Privacy-Preserving Systemic Portfolio Analysis and Repair

Martin Zbudila*, Roozbeh Sarenche*, Aysajan Abidin, and Bart Preneel

COSIC, KU Leuven, Belgium
firstname.lastname@esat.kuleuven.be

Abstract. Financial institutions can be exposed to systemic risk through overlapping portfolios and common liquidity shocks, yet identifying and reducing these dependencies requires precisely the portfolio information that institutions are unwilling to reveal. We present, to the best of our knowledge, the first privacy-preserving framework combining system-level portfolio analysis with coordinated portfolio repair. Our approach uses secure multi-party computation (MPC), allowing mutually distrustful institutions to jointly identify systemic vulnerabilities and determine concrete portfolio adjustments without disclosing their underlying holdings. We provide both an MPC realization and a federated variant that moves fund-local computation outside MPC while retaining secure global coordination, and we analyze the additional leakage introduced by this federated design. We implement and evaluate the complete framework on real, publicly available portfolios. Across representative network profiles, private analysis completes in roughly 1–4 minutes and portfolio repair takes 2.6–8.8 hours in MPC and 13–21 minutes in the federated setting. Our results show that both private systemic-risk analysis and coordinated intervention are practically feasible with current MPC technology.

Keywords: Privacy-preserving finance, multi-party computation, portfolio repair, hedge funds, fire-sale risk, federated optimization

1 Introduction

Financial institutions do not manage risk in isolation. Funds, asset managers, banks, and other institutional investors may hold overlapping portfolios or exposures to assets whose returns and liquidity are strongly related. Consequently, actions taken by one institution can affect others even without direct contractual links. During periods of stress, simultaneous liquidations can depress asset prices, impose mark-to-market losses on common holders, and amplify losses across the financial system. Understanding such dependencies therefore requires a system-wide view of portfolios rather than an institution-by-institution analysis.

This system-wide view is useful both for *measuring* risk and for *reducing* it. Joint portfolio information can reveal dependencies, crowded exposures, and structural concentrations that are not visible from an individual portfolio alone.

* These authors contributed equally.

It can also help identify which institutions are more strongly connected to the wider system and how losses may propagate when several institutions are exposed to the same market stress. Beyond identifying such vulnerabilities, a system-wide view can support coordinated preventative action, allowing institutions to adjust their portfolios and reduce the losses that could arise from future common shocks.

Obtaining this system-wide view, however, creates a fundamental confidentiality problem. Hedge funds provide a particularly clear example: their portfolio holdings and trading strategies are commercially sensitive and generally cannot be shared with competitors. This concern is well documented in the hedge-fund literature: funds seek confidential treatment precisely when disclosure could reveal proprietary positions or expose them to front-running and adverse price impact [6,12]. Yet the quantities needed for system-wide analysis and coordinated action depend on information that is distributed across multiple private portfolios. Conventional approaches do not resolve this tension: centralizing the portfolios requires trusting a data holder and creates a sensitive repository, while purely local analysis preserves confidentiality but cannot provide the joint information needed for system-wide analysis and coordination.

In this paper, we study how to bridge this gap. We develop a privacy-preserving framework in which multiple financial institutions can jointly analyze their portfolio dependencies and coordinate portfolio adjustments without revealing their underlying holdings to one another. Our primary construction uses *secure multi-party computation* (MPC), where private inputs and intermediate values remain secret-shared while the required joint computation is evaluated securely. This provides a strong privacy boundary, but secure optimization is substantially more expensive than its plaintext counterpart and becomes increasingly costly as the system grows. We therefore also investigate a complementary *federated optimization* architecture, where fund-local computation is performed in plaintext while the globally coupled coordination is jointly realized through MPC.

Our framework is designed to support both portfolio *analysis* and *rebalancing*. On the analysis side, it enables participating institutions to privately derive portfolio correlations and network structure, including a Minimum Spanning Tree (MST) that provides a sparse representation of the strongest dependencies among institutions, together with centrality measures and systemic-risk statistics. On the rebalancing side, it jointly determines portfolio adjustments and contingency liquidations that reduce modeled fire-sale losses while respecting institution-level portfolio constraints, turnover limits, and fairness requirements. Thus, rather than only identifying where systemic vulnerabilities exist, the framework also computes concrete measures for reducing them.

