

Robust Restaking Networks

Naveen Durvasula  

Columbia University, New York, NY, USA

Tim Roughgarden  

a16zcrypto, New York, NY, USA

Columbia University, New York, NY, USA

Abstract

We study the risks of validator reuse across multiple services in a restaking protocol. We characterize the robust security of a restaking network as a function of the buffer between the costs and profits from attacks. For example, our results imply that if attack costs always exceed attack profits by 10%, then a sudden loss of .1% of the overall stake (e.g., due to a software error) cannot result in the ultimate loss of more than 1.1% of the overall stake. We also provide local analogs of these overcollateralization conditions and robust security guarantees that apply specifically for a target service or coalition of services. All of our bounds on worst-case stake loss are the best possible. Finally, we bound the maximum-possible length of a cascade of attacks.

Our results suggest measures of robustness that could be exposed to the participants in a restaking protocol. We also suggest polynomial-time computable sufficient conditions that can proxy for these measures.

2012 ACM Subject Classification Networks → Network economics

Keywords and phrases Proof of stake, Restaking, Staking Risks

Digital Object Identifier 10.4230/LIPIcs.ITCS.2025.48

Related Version *Full Version:* <https://arxiv.org/abs/2407.21785>

Funding This research was supported in part by NSF awards CCF-2006737 and CNS-2212745, and research awards from the Briger Family Digital Finance Lab and the Center for Digital Finance and Technologies.

Acknowledgements We thank Tarun Chitra, Soubhik Deb, Sreeram Kannan, Mike Neuder, and Mallesh Pai for comments on earlier drafts of this paper. We thank Soubhik and Sreeram in particular for emphasizing the importance of local guarantees.

1 Introduction

1.1 Sharing Validators Across Services

Major blockchain protocols such as Bitcoin or Ethereum are “decentralized,” meaning that transaction execution is carried out by a large and diverse set of “validators.” Such protocols offer a form of “trusted computation,” in the sense that, because they are decentralized, no one individual or entity can easily interfere with their execution. A decentralized and Turing-complete smart contract platform such as Ethereum can then be viewed as a trusted programmable computer capable of performing arbitrary computations.

While Turing-complete, the computing functionality offered by Ethereum smart contracts suffers from limitations imposed by design decisions in the underlying consensus protocol. Most obviously, computation and storage in the Ethereum virtual machine is scarce, with perhaps 15–20 transactions processed per second. Could the Ethereum protocol be somehow bypassed, opening the door for different or more powerful computing functionality, while retaining at least some of the protocol’s decentralization? Or, what about applications that are not compatible with all Ethereum validators, perhaps due to demanding hardware requirements or regulatory constraints?



© Naveen Durvasula and Tim Roughgarden;

licensed under Creative Commons License CC-BY 4.0

16th Innovations in Theoretical Computer Science Conference (ITCS 2025).

Editor: Raghu Meka; Article No. 48; pp. 48:1–48:21

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

One natural approach to addressing these challenges is to allow the *reuse* of a blockchain protocol’s validators across multiple services, where a “service” is some task that could be carried out by some subset of validators. (The initial blockchain protocol can be viewed as the canonical service, performed by all validators.) For example, such services could include alternative consensus protocols (perhaps with higher throughput, a different virtual machine, or different consistency-liveness trade-offs), storage (“data availability”), or verifiable off-chain computation (“zk coprocessors”).¹

The obvious danger of validator reuse is an increased risk of a validator deviating from its intended behavior (e.g., due to overwhelming computational responsibilities). Our focus here is deliberate validator deviations in response to economic incentives, such as the profits that could be obtained by corrupting one or more services. The goal of this paper is to quantify such risks:

Under what conditions can validators be safely reused across multiple services?

1.2 Cryptoeconomic Security


We first review the usual “cryptoeconomic” approach to answering a more basic question: when is a blockchain protocol, without any additional services, “safe from attack”? The idea is to perform a cost-benefit analysis from the perspective of an attacker, and declare the protocol safe if the cost of carrying out an attack exceeds the profit that the attacker can expect from it; see also Figure 1. In the specific case of a proof-of-stake blockchain protocol with slashing (such as Ethereum), the cost can be estimated as the value of the validator stake that would be lost to slashing following an attack. For example, let V denote the set of validators of a proof-of-stake blockchain protocol and σ_v the stake of validator $v \in V$ (e.g., 32 ETH in Ethereum). In a typical PBFT-type protocol in which double-voting validators lose all of their stake, the cost of an attack (i.e., causing a consistency violation) can be bounded below by $\frac{1}{3} \sum_{v \in V} \sigma_v$. In this case, the protocol can be regarded as cryptoeconomically secure provided the estimated profit π of an attack is less than this quantity. To the extent that there is a “buffer” between $\frac{1}{3} \sum_{v \in V} \sigma_v$ and π , the protocol can be treated as “robustly secure,” meaning secure even after a sudden loss of some amount of stake (e.g., due to slashing caused by a software error).

1.3 Our Results: Robustly Secure Validator Reuse


Viewing the basic scenario of a blockchain protocol as a star graph (with one “service” representing the protocol connected to all the validators running it), the more complex scenario of validators reused across multiple services can be viewed as an arbitrary bipartite graph (Figure 2). As before, we suppose that each validator $v \in V$ has some stake σ_v that can be confiscated in the event that v participates in an attack on a service that it has agreed to support. There is now a set S of services, with π_s denoting the profit an attacker would obtain by compromising the service $s \in S$.² We assume that compromising a service $s \in S$ requires the participation of an α_s fraction of the overall validator stake supporting v (e.g., $\alpha_s = 1/3$). With this expanded network formalism to capture multiple services, when should we consider a network to be “secure”?

¹ Several projects, in various stages of production, are currently exploring this idea; the Eigenlayer restaking protocol [15] is perhaps the most well known of them. We stress that our goal here is to develop a model that isolates some of the fundamental challenges and risks of validator reuse, independent of any specific implementation of the idea.

² We follow [15] and assume that the estimates on attack profitability (the π_s ’s) are given. Developing tools to help produce such estimates in practice is an important open research direction.



■ **Figure 1** A blockchain protocol operated by a collection of validators, with π denoting the profit of successfully attacking the protocol and σ_v the amount of stake posted by a validator v as collateral.

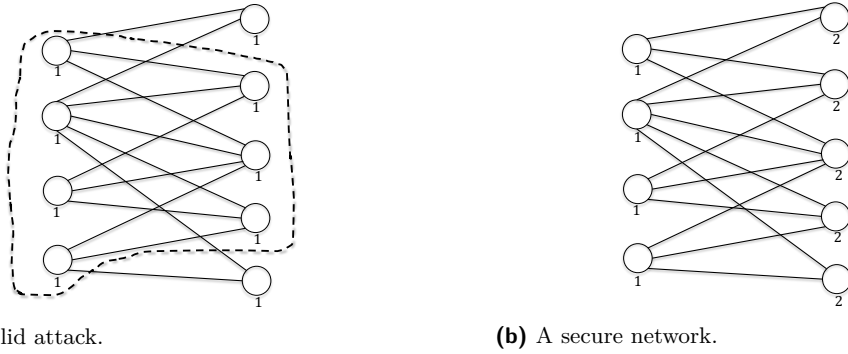


■ **Figure 2** A general restaking network, with validators reused across multiple services.

Intuitively, a network is insecure if there is a set of services that can be profitably attacked. This requires two conditions to be satisfied: the validators $B \subseteq V$ carrying out the attack must control sufficient stake to do so (for each service s in the attacked set A , the validators of B control at least a α_s fraction of the staked pledged to s), and they must profit from it (i.e., $\sum_{s \in A} \pi_s > \sum_{v \in B} \sigma_v$). We call such a pair (A, B) a *valid attack* and a network *secure* if there are no valid attacks; see also Figure 3. With multiple services, security is an inherently combinatorial (as opposed to binary) notion. For example, the computational problem of checking whether a network is secure is as hard as the (coNP-hard) problem of verifying the expansion of a bipartite graph (see [12]).

When is a network “robustly secure,” in the sense that the sudden loss of a small amount of stake cannot enable a catastrophic attack? Unlike the “all-or-nothing” version of this question with a single service, with multiple services, the following more fine-grained question is the appropriate one: given an initial shock in the form of a sudden loss of a ψ fraction of the overall stake, what is the total fraction of stake that might be lost following any consequent valid attacks?

