croman, Ari Juels, and Elaine Shi. 2016. Town Crier: An Authenticated Data Feed for Smart Contracts. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security.* ACM, 270–282. https://eprint.iacr.org/2016/168.pdf

[47] Guy Zyskind, Oz Nathan, and Alex Pentland. 2015. Enigma: Decentralized computation platform with guaranteed privacy. arXiv preprint arXiv:1506.03471. (2015). https://arxiv.org/pdf/1506.03471.pdf Accessed: 2017-08-22.

A CRYPTOECONOMIC PROTOCOLS

Figure 5: A cryptoeconomic protocol consists of three states. (1) an agent commits to the fulfilment of a specification ϕ. (2) the agent executes an action that either evaluates to true or false w.r.t. the specification. (3) the protocol is concluded by rewarding or punishing the agent for its action.

B DEPOSIT LAYERS

Figure 6: An example of a verifiable layered curated registry BALANCE with three layers. The more an agent contributes to the protocol the further it moves up the layers (from 1 to 3). The higher the layer, the lower the deposit needed.

C INCENTIVE PROOFS

The rational agent type $\mathcal{T}_r$ needs to decide which action to perform at the highest layer. From Figure 2, we see that at the highest layer the agent has either the choice to perform the desired action which results in a utility of $p - c_A - \mathrm{E}[rD_3]$ or an undesired action which is equivalent to a utility of $v - c_A - \mathrm{E}[rD_3] - D_3$. In this simplified figure, there are three layers. However, we can generalise this by assuming ω layers. By Definition 5, the rational type cannot come to a decision just considering these two utility functions. Reaching a decision requires $\mathcal{T}_r$ to consider several time-steps. By integrating BALANCE, we transform the decision process into a *sequential game*.

C.1 Decision boundary for type $\mathcal{T}_r$

The total utility received by an agent performing desired actions in the highest layer, for a sequence that lasts $t = \omega$ rounds, is expressed by:

$$u_A = \sum_{t=0}^{\omega-1} \left(\frac{\delta}{1+r}\right)^t \left(p - c_A - \mathrm{E}[rD_\omega]\right) \qquad (24)$$

Summing over cycles to infinity yields the following expression (τ swapped for t to express 'within' cycle time):

$$u_A = \sum_{\gamma=0}^{\infty} (\frac{\delta}{1+r})^{\gamma\omega} \left(\sum_{\tau=0}^{\omega-1} (\frac{\delta}{1+r})^\tau (p - c_A - \mathrm{E}[rD_\omega]) \right) \qquad (25)$$

The total utility corresponding to an agent who commits an undesired action in the highest layer ω, followed by $\omega - 1$ desired actions to return to that layer, can be expressed as a sequence with $t = \omega$ rounds as follows. We express the deposit D of layer $L_m \in \mathcal{L}$ at time t with D_t.

$$u_A = v - c_A - \mathrm{E}[rD_\omega] - D_\omega + \sum_{t=1}^{\omega-1} \left(\frac{\delta}{1+r}\right)^t \left(p - c_A - \mathrm{E}[rD_t]\right) \qquad (26)$$

Summing over cycles to infinity yields the following expression (τ swapped for t to express 'within' cycle time):

$$u_A = \sum_{\gamma=0}^{\infty} (\frac{\delta}{1+r})^{\gamma\omega} \left(v - c_A - D_\omega - \mathrm{E}[rD_\omega] \right.$$
$$\left. + \sum_{\tau=1}^{\omega-1} (\frac{\delta}{1+r})^\tau (p - c_A - \mathrm{E}[rD_\tau]) \right) \qquad (27)$$

The decision boundary for agents of type $\mathcal{T}_r$ is found by equating equations (25) and (27). This is as follows.

$$\sum_{\gamma=0}^{\infty} (\frac{\delta}{1+r})^{\gamma\omega} \left(v - c_A - D_\omega - \mathrm{E}[rD_\omega] \right.$$
$$\left. + \sum_{\tau=1}^{\omega-1} (\frac{\delta}{1+r})^\tau (p - c_A - \mathrm{E}[rD_\tau]) \right)$$
$$= \sum_{\gamma=0}^{\infty} (\frac{\delta}{1+r})^{\gamma\omega} \left(\sum_{\tau=0}^{\omega-1} (\frac{\delta}{1+r})^\tau (p - c_A - \mathrm{E}[rD_\omega]) \right) \qquad (28)$$

The following lemma simplifies this boundary condition.