Our main contributions are as follows:

- **MPC-friendly systemic portfolio analysis.** We develop private protocols for portfolio dependence and systemic-risk analysis, including a division- and square-root-free formulation of correlation comparisons.
- **Systemic portfolio-repair formulation.** We formulate an optimization problem that jointly determines portfolio adjustments and stress-contingent liquidations to achieve a prescribed systemic-risk reduction at minimum eco-

nomic cost, while guaranteeing individual improvements and enforcing fairness through a Shapley-based notion of structural responsibility.
- **Privacy-preserving portfolio repair.** We formulate cooperative portfolio repair as a bilinear optimization problem and adapt an iterative first-order method to solve its linear subproblems efficiently under MPC.
- **Federated repair.** We develop a federated variant that moves fund-local optimization outside MPC while retaining secure global aggregation and coordination. We discuss the additional leakage from iterative coordination signals and illustrate the effect of collusion empirically.
- **End-to-end implementation and evaluation.** We implement the complete framework and evaluate its accuracy, economic effectiveness, runtime, communication, and scalability on realistic portfolio data, comparing pure-MPC and federated repair against plaintext reference solutions. Our implementation is available at [10].

2 Related Work

Privacy-preserving computation in finance has received increasing attention, including applications of MPC, homomorphic encryption, federated learning, and related techniques [43]. Abbe et al. [5] use secure computation to derive aggregate risk statistics, including concentration measures and correlations, from private financial exposures, while Hastings et al. [31] develop privacy-preserving analytics for financial networks and systemic-risk assessment. Other work considers privacy-preserving portfolio pricing [13] and investor-level portfolio optimization [18,36,37]. More generally, secure linear optimization has been studied through MPC-based simplex methods [52,21] and privacy-preserving problem transformations [25,16], while practical deployments have demonstrated MPC on real economic and financial data [17]. Our work bridges two previously separate directions: privacy-preserving system-level analysis and systemic portfolio reallocation. We use MPC both to derive cross-fund dependencies from private holdings and to compute coordinated portfolio adjustments that reduce modeled fire-sale spillovers.

Correlation and network methods are established tools for analyzing financial dependence [42], and privacy-preserving graph computation has been studied for secure combinatorial algorithms and MSTs [9,38]. Our construction exploits the algebraic structure of portfolio correlation to obtain the MST edge ordering without explicitly evaluating pairwise divisions or square roots. On the financial side, Poledna et al. [46] quantify systemic risk from overlapping portfolios using granular security holdings, while Barucca et al. [14] show that common holdings across financial institutions are informative about vulnerability to fire sales. Such analyses generally assume that the relevant portfolio holdings are available to the analyst, although portfolio disclosure can reveal proprietary strategies and expose funds to front-running and adverse price impact [6,12,44]. Our protocols instead derive the required joint statistics without disclosing the underlying portfolios.

Portfolio reallocation has likewise been studied as a mechanism for reducing systemic risk. Pichler et al. [45] optimize overlapping bank portfolios to reduce

systemic risk while preserving institution-level objectives; related work examines trade-offs between individual and systemic risk [51] and the effect of portfolio choice on fire-sale spillovers [20]. These approaches assume access to the joint portfolio information required by the optimizer. We remove this assumption by jointly computing portfolio adjustments and contingency liquidations that reduce modeled fire-sale losses while respecting fund-level constraints and fairness, without revealing the participating funds' portfolios to one another.

3 System and Threat Model

This section introduces the financial setting considered throughout the paper, the system-level quantities used by our analysis, the techniques and the privacy assumptions under which the participating funds interact.