As in the case of a single service, some amount of “buffer” in stake (relative to attack profits) is necessary for robust security. We parameterize this overcollateralization factor via a parameter γ and suppose that, whenever $B \subseteq V$ is a subset of validators capable of



■ **Figure 3** Two restaking networks. Each service (left-hand side vertex) and validator (right-hand side vertex) is labeled with its profit-from corruption or stake, respectively. Assume that a service can be corrupted if and only if it is attacked by at least half of its validators (i.e., $\alpha_s = 1/2$ for every service s). The restaking network in (a) is not secure because there is a valid attack (indicated by the dotted line): three validators can earn a profit of 4 by corrupting all four services while losing only three units of stake. The restaking network in (b) is secure.

corrupting all the services in $A \subseteq S$, the total stake of B is at least a $1 + \gamma$ factor larger than the total profit from corrupting all of A . For example, this condition holds in the network in Figure 3(b) for $\gamma = 1/2$ (but not for larger values of γ).

Our first main result (Theorem 4) precisely characterizes the worst-case (over bipartite graphs and shocks, as a function of γ) fraction of the overall stake that can be lost due to a shock of size ψ : $\left(1 + \frac{1}{\gamma}\right) \psi$. Because the network was secure prior to the shock, the value of a $\left(1 + \frac{1}{\gamma}\right) \psi$ fraction of the overall stake is also an upper bound on the total profit obtained from all of the attacked services.

We also show that our result is tight in a strong sense (Theorem 6 and Theorem 7): for every ψ , γ , and ϵ greater than zero such that $0 \leq \left(1 + \frac{1}{\gamma}\right) \psi - \epsilon \leq 1$, there exists a restaking graph in which a ψ fraction of the overall stake can disappear in a shock that results in the loss of at least a $\left(1 + \frac{1}{\gamma}\right) \psi - \epsilon$ fraction of the overall stake.³

Qualitatively, this result implies that a constant-factor strengthening of the obvious necessary condition for security automatically implies robust security. For example, if attack costs always exceed attack profits by 10%, then a sudden loss of .1% of the overall stake cannot result in the ultimate loss of more than 1.1% of the overall stake.

Our result suggests a “risk measure” that could be exposed to the participants in a restaking protocol, namely the maximum value of the buffer parameter γ that holds with respect to the current restaking network. We also suggest easy-to-check sufficient conditions that can proxy for this risk measure (Corollary 5). These conditions are similarly tight, as shown in Theorem 6 and Theorem 7.

1.4 Our Results: Local Robust Security Guarantees

The results described in Section 1.3 are “global” with respect to the network structure, in three distinct senses: (i) the overcollateralization condition is assumed to hold for every subset $B \subseteq V$ of validators and $A \subseteq S$ of services that B is capable of corrupting; (ii) the

³ After slightly reducing the validator stakes, the network in Figure 3(b) already shows that the bound is tight for the special case in which $\psi = 1/5$ and γ is arbitrarily close to $1/2$. (Consider a shock that knocks out the validator that is connected to all four services.)

initial shock can affect any subset of validators, subject to the assumed bound of ψ on the total fraction of stake lost; and (iii) any subset of validators might lose stake following the initial shock, subject to our upper bound of $(1 + \frac{1}{\gamma})\psi$ on the total fraction of lost stake.

Local guarantees

We next pursue more general “local” guarantees, which are parameterized by a set C of services. (The global guarantees will correspond to the special case in which $C = S$.) For example, C might be a set of closely related services that share a number of dedicated validators. The operators of such a set C might object to both the assumptions and the conclusion of our global guarantee:


- How can we be sure that the overcollateralization factor holds for services and validators that we know nothing about?
- And even if we could, how can we be sure that random validators that we have nothing to do with won’t suddenly lose their stake (e.g., because they supported a malicious or buggy service), resulting in an initial shock that causes the loss of more than a ψ fraction of the overall stake?
- And even if we could, how can we be sure that our validators won’t be the ones that lose their stake following a shock that is purely the fault of other services and/or validators?

To address these concerns, we next consider a refined version of the basic model, parameterized by a set C of services. We denote by $\Gamma(C)$ the validators that are exclusive to C , meaning that they contribute to no services outside C . (In the special case in which $C = S$, $\Gamma(C) = V$ and we recover the original model.) Intuitively, the validators in $\Gamma(C)$ are the ones that services in C are “counting on”; other validators support (potentially malicious or buggy) services outside C and, from C ’s perspective, could disappear at any time. We then restrict attention to initial shocks in which at most a ψ fraction of the total stake controlled by the validators of $\Gamma(C)$ is lost. (The shock can affect validators outside of $\Gamma(C)$ arbitrarily.) The goal is then to identify overcollateralization conditions guaranteeing that, no matter what the initial shock and subsequent attacks, the fraction of stake ultimately lost by the validators in $\Gamma(C)$ is bounded (by a function of the shock size ψ and an overcollateralization parameter γ).

Generalizing our global guarantees to local guarantees is not straightforward and, as the next two examples show, requires additional compromises. The first example shows that protection can be guaranteed only against a subset of valid attacks, a natural and well motivated subset that we call “stable attacks.” The second example shows that, even when restricting attention to stable attacks, overcollateralization is required not only for potential attacks, but more generally for what we call “attack headers.”

Necessity of restricting to stable attacks

In more detail, in the first example (depicted on the left of Figure 4), we consider a restaking network with two services and three validators. (Vertices in this figure correspond to validators, which were previously represented as the right-hand side vertices of a bipartite graph; each service is now identified with its neighborhood of validators.) The service highlighted in green on the left has 100 times as much stake securing it as its profit from corruption, and shares no validators with other services, so one might hope that it would be protected from any shocks that affect only the other validators. However, if the validator outlined in red disappears, the two validators outlined in yellow can attack both services for a profit of $102 - 101 > 0$. But this example is unsatisfying: the validator with stake 1



■ **Figure 4 Simple overcollateralization is insufficient in the local setting.** There are two restaking networks shown above. In each, the validators are denoted along with their corresponding stakes by the yellow circles. Services and their profits from corruption are denoted by the blue rounded squares. In each of these networks, the service outlined in green is overcollateralized. However, despite being unrelated to the overcollateralized service, if the validator outlined in red disappears, the validators in yellow can attack all services (including the overcollateralized one).

could have attacked the service with profit from corruption 101 on its own to yield a profit of 100. The addition of the validator with stake 100 added 100 to the cost of the attack, but only added 1 (from the corruption of the service highlighted in green) to the profit. We show in Theorem 8 that the issue suggested by this example is fundamental: without further restrictions on attacks that rule out contrived examples such as this one, there does not exist a local condition that guarantees local security. In response, we confine attention to sequences of what we call *stable* attacks in which all of the attacking validators contribute positively to the profit of the attack (as opposed to free riding on the profits attributable to other attacking validators). The attack in the example is not stable, as the validator with stake 100 was not a profitable addition to the attack that could have been carried out by the validator with stake 1.

Necessity of overcollateralizing attack headers

Even if we restrict our attention to stable attacks, simple overcollateralization is insufficient to guarantee local security. To appreciate the issue, consider the restaking network shown on the right in Figure 4. In this example, there are three services, each with a profit from corruption of 1, and three services each with stake 1. Each validator is used to secure two different services. The service outlined in green is overcollateralized in that its profit from corruption is 1, but two units of stake are securing it. Despite this, if the unrelated validator outlined in red disappears, the two validators outlined in yellow can attack all three services for a profit of $3 - 2 > 0$. Furthermore, if all validators are required to attack each service (i.e., $\alpha_s = 1$), then this attack is stable, as the inclusion of both validators in the attack is profitable. Thus, despite being overcollateralized, a stable attack can be launched on the service highlighted in green even if stake unrelated to the service disappears in a shock.

We find that to guarantee robust security for a coalition of services C using only “local” information, it is necessary to overcollateralize not only pairs (A, B) where B is a set of validators capable of attacking all services in A , but a more general collection of pairs that we

call *attack headers*. Informally, this amounts to requiring that there is some “buffer” in stake for any potential attack on some services in C even if we were to allow every validator that is also securing a service outside of C to join the attack without considering their profitability. Our main result here (Theorem 10) formally provides a local condition guaranteeing that, whenever an initial shock knocks out at most a ψ fraction of the stake that provides security exclusively to C , the worst-case loss of such stake, after an arbitrary sequence of stable attacks, is at most a $(1 + \frac{1}{\gamma})\psi$ fraction. We show that our bounds are tight and indeed require overcollateralization of all attack headers (Corollary 14 and Theorem 15), and again provide easily computable sufficient (and similarly tight) conditions that can proxy for the overcollateralization condition (Corollary 13). Our local condition generalizes the global overcollateralization condition, with the latter corresponding to the special case of the former in which $C = S$.