Lemma 2. Comparing payoffs of desired as opposed to undesired actions into the infinite horizon, with $\gamma \rightarrow \infty$, is the same as comparing the payoffs of these actions over a single cycle.

PROOF. Inspection of (28) shows that the two sides of the equation are equal when the terms inside the summation operator (summing from $\gamma = 0$ to ∞) are equal. □

Therefore we can equation becomes

$$v - c_A - D_\omega - \mathrm{E}[rD_\omega] + \sum_{t=1}^{\omega-1}(\frac{\delta}{1+r})^t(p - c_A - \mathrm{E}[rD_t])$$
$$= \sum_{t=0}^{\omega-1}(\frac{\delta}{1+r})^t(p - c_A - \mathrm{E}[rD_\omega]) \tag{29}$$

Which simplifies to:

$$v = p + D_\omega + \sum_{t=1}^{\omega-1}(\frac{\delta}{1+r})^t(\mathrm{E}[rD_t] - \mathrm{E}[rD_\omega]) \tag{30}$$

If we equate (24) and (26) we get the decision boundary for an agent to perform a desired or undesired action. We express this by the factor f. By (7), f_m determines the deposit of a layer m in relation to the base deposit D_{base}.

Lemma 3. Introducing BALANCE creates a sequential game in which factors f determine the deposit at each layer. The factor of the highest layer f_ω in a sequential game with at least two layers can be set lower than f_1 and still achieves a higher utility for a desired action compared to an undesired action.

PROOF. First, we equal (24) and (26) to describe the decision boundary. Next, we set $D_\omega = f_\omega D_{\mathrm{base}}$ and $D_t = f_t D_{\mathrm{base}}$ (where t and m are are substitutes for a single cycle). This leaves us with the following equality.

$$\sum_{t=0}^{\omega-1}\left(\frac{\delta}{1+r}\right)^t\left(p - c_A - \mathrm{E}[rf_\omega D_{\mathrm{base}}]\right)$$
$$= v - c_A - \mathrm{E}[rf_\omega D_{\mathrm{base}}] - (f_\omega D_{\mathrm{base}}) \tag{31}$$
$$+ \sum_{t=1}^{\omega-1}\left(\frac{\delta}{1+r}\right)^t\left(p - c_A - \mathrm{E}[rf_t D_{\mathrm{base}}]\right)$$

We assume that the agent is able to progress to the next layer after every time step t. Also, we apply our initial assumptions that $p - c_A$ for performing a desired action is 0, and that c_A for the undesired action is 0. We then simplify (31).

$$\sum_{t=1}^{\omega-1}\left(\frac{\delta}{1+r}\right)^t\left(-\mathrm{E}[rf_\omega D_{\mathrm{base}}]\right)$$
$$= v - f_\omega D_{\mathrm{base}} + \sum_{t=1}^{\omega-1}\left(\frac{\delta}{1+r}\right)^t\left(-\mathrm{E}[rf_t D_{\mathrm{base}}]\right) \tag{32}$$

$$v = f_\omega D_{\mathrm{base}} + \sum_{t=1}^{\omega-1}\left(\frac{\delta}{1+r}\right)^t\left(\mathrm{E}[rD_{\mathrm{base}}(f_t - f_\omega)]\right) \tag{33}$$

In (33), we can re-arrange for f_ω. When we equate v and D_1, we imply the decision boundary without BALANCE. We can then express f_ω in relative terms by setting $D_{\mathrm{base}} = 1$ and $v = D_1 = f_1 * D_{\mathrm{base}}$. Also, we adjust the term to f_ω.

$$f_\omega = \frac{\mathrm{E}[rf_t]\sum_{t=1}^{\omega-1}\left(\frac{\delta}{1+r}\right)^t}{\mathrm{E}[r]\sum_{t=1}^{\omega-1}\left(\frac{\delta}{1+r}\right)^t} \tag{34}$$

From (34), we observe that if f_t is smaller than 1, then f_ω would become greater than 1. This would allow the agent to perform undesired actions in the first layer at a positive utility since v would become larger than $D_1 f_1 - p$. Therefore, the factor at the lowest layer must be greater than 1. If we set any factor f_t equal to 1, f_ω becomes also 1. Consequently, allowing a factor f_t to be greater than 1 allows us to set a deposit factor f_ω smaller than f_1. □

Lemma 4 (Minimum deposit at the lowest layer). The deposit factor at the lowest layer f_1 must be greater than 1 to allow the factor at the highest layer f_ω to be less than 1.