3.1 Financial System

We consider a set of N financial institutions $\mathcal{F} = \{F_1, \ldots, F_N\}$, such as hedge funds, asset managers, banks, or other institutional investors, that invest in a common universe of M assets $\mathcal{A} = \{A_1, \ldots, A_M\}$. For concreteness, we refer to the participating institutions as *funds* throughout the paper. Fund F_i holds a private portfolio $\mathbf{h}_i = (h_{i1}, \ldots, h_{iM})^\top \in \mathbb{R}_+^M$, where $h_{ij} \geq 0$ denotes its monetary exposure to asset A_j, together with a cash position $c_i \geq 0$. The asset holdings of all funds are collected in the holdings matrix $\mathbf{H} = [h_{ij}] \in \mathbb{R}_+^{N \times M}$. Each fund F_i observes only its own holdings vector $\mathbf{h}_i$ and cash position c_i; the complete matrix $\mathbf{H}$ and the positions of the other funds are private. The total value of fund F_i is $V_i = c_i + \sum_{j=1}^M h_{ij}$. We denote the aggregate asset holdings of all participating funds by $\mathbf{h}^{\text{agg}} := \sum_{i=1}^N \mathbf{h}_i$, where $h_j^{\text{agg}} = \sum_{i=1}^N h_{ij}$ is the total monetary exposure of the system to asset A_j. We measure direct portfolio overlap by cosine similarity $s_{i\ell}^H := h_i^\top h_\ell / (\|h_i\|_2 \|h_\ell\|_2)$ and report its mean over $i < \ell$. We detail the operational scenario and system requirements in App. A.

Asset and fund correlations. Funds may be economically related through both common holdings and dependence among the underlying assets. Direct dependence arises when two funds hold the same assets, while indirect dependence may arise when they hold different assets whose returns co-move. Let $P_{j,t} > 0$ denote the price of asset A_j at time $t = 0, \ldots, T$, and define its continuously compounded return as $r_{j,t} = \log(P_{j,t}/P_{j,t-1})$, for $t = 1, \ldots, T$. Let $\bar{r}_j = T^{-1} \sum_{t=1}^T r_{j,t}$. The sample covariance between the returns of assets A_j and A_k is $\Sigma_{jk}^A = \frac{1}{T-1} \sum_{t=1}^T (r_{j,t} - \bar{r}_j)(r_{k,t} - \bar{r}_k)$, and the corresponding asset-return covariance matrix is $\boldsymbol{\Sigma}^A = [\Sigma_{jk}^A]_{j,k=1}^M$. The Pearson correlation between assets A_j and A_k is $\rho_{jk}^A = \Sigma_{jk}^A / \sqrt{\Sigma_{jj}^A \Sigma_{kk}^A}$, and the pairwise correlations form the asset-correlation matrix $\mathbf{R}^A = [\rho_{jk}^A]_{j,k=1}^M$.

To characterize dependence between funds, we combine their portfolio allocations with the covariance structure of the underlying assets. For funds F_i and F_ℓ,

their portfolio correlation is defined as

$$\rho^F_{i\ell} = \frac{\mathbf{h}_i^\top \boldsymbol{\Sigma}^A \mathbf{h}_\ell}{\sqrt{\mathbf{h}_i^\top \boldsymbol{\Sigma}^A \mathbf{h}_i}\sqrt{\mathbf{h}_\ell^\top \boldsymbol{\Sigma}^A \mathbf{h}_\ell}}\,. \tag{1}$$

The resulting fund-correlation matrix is $\mathbf{R}^F = [\rho^F_{i\ell}]_{i,\ell=1}^N$. This measure captures dependence arising both from overlapping holdings and from exposures to different assets whose returns co-move. The matrix $\mathbf{R}^F$ provides the basis for the system-level analysis developed in Sec. 4.