1.5 Our Results: Cascading Attacks

An attack on a restaking network results in a loss of stake (of the attacking validators), and this may introduce new opportunities for other sets of validators to carry out profitable attacks. That is, an initial shock may set off an entire *cascade* of attacks. (All of the bounds on stake loss described in Sections 1.3 and 1.4 hold for cascades of attacks of arbitrary length.) Our final result concerns the maximum-possible length of such a cascade, and shows that this quantity is also governed in part by the overcollateralization factor γ .

Precisely, we define the *reference depth* of a cascade of attacks as a measure of the “long-range dependence” between different attacks in an attack sequence. For example, if each attack in the sequence is directly enabled by the loss of the validators slashed in the previous attack, then the reference depth of the sequence is 1. Our main result here (Theorem 16) bounds the maximum-possible attack length as a function of the reference depth k , the shock size ψ , the overcollateralization factor γ , and the minimum stake ϵ held by a validator: $k(1 + \log_{1+\gamma}(\frac{\psi \cdot \lceil \text{total stake} \rceil}{\epsilon \gamma}))$. For example, in the case of constant reference depth and equal validator stake amounts, the worst-case attack length is logarithmic in the number of validators with overcollateralization and linear without it.

1.6 Related Work

Our focus on the risks of cascading failures following a small shock echoes some of themes in the literature on systemic risk in financial networks. For example, Eisenberg and Noe [9] that study the existence and structure of inter-firm payments in a financial network following a default. This work is extended by Glasserman and Young [11], who study how “connectedness,” meaning the fraction of liabilities that a firm externally owes, affects contagion risk. Acemoglu et al. [2] build further on this work and study network “stability,” meaning the propensity for shocks to propagate; they show that connectivity initially improves stability, but then at a phase transition, denser connectivity leads to increased shock propagation. In a subsequent paper, Acemoglu et al. [1] unify a number of the preceding results.

A separate line of work, beginning with Chen et al. [6], aims to axiomatically characterize systemic risk measures. The model in [6] can capture, in particular, a contagion model characterized by a matrix of profits and losses over different firms and outcomes. The results in [6] characterize the global measures of risk (operating on the matrix) that satisfies certain sets of desirable axioms. This work is expanded upon in Kromer et al. [13], where the authors consider general outcome measure spaces, as well as in Feinstein et al. [10], where the authors consider set-valued risk measures. Battison et al. [4] study systemic risk measurement

when a regulator has limited information about contracts made between financial network participants (e.g., the fraction of the face value that can likely be recovered if a counterparty defaults), and show how small errors in knowledge can lead to large errors in systemic risk measurement.

Our work also shares some conceptual similarity with the well-known work of Diamond and Dybvig [8] on bank runs and of Brunnermeier et al. [5] on the risks of re-hypothecation.

The model in the present work differs substantially from those considered in the aforementioned papers, in large part because of the idiosyncrasies of restaking networks, including their combinatorial and bipartite nature and their susceptibility to economically motivated attacks.

Restaking has also, to a limited extent, been studied in its own right. The EigenLayer team introduces restaking in [15]. The framing of an economically motivated attacker that trades off stake loss with profits from corruption appears in [15], and is used also by Deb et al. [7]. Chitra and Neuder [14] discuss restaking risk from a validator perspective, comparing restaking with investments in other financial instruments (e.g. bonds). Alexander [3] considers the interplay between existing leveraging schemes and liquid restaking tokens, which can amplify large-scale credit risk.

2 Model

Validators and Services

We consider a setting in which there is a set V of validators and a set S of services. Each service $s \in S$ has some profit from corruption π_s , and each validator $v \in V$ has some stake σ_v . We also associate with each service s a parameter α_s that denotes the fraction of stake required to corrupt/launch an attack on s . We call a bipartite graph $G = (S, V, E, \pi, \sigma, \alpha)$ a *restaking graph*; an edge is drawn between a validator $v \in V$ and a service $s \in S$ if v is restaking for s . For a given set of vertices A in a graph G , we use the notation $N_G(A)$ to denote the neighbors of A .

Attack Dynamics

For simplicity, we assume that validators lose their full stake σ_v if they launch an attack on a service. As such, for a given collection of services $A \subseteq S$, and a given collection of validators $B \subseteq V$ restaking for those services, we say that (A, B) is an *attacking coalition* for a restaking graph G if the validators in B possess enough stake to corrupt the services A :

$$\underbrace{\sum_{v \in B \cap N_G(\{s\})} \sigma_v}_{\text{Total stake in } s \text{ owned by validators } B} \geq \alpha_s \cdot \underbrace{\sum_{v \in N_G(\{s\})} \sigma_v}_{\text{Total amount restaked in } s} \quad \forall s \in A \quad (1)$$

We further say that (A, B) is a *valid attack* if it is an attacking coalition that has an incentive to launch an attack:

$$\underbrace{\sum_{s \in A} \pi_s}_{\text{Total profit from corrupting } A} > \underbrace{\sum_{v \in B} \sigma_v}_{\text{Total stake owned by validators } B} \quad (2)$$

If a valid attack (A, B) is carried out, we denote by $G \searrow B$ the induced subgraph $G[S, V \setminus B]$. The graph $G \searrow B$ denotes the state of the restaking graph after the attack is carried out. If no valid attacks exist on the graph, we call it *secure*. To simplify notation, for any subset of validators $B \subseteq V$, we will use the shorthand σ_B to denote $\sum_{v \in B} \sigma_v$. Similarly, for any $A \subseteq S$, we will use π_A as shorthand for $\sum_{s \in A} \pi_s$.

EigenLayer sufficient conditions

We note in passing that, in their whitepaper [15], EigenLayer proposes some efficiently verifiable sufficient conditions for security that they check to ensure that an attack does not exist.

▷ **Claim 1** (EigenLayer Sufficient Conditions, from Appendix B.1 of the EigenLayer Whitepaper [15]). A restaking graph G is secure if for each validator $v \in V$,

$$\sum_{s \in N_G\{v\}} \frac{\sigma_v}{\sigma_{N_G\{s\}}} \cdot \frac{\pi_s}{\alpha_s} \leq \sigma_v \quad (3)$$

Proof. It is shown in Appendix B.1 of the whitepaper that if Equation (3) holds for G , then the graph is secure (i.e. no valid attacks (A, B) satisfying Equations (1) and (2) exist). ◁

Cascading attacks

Our goal is to understand when a small shock can result in the loss of a large fraction of the overall stake. Small shocks can turn into large shocks by means of a cascading attack. Formally, we say that a disjoint sequence $(A_1, B_1), \dots, (A_T, B_T) \in 2^S \times 2^V$ is a *valid cascade of attacks* on a restaking graph $G = (S, V, E, \pi, \sigma, \alpha)$ if for each $t \in [T]$, (A_t, B_t) is a valid attack on $G \searrow \bigcup_{i=1}^{t-1} B_i$. We denote by $\mathcal{C}(G)$ the set of all such sequences of valid cascading attacks.

Worst-case stake loss

We now define a metric that measures the total potential loss of stake due to a sequence of cascading attacks. In our model, we first suppose that an initial small shock decreases the amount of stake. Formally, we define, for a given restaking graph $G = (S, V, E, \pi, \sigma, \alpha)$,

$$\mathbb{D}_\psi(G) := \left\{ D \subseteq V \mid \frac{\sigma_D}{\sigma_V} \leq \psi \right\} \quad (4)$$

to be the set of validator coalitions that constitute at most a ψ -fraction of all stake. Given some $D \in \mathbb{D}_\psi(G)$, we use the notation $G \searrow D := G[S, V \setminus D]$ to denote the induced subgraph of the restaking graph when we delete the validators D . We now define

$$R_\psi(G) := \underbrace{\psi}_{\text{Initial shock}} + \max_{D \in \mathbb{D}_\psi(G)} \max_{(A_1, B_1), \dots, (A_T, B_T) \in \mathcal{C}(G \searrow D)} \underbrace{\frac{\sigma_{\bigcup_{t=1}^T B_t}}{\sigma_V}}_{\text{Stake lost from cascading attacks}} \quad (5)$$

This quantity represents the worst-case total fraction of stake lost due to an initial ψ -fraction of the stake disappearing. By construction, $\psi \leq R_\psi(G) \leq 1$.

3 Overcollateralization Provides Robust Security

In this section, we show that “scaling up” the definition of security automatically results in robust security, meaning bounded losses from cascading attacks that follow an initial shock. We first show that, without loss of generality, it suffices to consider single valid attacks $(A, B) \in \mathcal{C}(G \searrow D)$ instead of more general cascading attacks.