PROOF. The proof follows from the proof of Theorem 3. □

C.2 Linear factor adjustment

We assume a linear relationship between the smallest factor f_ω and the other factors f_t. Next, we assume that at D_1 the factor f_1 is > 1 and that at D_ω the factor $f_\omega \geq 1$ is the lowest factor. From this, we can calculate linear relation between f_1 and f_ω as follows.

$$f_t = f_1 - \left(\frac{f_1 - f_\omega}{\omega - 1}\right)(t - 1) \tag{35}$$

We can then replace f_t in (34) by (35) to express the term purely in terms of f_ω. Further, we assume that the valuation v is lower than the deposit D_1, i.e. the deposit is higher than the desire for performing an undesired action. This leaves us with the following equation:

$$f_{\bar{\omega}} = \frac{v\omega - v - \sum_{t=1}^{\omega-1}\left(\frac{\delta}{1+r}\right)^t \mathrm{E}[r(f_1\omega - f_1 t)]}{\omega - 1 - r\sum_{t=1}^{\omega-1}\left(\frac{\delta}{1+r}\right)^t \mathrm{E}[r(\omega - t)]} \tag{36}$$

D XCLAIM

D.1 Currency fluctuations and valuations

ETH-BTC fluctuation. In order to determine the buffer for the deposit to account for exchange rate fluctuations we collect the daily high and low prices of BTC to USD and ETH to USD from a period of one year (May 3, 2018 to May 3, 2019). The data is collected from the Poloniex exchange API[15]. We are interested in the *drop* of the exchange rate, i.e. from a the high price to the low price. Therefore we calculate the drop as $\frac{p_{\mathrm{low}} - p_{\mathrm{high}}}{p_{\mathrm{high}}}$. Using the exchange data, we find that the highest drop within the year is 17.06% and the average per day drop is 3.90% with a standard deviation of 2.73%. The maximum daily exchange rate drop is visualised in Figure 7.

---

[15]Taken from https://docs.poloniex.com/

Figure 7: Using daily data from the Poloniex exchange, this figure plots the percentage decrease (in absolute terms) of the low price relative to the high price for a given day.


Figure 8: Boundary of f_ω for the initial factor of $f_1 = \{2.06\}$ with an additional buffer for exchange rate fluctuations and valuation differences. The function is the decision boundary for an economically rational agent to decide between a desired and undesired action given a number of time steps t. Assuming $r = 0.05$, $\delta = 0.9$, $\omega = 12$.

ETH-BTC valuations. Agents can express a different preference for cryptocurrencies. To measure this difference we scanned the order books of the Binance exchange for bids and asks for a period of one week from May 3 to May 9, 2019[16]. We calculated the valuation difference by comparing the bid/ask price to the exchange rate at the time the order is placed. We note that most agents place orders close to the actual exchange rate as the mean difference between price and exchange rate is 0.64% with a standard deviation of 2.73%. However, there are significant outliers with up to 76.34% in this data. The valuation difference serves as an orientation to calculate the security buffer for the cross-currency deposit.

We note several shortcomings of our approach. Agents might place extreme orders to influence the exchange rate price instead of actually valuation the currency at that level. Further, we did not consider a *weighted* approach where the difference between price and exchange rate also considers the size of the order. Further, large difference can also be caused by human error when entering the order.

D.2 Factor calculation

We can calculate the factors taking the highest deviations from the exchange rate fluctuation and the valuation differences. We define the highest exchange rate fluctuation as $\Delta_{ex} = 0.1706$ and the highest valuation difference as $\Delta_v = 0.7634$. We calculate a buffer b_{deposit}:

$$b_{\mathrm{deposit}} = (1 + \Delta_{ex}) * (1 + \Delta_v) \tag{37}$$

We then determine the factor f_1 by using a base deposit of 1: $f_1 = 1 * b_{\mathrm{deposit}} = 2.06424$. We assume $r = 0.05$, $\delta = 0.9$, $\omega = 12$ and linear relationship between factors. Further, we apply the buffer b_{deposit} to all subsequent factors f. We apply these assumptions and results to (36). This yields $f_\omega = 1.85$. Our results are visualised in Figure 8.

D.3 Security arguments

We discuss possible effects on XCLAIM by integrating BALANCE.