Stress and fire-sale losses. Under market stress, funds may be required to liquidate part of their asset holdings. Let $\mathbf{x}_i = (x_{i1}, \ldots, x_{iM})^\top \in \mathbb{R}^M_+$ denote the stressed liquidation vector of fund F_i, where x_{ij} is the monetary amount of asset A_j sold by that fund. Aggregate stressed sales are $\mathbf{q} := \sum_{i=1}^N \mathbf{x}_i$. Price impact is represented by $\boldsymbol{\Lambda} = [\Lambda_{jk}] \in \mathbb{R}^{M\times M}$, where Λ_{jk} measures the price decline of asset A_j induced by sales of asset A_k. The resulting stressed price-decline vector is $\boldsymbol{\delta} := \boldsymbol{\Lambda}\mathbf{q}$. The diagonal entries of $\boldsymbol{\Lambda}$ capture direct price impact, while the off-diagonal entries capture cross-asset spillovers. Following standard fire-sale models [30,23], we measure fire-sale losses as the decline in portfolio value induced by aggregate stressed sales.

Definition 1 (Fire-sale loss). *Given aggregate holdings* $\mathbf{h}^{\mathrm{agg}}$ *and aggregate stressed sales* $\mathbf{q}$*, the* system-wide fire-sale loss *is* $L(\mathbf{h}^{\mathrm{agg}}, \mathbf{q}) := (\mathbf{h}^{\mathrm{agg}})^\top \boldsymbol{\delta} = (\mathbf{h}^{\mathrm{agg}})^\top \boldsymbol{\Lambda}\mathbf{q}$*. The corresponding fire-sale loss of fund* F_i *is*

$$L_i(\mathbf{h}_i, \mathbf{q}) := \mathbf{h}_i^\top \boldsymbol{\delta} = \mathbf{h}_i^\top \boldsymbol{\Lambda}\mathbf{q}\,.$$

By construction, $L(\mathbf{h}^{\mathrm{agg}}, \mathbf{q}) = \sum_{i=1}^N L_i(\mathbf{h}_i, \mathbf{q})$, so the system-wide loss is exactly the sum of the individual fund losses. These quantities characterize the effect of a given stress scenario and form the basis of the portfolio-rebalancing problem introduced in Sec. 5.

3.2 Threat and Privacy Model

Multi-party computation (MPC) allows a set of mutually distrustful institutions to jointly evaluate a function over their private inputs without revealing the inputs or intermediate values beyond what can be inferred from the prescribed outputs. We consider a semi-honest adversary, where corrupted institutions follow the prescribed protocol but may inspect their local views and authorized outputs to infer information about honest participants. Private portfolio holdings, liquidation decisions, and intermediate values are therefore protected unless explicitly designated for release.

For the pure-MPC protocols, we target *dishonest-majority* with full-threshold security, thus any proper subset of the institutions may be corrupted. The preprocessing required is provided by a dealer that receives no portfolio inputs and is assumed not to collude with an online party. This dealer can be realized by a semi-honest independent party, or distributively [50,4].

The federated construction provides a weaker privacy guarantee because intermediate aggregation results are intentionally revealed across iterations. We therefore assume bounded collusion for the federated functionality. Appendix D formalizes the resulting output and transcript leakage.

Remark 1 (Regulatory accountability). In regulated deployments, the semi-honest cryptographic assumption may be complemented by institutional deterrence: participating institutions have known identities and may face contractual, financial or regulatory consequences for detected deviations [29,28]. Generic compilers from passive to covert security with public verifiability already exist [24,48]. Input authenticity could further be established by having holdings authenticated by an authoritative external source, such as a custodian, prime broker, exchange, or regulatory reporting system, and privately checking consistency between the authenticated commitment and the value supplied to the computation [7]. Both protocol and input verification could be distributed across multiple supervisory authorities rather than entrusted to a single regulator, consistent with established supervisory-college structures [8]. We leave the integration of these accountability and input-authentication mechanisms to future work.

4 Privacy-Preserving System-Level Analysis

This section provides a system-level view of the hedge funds by quantifying how their portfolios are concentrated, how strongly the funds are interconnected, and which institutions occupy structurally important positions within the market. The metrics help identify common exposures, clusters, central funds, and broader forms of portfolio crowding that are not visible from an individual portfolio alone. We first introduce the metrics used to characterize these properties and then describe how they can be computed in a privacy-preserving manner.