► **Lemma 2.** *Let $G = (S, V, E, \pi, \sigma, \alpha)$ be an arbitrary restaking graph, and further suppose that (A, B) is an attacking coalition on $G \searrow D$, where $D \subseteq V$. Then, $(A, B \cup D)$ is an attacking coalition on G .*

48:10 Robust Restaking Networks

Proof. Because (A, B) is an attacking coalition on $G \searrow D$, we must have by Equation (1) that

$$\sigma_{B \cap N_G \{s\}} \geq \alpha_s \cdot \sigma_{N_G \{s\} \setminus D} \quad \forall s \in A \quad (6)$$

It follows that for any $s \in A$,

$$\sigma_{(B \cup D) \cap N_G \{s\}} = \sigma_{B \cap N_G \{s\}} + \sigma_{D \cap N_G \{s\}} \quad (7)$$

$$\geq \alpha_s \cdot \sigma_{N_G \{s\} \setminus D} + \sigma_{D \cap N_G \{s\}} \quad (8)$$

$$\geq \alpha_s \cdot \sigma_{N_G \{s\}} \quad (9)$$

and the desired result follows. \blacktriangleleft

► **Corollary 3.** *Let $G = (S, V, E, \pi, \sigma, \alpha)$ be an arbitrary restaking graph, and further suppose that $(A_1, B_1), \dots, (A_T, B_T) \in \mathcal{C}(G)$ is a valid sequence of cascading attacks on G . Then, $(\bigcup_{t=1}^T A_t, \bigcup_{t=1}^T B_t)$ is also a valid attack on G .*

Proof. By repeatedly applying Lemma 2, we find that for each $t \in [T]$, $(A_t, \bigcup_{i=1}^t B_t)$ is an attacking coalition on G . It follows by inspection of Equation (1) that we must therefore have that $(\bigcup_t A_t, \bigcup_t B_t)$ is an attacking coalition on G . To finish the result, it suffices to show that Equation (2) holds for this attacking coalition on the original graph G . This follows from the disjointness of the A_t 's and of the B_t 's:

$$\pi_{\bigcup_t A_t} = \sum_{t=1}^T \pi_{A_t} > \sum_{t=1}^T \sigma_{B_t} = \sigma_{\bigcup_t B_t} \quad (10)$$

where in the inner inequality we use that for each $t \in [T]$, $\pi_{A_t} > \sigma_{B_t}$ by Equation (2), as (A_t, B_t) is a valid attack on $G \searrow \bigcup_{i=1}^{t-1} B_i$. It follows that $(\bigcup_t A_t, \bigcup_t B_t)$ is a valid attack on G . \blacktriangleleft

Adding Multiplicative Slack

Our condition is given by adding multiplicative slack to Equation (2). Formally, we say that a restaking graph G is *secure with γ -slack* if for all attacking coalitions (A, B) on G ,

$$(1 + \gamma) \underbrace{\pi_A}_{\text{Total profit from corrupting } A} \leq \underbrace{\sigma_B}_{\text{Total stake owned by validators } B} \quad (11)$$

► **Theorem 4.** *Suppose that a restaking graph $G = (S, V, E, \pi, \sigma, \alpha)$ is secure with γ -slack for some $\gamma > 0$. Then, for any $\psi > 0$, $R_\psi(G) < \left(1 + \frac{1}{\gamma}\right) \psi$.*

Proof. Take any $\psi > 0$ and any $D \in \mathbb{D}_\psi(G)$ for some restaking graph G where (11) holds. Let $(A_1, B_1), \dots, (A_T, B_T) \in \mathcal{C}(G \searrow D)$ be arbitrary. Applying Corollary 3, we must have that $(\bigcup_t A_t, \bigcup_t B_t) \in \mathcal{C}(G \searrow D)$ as well. Defining $A := \bigcup_t A_t$ and $B := \bigcup_t B_t$, we must therefore have that (A, B) is an attacking coalition on $G \searrow D$, and furthermore that

$$\pi_A > \sigma_B \quad (12)$$

By Lemma 2, we must also have that $(A, B \cup D)$ is an attacking coalition on the original graph G . It then follows that as G is secure with γ -slack, Equation (11) must hold on $(A, B \cup D)$, whence

$$(1 + \gamma) \pi_A \leq \sigma_{B \cup D} = \sigma_B + \sigma_D \quad (13)$$

Putting this together with Equation (12), we find that

$$(1 + \gamma)\sigma_B < (1 + \gamma)\pi_A \leq \sigma_B + \sigma_D \leq \sigma_B + \psi \cdot \sigma_V \quad (14)$$

It follows that

$$\gamma \cdot \sigma_B < \psi \cdot \sigma_V \implies \frac{\sigma_B}{\sigma_V} < \frac{\psi}{\gamma} \quad (15)$$

$$\implies \psi + \frac{\sigma_B}{\sigma_V} < \left(1 + \frac{1}{\gamma}\right) \psi \quad (16)$$

As we took A_1, \dots, A_T to be arbitrary, we find that $R_\psi(G) < \left(1 + \frac{1}{\gamma}\right) \psi$, as desired. ◀

The EigenLayer sufficient conditions (3) can be similarly “scaled up” to yield efficiently checkable sufficient conditions for security with γ -slack (and hence, by Theorem 4, robustness to cascading attacks).

► **Corollary 5.** *Let G be a restaking graph such that, for all validators $v \in V$,*

$$\sum_{s \in N_G(\{v\})} \frac{\sigma_v}{\sigma_{N_G(\{s\})}} \cdot \frac{(1 + \gamma)\pi_s}{\alpha_s} \leq \sigma_v \quad (17)$$

Then, $R_\psi(G) < \left(1 + \frac{1}{\gamma}\right) \psi$.

Proof. This follows from Theorem 4 and Claim 1. Noting that the γ -slack condition holds precisely iff no valid attacks exist when profits from corruption π_s are inflated by a multiplicative factor of $(1 + \gamma)$, it suffices to apply EigenLayer’s sufficient conditions from Claim 1 with modified profits from corruption $(1 + \gamma)\pi_s$. ◀

Note that, given a restaking network, it is straightforward to compute the minimum value of γ such that the condition in (17) holds. This value can then be interpreted as an easily computed “risk measure” of such a network.

4 Lower Bounds for Global Security

In this section, we show that the upper bounds from the previous section are tight. We first show that if there is no multiplicative slack (i.e. $\gamma = 0$), then very small shocks can cause all stake to be lost in the worst case. This holds even under the EigenLayer conditions (3)⁴.

► **Theorem 6.** *For any $0 < \epsilon < 1$, there exists a restaking graph G that is secure and meets the EigenLayer condition (3), but has $R_\psi(G) = 1$ for all $\psi \geq \epsilon$.*

Proof. We construct a restaking graph $G = (S, V, E, \pi, \sigma, \alpha)$ with one service $S = \{x\}$ and two validators $V = \{a, b\}$, where an edge exists between each validator and the service (i.e. $E := (\{x\}, \{a\}), (\{x\}, \{b\})$). We then let $\sigma_a := \epsilon$, $\sigma_b := 1 - \epsilon$, $\pi_x := 1$, and $\alpha_x := 1$. Without loss of generality, we may assume $\psi < 1$ as $R_1(G) = 1$ for any restaking graph G . This graph satisfies (3) as

⁴ The graph exhibited in the proof of Theorem 6 has $\pi_x/\sigma_a \rightarrow \infty$ as $\epsilon \rightarrow 0$. It is possible to construct a counterexample with similar properties while maintaining that π_s/σ_v is greater than some universal constant for any $s \in S$ and $v \in V$. This is done in the full version of the paper.

48:12 Robust Restaking Networks

$$\frac{\sigma_a}{\sigma_a + \sigma_b} \cdot \frac{\pi_x}{\alpha_x} = \sigma_a \quad (18)$$

$$\frac{\sigma_b}{\sigma_a + \sigma_b} \cdot \frac{\pi_x}{\alpha_x} = \sigma_b \quad (19)$$

whence the graph is also secure by Claim 1. We now consider an initial shock $D = \{a\}$. As $\sigma_a/\sigma_V = \epsilon \leq \psi$, it follows that $D \in D_\psi(G)$. The pair $(\{x\}, \{b\})$ is a valid attack on $G \searrow D$, since $\{b\} = N_{G \searrow D} \{x\}$ whence it is an attacking coalition, and $\sigma_b < \pi_x$. It follows that $R_\psi(G) \geq \frac{\sigma_a + \sigma_b}{\sigma_V} = 1$ as desired. \blacktriangleleft

Next, we show that the bound we give in Theorem 4 (indeed, more strongly, the condition given in Corollary 5) is tight for all $\psi, \gamma > 0$.