---

[16]Taken from https://github.com/binance-exchange/binance-official-api-docs/blob/master/rest-api.md

Vaults A take multiple roles to reduce their deposit. There is no guarantee that A has an opportunity to perform a desired action within a given period t. A would thus fall to a lower layer. To prevent this, A could ensure a constant flow of redeem requests by creating an artificial demand as A can act as CbA Requester, Vault, and CbA Redeemer, called strategy s_{boosting} as described in Section 6.2. Assuming that expected interest on the deposit $\mathrm{E}[r] * D_{\mathrm{layer}(A)}$ is higher than the cost $c_A(s_{\mathrm{boosting}})$, this would be a dominant strategy. This lead to an issue: A could ensure that its actions reduce its deposit and could perform the undesired action, i.e. not executing a redeem request of B with a lower deposit and punishment. XCLAIM can prevent this by requiring a *maximum deposit* per vault A. Given an expected interest $\mathrm{E}r$, the maximum allowed deposit is defined by $D_{max} = \frac{c_A(s_{\mathrm{demand}})}{\mathrm{E}[r]}$. This also holds true under Sybil identities, since deposit reductions are not transitive between multiple identities.

Conflicting minimum and maximum deposits. XCLAIM suggests to prevent Sybil attacks by either (i) requiring a minimum deposit from vaults through the iSC, or (ii) a fee for issuing based on the total amount and not per issue request. The minimum deposit amount is in conflict with the Sybil resistance requirement for BALANCE as defined in Section 6.2. Hence, XCLAIM would need to adopt a fee model based on issue amounts rather than requiring a minimum deposit per vault when integrating BALANCE. As such, Sybil resistance in both XCLAIM and BALANCE is maintained.

Length of time period t motivates agents to delay actions by a vault. Assuming that an agent can fulfil more than one desired action within a time period t and A has already a high enough score to progress to a higher layer or remain in the highest. In that case, A can play a strategy s_{delay} to delay performing redeem requests ($\mathcal{A}_2$) and balancing its deposit ($\mathcal{A}_3$) until the next period (see also Section 6.3). In XCLAIM's case s_{delay} is limited by enforcing time limits on agreements $\mathcal{A}_{2,3,4}$. Thus, vaults do not delay their actions.

Artificial redeem requests. Since the score for the redeem request $\mathcal{A}_2$ is positive, A might try to obtain $i(b)$ to execute redeem requests

to maintain it's ranking. This would require that the cost for redeem is lower than the opportunity cost of remaining at a higher layer. This argumentation is similar to Section 6.2 as a maximum deposit will prevent a vault from having a higher gain in opportunity cost compared to the cost of executing redeem requests.

Deposit adjustments to improve score. The score to resolve a state where $D_A < \lceil D_A \rceil$ is positive. As such, A has an incentive to let its deposit fall below $\lceil D_A \rceil$ and resolve this state by adding more deposit, expressed by strategy $s_{\text{deposit-bouncing}}$. By having a positive score, BALANCE encourages $s_{\text{deposit-bouncing}}$. However, we can prevent the impact of this strategy by setting the score for such an action to be below a lower bound of a given layer L_m. For example, XCLAIM might allow A to improve its score with $s_{\text{deposit-bouncing}}$, but sets the score such that A is only able to reach a layer $L_m < L_\omega$. Thereby, it *limits* $s_{\text{deposit-bouncing}}$ but does not prevent it.

E   NOTATION AND SYMBOLS

Table 1: Overview of notation and symbols.

Symbol	Definition
π	Cryptoeconomic protocol
$P = \{1, ..., n\}$	Set of n agents
$\sigma \in \{\Sigma_d, \Sigma_u\}$	Desired and undesired action
u_i	Utility of agent $i \in P$
v_i	Private valuation agent $i \in P$ attaches to the realisation of a particular outcome
c_i	Costs for agent $i \in P$
r	Interest rate
$E[..]$	Expected value
$\mathcal{A} = \langle \phi, p, \mathcal{D} \rangle$	Agreements with specification ϕ, payment p, and set of deposits $\mathcal{D}$
π_{BALANCE}	Cryptoeconomic protocol with BALANCE
$\mathcal{A} = \langle \phi, p, s, \mathcal{D} \rangle$	Agreements with additional score s
$\mathcal{L} = \{L_1 \prec ... \prec L_\omega\}$	An ordered list of ω layers
$L_m = \langle l, u, f \rangle$	Layer with a lower bound l, upper bound u, and a deposit factor f for layer $L_m \in \mathcal{L}$
D_m	Deposit at a specific layer $L_m \in \mathcal{L}$, e.g. D_1 for layer 1
f_m	Deposit factor for layer $L_m \in \mathcal{L}$
δ	Discount factor
t	Time step