4.1 System-Level Financial Metrics

We use a compact set of metrics to characterize concentration, dependence, and connectivity across the participating funds. We begin with system-wide concentration, then examine pairwise fund dependence and the network structure induced by it, and extract the strongest dependence backbone through an MST.

System-wide concentration indices. We measure how total risky holdings are distributed across both assets and funds using the Herfindahl–Hirschman index (HHI) [33,32]. System-wide asset concentration is measured by $\mathrm{HHI}^A := \sum_{j=1}^{M} \left(h_j^{\mathrm{agg}} / \sum_{k=1}^{M} h_k^{\mathrm{agg}} \right)^2$, with $N_{\mathrm{eff}}^A := 1/\mathrm{HHI}^A$. System-wide fund concentration is measured by $\mathrm{HHI}^F := \sum_{i=1}^{N} \left(\sum_{j=1}^{M} h_{ij} / \sum_{k=1}^{M} h_k^{\mathrm{agg}} \right)^2$, with $N_{\mathrm{eff}}^F := 1/\mathrm{HHI}^F$. A high HHI^A indicates that aggregate exposure is concentrated in a small number of assets, whereas a high HHI^F indicates that a small number of funds account for a large share of total risky holdings. The effective numbers provide the corresponding equivalent numbers of equally weighted assets or funds.

Inter-fund dependence. The fund-correlation matrix $\mathbf{R}^F = [\rho^F_{i\ell}]_{i,\ell=1}^N$, defined in Eqn. (1), describes the pairwise dependence between the participating portfolios. For each fund F_i, we consider its average correlation $\bar{\rho}^F_i := (N-1)^{-1} \sum_{\ell \neq i} \rho^F_{i\ell}$ and average absolute correlation $\bar{\rho}^F_{i,\mathrm{abs}} := (N-1)^{-1} \sum_{\ell \neq i} |\rho^F_{i\ell}|$. The former indicates whether a fund tends to move with or against the rest of the system, while the latter captures the overall magnitude of its dependence regardless of sign. At the system level, we analogously define the mean pairwise correlation $\bar{\rho}^F := \binom{N}{2}^{-1} \sum_{i<\ell} \rho^F_{i\ell}$ and mean absolute pairwise correlation $\bar{\rho}^F_{\mathrm{abs}} := \binom{N}{2}^{-1} \sum_{i<\ell} |\rho^F_{i\ell}|$.

Correlation network and centrality. Pairwise fund correlations can be represented as a network in which funds are nodes and sufficiently strong positive correlations form edges, as commonly done in financial correlation networks [42]. For a correlation threshold $\tau \geq 0$, we define the undirected graph $\mathcal{G}_\tau = (\mathcal{F}, \mathcal{E}_\tau)$, where $\mathcal{E}_\tau := \{(F_i, F_\ell) : i < \ell,\ \rho^F_{i\ell} \geq \tau\}$ is the edge set connecting fund pairs whose correlation is at least τ. The threshold-network degree of fund F_i is $\deg_\tau(F_i) := \sum_{\ell \neq i} \mathbb{1}\{\rho^F_{i\ell} \geq \tau\}$, and its normalized degree centrality is $\deg_\tau(F_i)/(N-1)$. This measures the fraction of other funds to which F_i is strongly positively correlated. To account also for the strength of these connections, we define the node strength $\mathrm{str}_\tau(F_i) := \sum_{\ell \neq i} \mathbb{1}\{\rho^F_{i\ell} \geq \tau\} \rho^F_{i\ell}$, with normalized strength centrality $\mathrm{str}_\tau(F_i)/(N-1)$. A high value therefore indicates a fund with strong correlations to a substantial portion of the system. At the system level, the network density $\mathrm{dens}(\mathcal{G}_\tau) := |\mathcal{E}_\tau| / \binom{N}{2}$ measures the fraction of possible edges retained at threshold τ. Connected components and their sizes characterize the fragmentation of $\mathcal{G}_\tau$ at the chosen threshold.