► **Theorem 7.** *For any $\psi, \gamma, \epsilon > 0$ such that*

$$0 \leq \left(1 + \frac{1}{\gamma}\right) \psi - \epsilon \leq 1, \quad (20)$$

there exists a restaking graph G that satisfies the condition (17) from Corollary 5 but has $R_\psi(G) \geq \left(1 + \frac{1}{\gamma}\right) \psi - \epsilon$.

Proof. We construct a restaking graph $G = (S, V, E, \pi, \sigma, \alpha)$ with three validators $V = \{a, b, c\}$ and one service $S = \{x\}$, where an edge exists between each of the validators a and b and the service x (i.e. the edge set $E := \{(x, a), (x, b)\}$). Without loss of generality, suppose that $\epsilon \leq \psi/\gamma$. Let $\sigma_a > 0$ be any positive constant. We define

$$\sigma_b := \sigma_a \left(\frac{1}{\gamma} - \frac{\epsilon}{\psi} \right) \quad (21)$$

$$\sigma_c := \sigma_a \left(\frac{1 - \psi + \epsilon}{\psi} - \frac{1}{\gamma} \right) \quad (22)$$

$$\pi_x := \frac{\left(1 + \frac{1}{\gamma}\right) \psi - \epsilon}{1 + \gamma} \cdot \sigma_V \quad (23)$$

$$\alpha_s := 1 \quad (24)$$

Notice first that $\sigma_b \geq 0$ as we have taken $\epsilon \leq \psi/\gamma$. Next, observe that

$$\sigma_c \geq 0 \iff \frac{1 - \psi + \epsilon}{\psi} \geq \frac{1}{\gamma} \quad (25)$$

$$\iff 1 \geq \left(1 + \frac{1}{\gamma}\right) \psi - \epsilon \quad (26)$$

whence $\sigma_c \geq 0$ by Equation (20). Finally, we must also have that $\pi_x \geq 0$ by Equation (20) as well. Next, notice by construction that

$$\sigma_V = \sigma_a + \sigma_b + \sigma_c = \sigma_a \left(1 + \frac{1}{\gamma} - \frac{\epsilon}{\psi} + \frac{1 - \psi + \epsilon}{\psi} - \frac{1}{\gamma} \right) = \frac{\sigma_a}{\psi} \quad (27)$$

This graph meets condition (17) from Corollary 5, as

$$\sum_{s \in N_G \{a\}} \frac{\sigma_a}{\sigma_{N_G \{s\}}} (1 + \gamma) \pi_s = \frac{\sigma_a}{\sigma_a + \sigma_b} \left[\left(1 + \frac{1}{\gamma}\right) \psi - \epsilon \right] \sigma_V \quad (28)$$

$$= \frac{\sigma_a}{\sigma_a \left(1 + \frac{1}{\gamma} - \frac{\epsilon}{\psi}\right)} \left[\left(1 + \frac{1}{\gamma}\right) \psi - \epsilon \right] \frac{\sigma_a}{\psi} \quad (29)$$

$$= \sigma_a \quad (30)$$

A similar argument shows that

$$\sum_{s \in N_G \{b\}} \frac{\sigma_b}{\sigma_{N_G \{s\}}} (1 + \gamma) \pi_s = \frac{\sigma_b}{\sigma_a \left[\left(1 + \frac{1}{\gamma}\right) \psi - \epsilon \right]} \left[\left(1 + \frac{1}{\gamma}\right) \psi - \epsilon \right] \sigma_a = \sigma_b \quad (31)$$

Finally, as the validator c has no neighbors, it also satisfies the condition, whence the graph indeed satisfies (17). We now consider an initial shock $D := \{a\}$ to the graph. Because

$$\frac{\sigma_a}{\sigma_V} = \frac{\sigma_a}{\sigma_a / \psi} = \psi, \quad (32)$$

the shock $D \in \mathbb{D}_\psi$ constitutes a ψ fraction of the total stake as desired. The attack $(\{x\}, \{b\})$ is a valid attack on $G \searrow D$. To see this, notice first that $(\{x\}, \{b\})$ is an attacking coalition on $G \searrow D$ as $\{b\} = N_{G \searrow D} \{x\}$. Furthermore, $\{b\}$ is incentivized to attack since

$$\pi_x - \sigma_b = \frac{\left(1 + \frac{1}{\gamma}\right) \psi - \epsilon}{1 + \gamma} \cdot \sigma_V - \sigma_b \quad (33)$$

$$= \left[\frac{1 + \frac{1}{\gamma} - \frac{\epsilon}{\psi}}{1 + \gamma} - \left(\frac{1}{\gamma} - \frac{\epsilon}{\psi}\right) \right] \sigma_a \quad (34)$$

$$= \left[\frac{\left(1 + \frac{1}{\gamma}\right) \psi + \gamma \epsilon}{(1 + \gamma) \psi} - \frac{1}{\gamma} \right] \sigma_a \quad (35)$$

$$= \frac{\gamma \epsilon \sigma_a}{(1 + \gamma) \psi} > 0 \quad (36)$$

whence Equation (2) is satisfied for the pair $(\{x\}, \{b\})$. It follows that

$$R_\psi(G) \geq \frac{\sigma_a + \sigma_b}{\sigma_V} = \frac{\left(1 + \frac{1}{\gamma} - \frac{\epsilon}{\psi}\right) \sigma_a}{\sigma_a / \psi} = \left(1 + \frac{1}{\gamma}\right) \psi - \epsilon \quad (37)$$

as desired. ◀

5 Local Security

The bound in Theorem 4 on the worst-possible stake loss from cascading attacks is reassuring from a global perspective, but less so from the perspective of one or a small number of services who would like an assurance that they will not be among those affected by such attacks. Beginning with this section, we focus on a specific coalition of services $C \subseteq S$ that seeks to insulate their shared security $\Gamma(C)$ against shocks and resulting cascading attacks that may come about due to the decisions of other services and validators. Formally, we denote by

$$\Gamma(C) := \{v \in V \mid N_G \{v\} \subseteq C\} \quad (38)$$

the set of validators that exclusively provide security for services in C .

Worst-case stake loss (local version)

As before, we first suppose that an initial shock affects the restaking graph. Whereas previously, we considered shocks for which the total stake in the shock was bounded, we now consider shocks for which the total stake that impacts the exclusive security of some coalition of services C (i.e., stake that secures services from C and only services from C) is bounded. Formally, for any coalition of services $C \subseteq S$ within some restaking graph G , we let

$$\mathbb{D}_\psi(C, G) := \left\{ D \subseteq V \mid \frac{\sigma_{D \cap \Gamma(C)}}{\sigma_{\Gamma(C)}} \leq \psi \right\} \quad (39)$$

denote the set of all validator coalitions that provide at most ψ stake to the aggregate security of the coalition of services C . Notice that shocks $D \in \mathbb{D}_\psi(C, G)$ may have much more total stake σ_D than a ψ fraction of the graph. We are instead only guaranteed that the impact of the shock on stake that is being used exclusively for members in C is small. We are now interested in the potential cascading losses that can affect the stake that is exclusively utilized by the coalition C after a shock occurs that destroys at most ψ stake from the aggregate security of C . Formally, we study the quantity

$$R_\psi(C, G) := \underbrace{\psi}_{\text{Initial Shock}} + \max_{D \in \mathbb{D}_\psi(C, G)} \max_{(A_1, B_1), \dots, (A_T, B_T) \in \mathcal{C}(G \setminus D)} \underbrace{\frac{\sigma_{\bigcup_{t=1}^T B_t \cap \Gamma(C)}}{\sigma_{\Gamma(C)}}}_{\text{Stake lost from cascading attacks}} \quad (40)$$

Local security conditions

We seek sufficient conditions on a restaking graph G that guarantee a nontrivial upper bound on $R_\psi(C, G)$. Ideally, the sufficient condition would depend only the neighborhood/choices of the coalition C to provide a guarantee that holds regardless of the choices made by other validators and services (i.e., the services of C can attempt to “control their own destiny” by ensuring that the locally defined sufficient condition holds). Formally, given a restaking graph $G = (S, V, E, \pi, \sigma, \alpha)$ and a coalition of services $C \subseteq S$, we call a restaking graph $G' = (S', V', E', \pi', \sigma', \alpha')$ a *C-local variant* of G if C cannot distinguish G' from G on the basis of local information: $C \subseteq S'$, $N_G C = N_{G'} C$, and

$$(\pi_s, \alpha_s) = (\pi'_s, \alpha'_s) \quad \forall s \in C \quad (41)$$

$$(\sigma_v, N_G \{v\}) = (\sigma'_v, N_{G'} \{v\}) \quad \forall v \in N_G C \quad (42)$$

We then define a *local security condition* $f : (C, G) \mapsto \{0, 1\}$ to be a Boolean function that takes as input a restaking graph G and a coalition of services $C \subseteq S$ such that $f(C, G)$ must be equal to $f(C, G')$ for all C -local variants G' . The intuition behind this definition is that the condition should only depend on service-level information (e.g. profits from corruption, security thresholds) for services in the coalition, and validator-level information for validators in $N_G C$.

Local security impossibility

Unfortunately, without further restrictions on the attacks under consideration (like those defined later in this section), it is impossible to construct any nontrivial local security condition that yields any nontrivial upper bound on $R_\psi(C, G)$.

► **Theorem 8.** *For any local security condition f , any secure restaking graph G and coalition of services $C \subseteq S$ such that $f(C, G) = 1$, there exists a secure C -local variant G' of G such that $R_0(C, G') = 1$.*

Proof. Take any f , C , and secure G such that $f(C, G) = 1$. Define

$$\Delta := \sigma_{N_G C} - \pi_C \quad (43)$$

to be the total overcollateralization of C in aggregate. Next, we define G' to be an augmented version of G , where we add a new service s^* that has a profit from corruption $\pi_{s^*} = \Delta + 2\epsilon$ where $\epsilon > 0$. We further add 2 validators a and b to the graph who are adjacent only to s^* . We let $\sigma_a := \Delta + \epsilon$ and $\sigma_b := \epsilon$. As $\sigma_a + \sigma_b \geq \pi_{s^*}$, the graph G' must be secure as G was secure. Next, notice that as the validator a is not path-connected to C , G' must be a C -local variant of G . However, by construction, the attack $(C \cup \{s^*\}, N_G C \cup \{b\})$ is valid on the graph $G' \searrow \{a\}$. It follows that $R_0(C, G') = 1$. ◀

Stable attacks

While the above impossibility appears to be quite strong, it is somewhat contrived. At the heart of the impossibility is that under the definition of a valid attack (i.e. Equations (1) and (2)), not every validator must be productive in carrying out the attack. There may be a subset of validators in the attack that can yield more net profit than the full coalition. In what follows, we show that if we assume that malicious validator coalitions will choose to add others to their ranks only if it is profitable for them in net to do so, then a local security condition with guarantees similar to those in Theorem 4 does indeed exist. Formally, for $A \subseteq S$ and $B \subseteq V$, we say that an attack (A, B) is *stable* if it is valid (i.e. Equations (1) and (2) hold), and for all $A' \subseteq A$ and $B' \subseteq B$ such that (A', B') is valid,

$$\sigma_{B \setminus B'} < \pi_{A \setminus A'} \quad (44)$$

Intuitively, if (44) did not hold, then the validators of B' would be better off ditching those in $B \setminus B'$ and attacking only the services in A' .

We say that a disjoint sequence $(A_1, B_1), \dots, (A_T, B_T) \in 2^S \times 2^V$ is a cascade of stable attacks on a restaking graph $G = (S, V, E, \pi, \sigma, \alpha)$ if for each $t \in [T]$, (A_t, B_t) is a stable attack on $G \searrow \bigcup_{i=1}^{t-1} B_i$. We denote by $\mathcal{S}(G)$ the set of all such sequences of cascading stable attacks. In light of Theorem 8, we redefine our notion of worst-case stake loss using stable attacks:

$$R_\psi(C, G) := \psi + \max_{D \in \mathbb{D}_\psi(C, G)} \max_{(A_1, B_1), \dots, (A_T, B_T) \in \mathcal{S}(G \searrow D)} \frac{\sigma_{\bigcup_t B_t \cap \Gamma(C)}}{\sigma_{\Gamma(C)}} \quad (45)$$

Unlike valid attacks, unions of sequences of stable cascading attacks need not be stable (which will complicate the proof of Theorem 10 in the next section).

▷ **Claim 9.** There exists a restaking graph G and a sequence $(A_1, B_1), (A_2, B_2) \in \mathcal{S}(G)$ such that $(A_1 \cup A_2, B_1 \cup B_2) \notin \mathcal{S}(G)$.

Proof. Consider restaking graph where there are two services $S = \{x, y\}$ and two validators $V = \{a, b\}$ where both validators are restaking in both services. Furthermore, $\pi_x = \pi_y = 2$, $\alpha_x = \alpha_y = \frac{1}{2}$, and $\sigma_a = \sigma_b = 1$. In this case, the sequence $(\{x\}, \{a\}), (\{y\}, \{b\}) \in \mathcal{S}(G)$. However, $(\{x, y\}, \{a, b\}) \notin \mathcal{S}(G)$ because the attack $(\{x, y\}, \{a\})$ is a valid attack. ◀

6 A Local Security Condition for Stable Attacks

In this section, we give a family of local security conditions that yield guarantees on the local worst-case stake loss $R_\psi(C, G)$ that resemble our result from Theorem 4 for global security. In Theorems 4 and 7, we showed that security with γ -slack was necessary and sufficient in order to obtain a $(1 + \frac{1}{\gamma})\psi$ upper-bound on $R_\psi(G)$. In other words, it was both necessary and sufficient to ensure that all attacking coalitions (A, B) were overcollateralized

multiplicatively by a factor of $(1 + \gamma)$. Our condition for local security is similar: we must ensure that certain *attack headers* (X, Y) are overcollateralized multiplicatively by a factor of $(1 + \gamma)$.

Attack header

Formally, we say that (X, Y) is an *attack header*, for $X \subseteq S$ and $Y \subseteq \Gamma(X)$, if there exists a set of validators $B \subseteq V$ satisfying

$$B \cap \Gamma(X) = \emptyset \quad (46)$$

such that $(X, B \cup Y)$ is an attacking coalition. In other words, (X, Y) is an attack header if Y can be appended to a collection of validators B that may be slashed without attacking services in X to form an attacking coalition that attacks the services X . An attacking coalition is a special case of an attack header (in which B can be taken as \emptyset).

► **Theorem 10.** *Let $G = (S, V, E, \pi, \sigma, \alpha)$ be a restaking graph and $C \subseteq S$ be a coalition of services. If, for all attack headers (X, Y) where $X \subseteq C$,*

$$(1 + \gamma)\pi_X \leq \sigma_Y \quad (47)$$

then $R_\psi(C, G) < (1 + \frac{1}{\gamma})\psi$. Furthermore, the Boolean function that checks whether Equation (47) holds for all attack headers is a local security condition.

Proof. Let $D \in \mathbb{D}_\psi(C, G)$ and $(A_1, B_1), \dots, (A_T, B_T) \in \mathcal{S}(G \searrow D)$ be arbitrary. For each $t \in [T]$, define

$$L_t := B_t \cap \Gamma(C) \quad (48)$$

to be the set of all validators exclusively securing C that were lost in the t^{th} attack. Next, define

$$A'_t := \left\{ s \in A_t \mid \sigma_{B_t \setminus L_t} \geq \alpha_s \cdot \sigma_{N_G\{s\} \setminus (\bigcup_{i=1}^{t-1} B_i \cup D)} \right\} \quad (49)$$

to be the maximal set of services such that $(A'_t, B_t \setminus L_t)$ is an attacking coalition on $G \searrow (\bigcup_{i=1}^{t-1} B_i \cup D)$. Notice also that because (A_t, B_t) is a stable attack on $G \searrow (\bigcup_{i=1}^{t-1} B_i \cup D)$, we must have that $A_t \setminus A'_t$ is nonempty, and furthermore must be a subset of C .

▷ **Claim 11.**

$$\sigma_{\bigcup_t B_t \cap \Gamma(C)} < \pi_{\bigcup_t A_t \setminus A'_t} \quad (50)$$

Proof. From stability, we have that

$$\sigma_{L_t} = \sigma_{B_t \setminus (B_t \setminus L_t)} < \pi_{A_t \setminus A'_t} \quad (51)$$

To see why this holds, notice that there are two cases. If the attacking coalition $(A'_t, B_t \setminus L_t)$ is a valid attack, then the inequality follows directly from the stability definition. If instead $(A'_t, B_t \setminus L_t)$ is not a valid attack despite being an attacking coalition, the inequality must still hold because the original attack (A_t, B_t) is valid and therefore satisfies Equation (2). Iterating over t , we find that

$$\sigma_{\bigcup_t B_t \cap \Gamma(C)} = \sum_{t=1}^T \sigma_{L_t} < \sum_{t=1}^T \pi_{A_t \setminus A'_t} = \pi_{\bigcup_t A_t \setminus A'_t} \quad (52)$$

◁

▷ Claim 12. $(\bigcup_t (A_t \setminus A'_t), (\bigcup_t B_t \cup D) \cap \Gamma(C))$ is an attack header on G , and $\bigcup_t (A_t \setminus A'_t) \subseteq C$.