Minimum spanning tree. While the threshold network retains all correlations above τ, it may remain dense and obscure the system's strongest dependence structure. Following standard correlation-network constructions [41,42], we map each fund correlation to the distance $\mathrm{dist}^F_{i\ell} := \sqrt{2(1-\rho^F_{i\ell})}$, so that more strongly correlated funds are closer. We define a minimum spanning tree (MST) as $\mathcal{T} = (\mathcal{F}, \mathcal{E}_{\mathrm{MST}})$, where $\mathcal{E}_{\mathrm{MST}}$ minimizes $\sum_{(F_i,F_\ell)\in\mathcal{E}'} \mathrm{dist}^F_{i\ell}$ over all edge sets $\mathcal{E}'$ for which $(\mathcal{F}, \mathcal{E}')$ is connected and acyclic. Since any connected acyclic graph on N vertices has exactly $N-1$ edges, the resulting tree spans all N funds using $N-1$ edges. The MST degree $\deg_{\mathrm{MST}}(F_i) := \sum_{\ell \neq i} \mathbb{1}\{(F_i, F_\ell) \in \mathcal{E}_{\mathrm{MST}}\}$ identifies funds that act as hubs in this backbone, while the total tree length $\mathrm{len}(\mathcal{T}) := \sum_{(F_i,F_\ell)\in\mathcal{E}_{\mathrm{MST}}} \mathrm{dist}^F_{i\ell}$ summarizes how tightly connected these strongest relationships are.

4.2 MPC-Based Metric Calculation

The metrics of Sec. 4.1 can be evaluated on secret-shared portfolios, but a direct implementation may require costly nonlinear operations such as secure divisions and square roots, often repeated for many fund or asset pairs. We therefore reformulate the computations to reuse common terms and, whenever only comparisons or graph structure are needed, avoid explicit normalization altogether.

System-wide concentration. Both concentration indices share the same denominator. Let $B := \sum_{j=1}^{M} h_j^{\mathrm{agg}} = \sum_{i,j} h_{ij}$. We compute $\mathrm{HHI}^A = B^{-2} \sum_j (h_j^{\mathrm{agg}})^2$ and $\mathrm{HHI}^F = B^{-2} \sum_i (\sum_j h_{ij})^2$, so only a single secure reciprocal of B is needed for both indices; the remaining operations are secure sums and multiplications.

Correlation statistics. Define $\kappa_{i\ell} := \mathbf{h}_i^\top \boldsymbol{\Sigma}^A \mathbf{h}_\ell$ and $\nu_i := \kappa_{ii}$, so that $\rho_{i\ell}^F = \kappa_{i\ell}/\sqrt{\nu_i \nu_\ell}$, consistently with Eqn. (1). We compute $\boldsymbol{\Sigma}^A \mathbf{h}_i$ once per fund and obtain the pairwise $\kappa_{i\ell}$ values through secure dot products. To avoid repeatedly evaluating reciprocal square roots, we compute $\alpha_i := 1/\sqrt{\nu_i}$ once per fund and then evaluate $\rho_{i\ell}^F = \kappa_{i\ell} \alpha_i \alpha_\ell$. This reduces the number of reciprocal square-root evaluations from $\Theta(N^2)$ to $\Theta(N)$. We then obtain the correlation averages, extrema, and strength statistics from the resulting shared correlation values through secure aggregation.

Division-free network construction. Several network metrics require only comparisons of correlations and therefore do not require the correlations themselves. In particular, since the denominator in Eqn. (1) is strictly positive, for the non-negative threshold τ used in $\mathcal{G}_\tau$,

$$\rho_{i\ell}^F \geq \tau \iff \kappa_{i\ell} \geq 0 \ \wedge \ \kappa_{i\ell}^2 \geq \tau^2 \nu_i \nu_\ell \,.$$

Threshold edges, degree centrality, and network density can thus be computed using multiplications and secure comparisons, without division or square roots.