Proof. Because each $A_t \setminus A'_t \subseteq C$, we must also have that $\bigcup_t A_t \setminus A'_t \subseteq C$ as well. Next, by applying Corollary 3 noting the disjointness of the (A_t, B_t) , we have that

$$\left(\bigcup_t A_t, \bigcup_t B_t \right) = \left(\bigcup_t (A_t \setminus A'_t) \cup \bigcup_t A'_t, \bigcup_t B_t \right) \quad (53)$$

must be an attacking coalition on $G \searrow D$. In particular, we must have that $(\bigcup_t (A_t \setminus A'_t), \bigcup_t B_t)$ is an attacking coalition on $G \searrow D$, whence by Lemma 3 we have that $(\bigcup_t (A_t \setminus A'_t), \bigcup_t B_t \cup D)$ is an attacking coalition on G . Rewriting the above as

$$\left(\bigcup_t (A_t \setminus A'_t), \left[\left(\bigcup_t B_t \cup D \right) \cap \Gamma(C) \right] \cup \left[\left(\bigcup_t B_t \cup D \right) \setminus \Gamma(C) \right] \right) \quad (54)$$

and noting by construction that

$$\left[\left(\bigcup_t B_t \cup D \right) \setminus \Gamma(C) \right] \cap \Gamma(C) = \emptyset, \quad (55)$$

we find that $(\bigcup_t (A_t \setminus A'_t), (\bigcup_t B_t \cup D) \cap \Gamma(C))$ is an attack header on G as desired. ◀

Putting these claims together yields the desired result. By Claim 12 and Equation (47), we find that

$$(1 + \gamma)\pi_{\bigcup_t A_t \setminus A'_t} \leq \sigma_{(\bigcup_t B_t \cup D) \cap \Gamma(C)} \leq \sigma_{\bigcup_t B_t \cap \Gamma(C)} + \psi \cdot \sigma_{\Gamma(C)} \quad (56)$$

Adding in Claim 11, we find that

$$(1 + \gamma)\sigma_{\bigcup_t B_t \cap \Gamma(C)} < \sigma_{\bigcup_t B_t \cap \Gamma(C)} + \psi \cdot \sigma_{\Gamma(C)} \quad (57)$$

$$\implies \frac{\sigma_{\bigcup_t B_t \cap \Gamma(C)}}{\sigma_{\Gamma(C)}} < \frac{\psi}{\gamma} \quad (58)$$

$$\implies R_\psi(C, G) < \left(1 + \frac{1}{\gamma}\right)\psi \quad (59)$$

as desired. To see that the Boolean function that checks whether Equation (47) holds for all attack headers is a local security condition, observe that it suffices to check that for all $X \subseteq C$ and $Y \subseteq N_G C \cap \Gamma(C) = \Gamma(C)$ such that $(X, Y \cup N_G C \setminus \Gamma(C))$ is an attacking coalition, Equation (47) holds. Thus, for any restaking graph G and C -local variant G' , this function will evaluate to the same output. ◀

The condition in (47) can be checked via enumeration, although the time required to do so grows exponentially in the number of services in C and the number of validators that contribute security exclusively to services in C . In the event that this is a prohibitive amount of computation, the same guarantee holds under a stronger, easily checked local analog of the EigenLayer sufficient condition (3) which, in effect, treats as malicious all validators that contribute security to any services outside of C .

48:18 Robust Restaking Networks

► **Corollary 13.** *For any restaking graph $G = (S, V, E, \pi, \sigma, \alpha)$ and coalition of services $C \subseteq S$, satisfaction of the local condition*

$$\sum_{s \in N_G\{v\}} \frac{\sigma_v}{\sigma_{N_G\{s\}}} \cdot \frac{(1 + \gamma)\pi_s}{\alpha'_s} \leq \sigma_v \quad (60)$$

for all $v \in N_G C$ with $N_G\{v\} \subseteq C$ guarantees that $R_\psi(C, G) < \left(1 + \frac{1}{\gamma}\right)\psi$, where for each $s \in C$,

$$\alpha'_s := \alpha_s - \frac{\sigma_{N_G\{s\} \setminus \Gamma(C)}}{\sigma_{N_G\{s\}}} \quad (61)$$

Proof. This follows from Theorem 10 and Claim 1. Notice that to guarantee that the condition (47) from Theorem 10 holds for all attack headers (X, Y) , it suffices to guarantee that no valid attacks exist, when (i) all profits from corruption for services in C are inflated by a multiplicative factor of $(1 + \gamma)$, and (ii) the fraction of stake required to corrupt a given service is offset by the fraction of stake in that service that is also restaking for services that do not belong to C . As Claim 1 provides a sufficient condition to guarantee the existence of no valid attacks, Equation (60) will guarantee that all attack headers are overcollateralized by a multiplicative factor of $(1 + \gamma)$. ◀

7 Lower Bounds for Local Security

We next show senses in which Theorem 10 (and more strongly, Corollary 13) is tight.

► **Corollary 14.** *For any $\psi, \gamma, \epsilon > 0$ such that*

$$0 \leq \left(1 + \frac{1}{\gamma}\right)\psi - \epsilon \leq 1, \quad (62)$$

there exists a restaking graph G that satisfies the condition (60) from Corollary 13 but has $R_\psi(C, G) \geq \left(1 + \frac{1}{\gamma}\right)\psi - \epsilon$.

Proof. Notice that if we take $C = S$, the condition (60) from Corollary 13 is identical to the condition (17) from Corollary 5. Furthermore, $R_\psi(S, G)$ is the same as $R_\psi(G)$ except for the fact that $R_\psi(S, G)$ considers only stable attacks. Repeating the argument from Theorem 7, and noting that the attack given in the proof of that result is stable, we obtain the desired result. ◀

► **Theorem 15.** *For any $\gamma > 0$, there exists a graph restaking graph $G = (S, V, E, \pi, \sigma, \alpha)$ that satisfies (17) from Corollary 5, but there exists a $C \subseteq S$ such that $R_0(C, G) = 1$.*

Proof. Let there be three services $S = \{x, y, z\}$, three validators $V = \{a, b, c\}$, and edges as follows:

$$E := \{(x, a), (x, b), (y, b), (y, c), (z, c), (z, a)\}. \quad (63)$$

Next, let $\alpha_s = 1$ for all $s \in S$, and let $\pi_x = \pi_y = \pi_z =: \pi > 0$ be an arbitrary positive constant. Next pick σ_a such that

$$\sigma_a < 2\pi, \quad (64)$$

let $\sigma_b := 2(1 + \gamma)\pi$, and let $\sigma_c := 2(1 + \gamma)\pi$. This graph satisfies (17) from Corollary 5 as

$$\frac{\sigma_a}{\sigma_a + \sigma_b} \cdot (1 + \gamma)\pi + \frac{\sigma_a}{\sigma_a + \sigma_c} \cdot (1 + \gamma)\pi < (1 + \gamma)\pi \left(\frac{1}{\sigma_b} + \frac{1}{\sigma_c} \right) \sigma_a = \sigma_a \quad (65)$$

$$\frac{\sigma_b}{\sigma_a + \sigma_b} \cdot (1 + \gamma)\pi + \frac{\sigma_b}{\sigma_b + \sigma_c} \cdot (1 + \gamma)\pi < \frac{3}{2}(1 + \gamma)\pi < \sigma_b \quad (66)$$

$$\frac{\sigma_c}{\sigma_a + \sigma_c} \cdot (1 + \gamma)\pi + \frac{\sigma_c}{\sigma_b + \sigma_c} \cdot (1 + \gamma)\pi < \frac{3}{2}(1 + \gamma)\pi < \sigma_c \quad (67)$$

Next, let $C := \{x, z\}$, and observe that the shock $D = \{b, c\}$ is an element of $\mathbb{D}_0(C, G)$. However, because $\sigma_a < \pi_x + \pi_z = 2\pi$, we must have that $(\{x, z\}, \{a\})$ is a stable attack on $G \searrow D$. As $\{a\} = \Gamma(C)$, it follows that $R_0(C, G) = 1$. \blacktriangleleft

8 Long Cascades

Cascade Structure

Although Corollary 3 shows that all long cascades can be made into a short, one-step attack, it is arguably more dangerous if large attacks can be made through long cascades of small attacks that each require less coordination. In this section, we show how adding γ -slack also enables us to upper-bound the length of a cascade in the worst case. Our results depend on what we call the *reference depth* of a cascading sequence of attacks. Formally, for a given restaking graph G and coalition of validators B_0 , we say that a valid cascade of attacks $(A_1, B_1), \dots, (A_T, B_T)$ on $\mathcal{C}(G \searrow B_0)$ has reference depth k if

$$k = \max \{i \in [T] \mid \exists t \in [T] \text{ s.t. } N_G(A_t) \cap B_{t-i} \neq \emptyset\} \quad (68)$$

In other words, a cascading sequence has reference depth k if the services attacked in a given time step are not affected by validators that were slashed more than k steps previous.

► **Theorem 16.** *Suppose that a restaking graph $G = (S, V, E, \pi, \sigma, \alpha)$ is secure with γ -slack for some $\gamma > 0$. Let $\epsilon = \min_{v \in V} \sigma_v$ denote the minimum stake held by a validator. Then, for any $\psi > 0$, $B_0 \in \mathbb{D}_\psi(G)$, and $(A_1, B_1), \dots, (A_T, B_T) \in \mathcal{C}(G \searrow B_0)$ with reference depth k ,*

$$T < k \left(1 + \log_{1+\gamma} \frac{\psi \cdot \sigma_V}{\epsilon \gamma} \right) \quad (69)$$

Proof. Because $(A_1, B_1), \dots, (A_T, B_T) \in \mathcal{C}(G \searrow B_0)$ has reference depth k , it follows that if we define

$$A'_i := \bigcup_{t=k(i-1)+1}^{ki} A_t \quad (70)$$

$$B'_i := \bigcup_{t=k(i-1)+1}^{ki} B_t \quad (71)$$

$$B'_0 := B_0 \quad (72)$$

for $i \in \{1, \dots, \lceil T/k \rceil\}$, we must have that for all $j < i - 1$,

$$N_G A'_i \cap B'_j = \emptyset \quad (73)$$

48:20 Robust Restaking Networks

Furthermore, by Corollary 3, we must have that the sequence $(A'_i, B'_i) \in \mathcal{C}(G \searrow B_0)$. Applying Corollary 3 again, we further have that for every i , $(\bigcup_{j=i+1}^{\lceil T/k \rceil} A'_j, \bigcup_{j=i+1}^{\lceil T/k \rceil} B'_j)$ is a valid attack on $G \searrow \bigcup_{j=0}^i B'_j$. It follows by Equation (2) that for any $i \in \{0, \dots, \lceil T/k \rceil - 1\}$,

$$\pi_{\bigcup_{j=i+1}^{\lceil T/k \rceil} A'_j} > \sigma_{\bigcup_{j=i+1}^{\lceil T/k \rceil} B'_j} = \sum_{j=i+1}^{\lceil T/k \rceil} \sigma_{B'_j} \quad (74)$$

Next, noting that

$$N_G \bigcup_{j=i+1}^{\lceil T/k \rceil} A'_j \cap \bigcup_{j=1}^{i-1} B'_j = \emptyset \quad (75)$$

by Equation (73), it follows that indeed $(\bigcup_{j=i+1}^{\lceil T/k \rceil} A'_j, \bigcup_{j=i+1}^{\lceil T/k \rceil} B'_j)$ is also a valid attack on $G \searrow B'_i$. Applying Lemma 2, we then find that $(\bigcup_{j=i+1}^{\lceil T/k \rceil} A'_j, B'_i \cup \bigcup_{j=i+1}^{\lceil T/k \rceil} B'_j)$ is an attacking coalition on the original graph G . Because G is secure with γ -slack, Equation (11) must now hold on this pair, whence for all $i \in \{0, \dots, \lceil T/k \rceil - 1\}$,

$$(1 + \gamma)\pi_{\bigcup_{j=i+1}^{\lceil T/k \rceil} A'_j} \leq \sigma_{\bigcup_{j=i}^{\lceil T/k \rceil} B'_j} = \sigma_{B'_i} + \sum_{j=i+1}^{\lceil T/k \rceil} \sigma_{B'_j} \quad (76)$$

Putting this together with Equation (74), we have that for all $i \in \{0, \dots, \lceil T/k \rceil - 1\}$,

$$(1 + \gamma) \sum_{j=i+1}^{\lceil T/k \rceil} \sigma_{B'_j} < \sigma_{B'_i} + \sum_{j=i+1}^{\lceil T/k \rceil} \sigma_{B'_j} \quad (77)$$

$$\implies \sigma_{B'_i} > \gamma \sum_{j=i+1}^{\lceil T/k \rceil} \sigma_{B'_j} \quad (78)$$

It can be shown inductively from Equation (78) that the sequence $\sigma_{B'_i} > X_i$, where we define the sequence X_i by

$$X_{\lceil T/k \rceil} := \epsilon \quad (79)$$

$$X_i := \gamma \sum_{j=i+1}^{\lceil T/k \rceil} X_j \quad i \in \{0, \dots, \lceil T/k \rceil - 1\} \quad (80)$$

Solving, we find that

$$X_0 = \epsilon \gamma (1 + \gamma)^{\lceil T/k \rceil - 1} \quad (81)$$

It then follows as desired that

$$\psi \cdot \sigma_V = \sigma_{B_0} > \epsilon \gamma (1 + \gamma)^{\lceil T/k \rceil - 1} \implies T < k \left(1 + \log_{1+\gamma} \frac{\psi \cdot \sigma_V}{\epsilon \gamma} \right) \quad (82)$$

◀

References

- 1 Daron Acemoglu, Asuman Ozdaglar, and Alireza Tahbaz-Salehi. Networks, shocks, and systemic risk. Technical report, National Bureau of Economic Research, 2015.
- 2 Daron Acemoglu, Asuman Ozdaglar, and Alireza Tahbaz-Salehi. Systemic risk and stability in financial networks. *American Economic Review*, 105(2):564–608, 2015.
- 3 Carol Alexander. Leveraged restaking of leveraged staking: What are the risks? *Available at SSRN 4840805*, 2024.
- 4 Stefano Battiston, Guido Caldarelli, Robert M May, Tarik Roukny, and Joseph E Stiglitz. The price of complexity in financial networks. *Proceedings of the National Academy of Sciences*, 113(36):10031–10036, 2016.
- 5 Markus K Brunnermeier, Gary Gorton, and Arvind Krishnamurthy. Risk topography. *Nber macroeconomics annual*, 26(1):149–176, 2012.
- 6 Chen Chen, Garud Iyengar, and Ciamac C Moallemi. An axiomatic approach to systemic risk. *Management Science*, 59(6):1373–1388, 2013. doi:10.1287/MNSC.1120.1631.
- 7 Soubhik Deb, Robert Raynor, and Sreeram Kannan. Stakesure: Proof of stake mechanisms with strong cryptoeconomic safety. *arXiv preprint*, 2024. doi:10.48550/arXiv.2401.05797.
- 8 Douglas W Diamond and Philip H Dybvig. Bank runs, deposit insurance, and liquidity. *Journal of political economy*, 91(3):401–419, 1983.
- 9 Larry Eisenberg and Thomas H Noe. Systemic risk in financial systems. *Management Science*, 47(2):236–249, 2001. doi:10.1287/MNSC.47.2.236.9835.
- 10 Zachary Feinstein, Birgit Rudloff, and Stefan Weber. Measures of systemic risk. *SIAM Journal on Financial Mathematics*, 8(1):672–708, 2017. doi:10.1137/16M1066087.
- 11 Paul Glasserman and H Peyton Young. Contagion in financial networks. *Journal of Economic Literature*, 54(3):779–831, 2016.
- 12 Subhash Khot and Rishi Saket. Hardness of bipartite expansion. In *24th Annual European Symposium on Algorithms (ESA 2016)*. Schloss-Dagstuhl – Leibniz Zentrum für Informatik, 2016. doi:10.4230/LIPIcs.ESA.2016.55.
- 13 Eduard Kromer, Ludger Overbeck, and Katrin Zilch. Systemic risk measures on general measurable spaces. *Mathematical Methods of Operations Research*, 84:323–357, 2016. doi:10.1007/S00186-016-0545-1.
- 14 Tarun Chitra Mike Neuder. The risks of lrts, 2024. URL: <https://ethresear.ch/t/the-risks-of-lrts/18799>.
- 15 EigenLayer Team. Eigenlayer: The restaking collective, 2023. URL: <https://docs.eigenlayer.xyz/overview/whitepaper>.