The same observation is particularly useful for the MST. Since $\mathrm{dist}_{i\ell}^F = \sqrt{2(1-\rho_{i\ell}^F)}$ is strictly decreasing in $\rho_{i\ell}^F$, the minimum spanning tree defined above is exactly the maximum spanning tree under fund correlation. Correlations can in turn be compared directly from adapted Eqn. (1). Correlations of different signs are ordered by their signs; for two correlations with the same sign, their magnitudes can be compared without normalization by squaring and cross-multiplying,

$$\frac{\kappa_{i\ell}^2}{\nu_i \nu_\ell} \gtreqless \frac{\kappa_{rs}^2}{\nu_r \nu_s} \iff \kappa_{i\ell}^2 \nu_r \nu_s \gtreqless \kappa_{rs}^2 \nu_i \nu_\ell \,,$$

with the comparison reversed when both correlations are negative. Hence, the edge ordering requires only multiplications and comparisons, without constructing $d_{i\ell}^F$ or evaluating pairwise divisions or square roots.

We construct the tree incrementally: starting from one fund, each step adds the highest-correlation edge connecting the current tree to an unselected fund. This requires $\Theta(N^2)$ secure comparisons and avoids sorting all $\binom{N}{2}$ candidate edges. Only if numerical edge distances or $\mathrm{len}(\mathcal{T})$ are required do we evaluate the distance expression for the selected $N-1$ edges.

5 Privacy-Preserving Portfolio Rebalancing

Section 4 focuses on analyzing the financial system and characterizing the relationships and risk exposures that arise across participating funds. In this section,

we move from *analysis* to *mitigation*: given a stressed market scenario, our goal is to determine how the funds can adjust their portfolios so as to reduce the resulting system-wide fire-sale losses. The rebalancing should remain economically practical, requiring only limited and low-cost changes to the original portfolios, while preserving each fund's investment constraints. At the same time, the resulting reduction in systemic loss should not be achieved at the expense of a subset of participants; the benefits of rebalancing should be distributed in a fair and economically defensible manner. We first formulate the rebalancing problem, then consider the privacy-preserving implementation.

5.1 Portfolio-Rebalancing Problem

We consider a stressed market scenario in which participating funds jointly adjust their portfolios to reduce system-wide fire-sale losses while minimizing adjustment cost, preserving fund-level constraints, and ensuring individually acceptable and fairly distributed benefits. Let h_{ij}^{base} denote the pre-rebalancing holding of asset A_j by fund F_i, corresponding to h_{ij} from Sec. 3.1, and collect these holdings in $\mathbf{H}^{\text{base}} = [h_{ij}^{\text{base}}]$. The post-rebalancing holdings are denoted by $\widetilde{h}_{ij}$ and collected in $\widetilde{\mathbf{H}} = [\widetilde{h}_{ij}]$. Rebalancing preserves each fund's baseline risky-asset budget $B_i^{\text{base}} := \sum_j h_{ij}^{\text{base}}$, so that $\sum_j \widetilde{h}_{ij} = B_i^{\text{base}}$, while respecting the position limits $\ell_{ij} \le \widetilde{h}_{ij} \le u_{ij}$.

Stress scenario and baseline. We model a common redemption shock by a rate $\rho_{\text{red}} \in [0,1]$, representing the fraction of the baseline fund value $V_i^{\text{base}} := c_i + B_i^{\text{base}}$ that must be returned to investors under stress. Cash is used first; hence, fund F_i must raise $D_i = [\rho_{\text{red}} V_i^{\text{base}} - c_i]^+$ through risky-asset sales, where $[z]^+ := \max\{z, 0\}$. Its stressed liquidation vector $\mathbf{x}_i$ must therefore satisfy $\sum_j x_{ij} = D_i$ and $0 \le x_{ij} \le \widetilde{h}_{ij}$. We collect the liquidation vectors in $\mathbf{X} = [x_{ij}]$ and define the associated aggregate sales by $\mathbf{q}(\mathbf{X}) := \sum_i \mathbf{x}_i$. For the post-rebalancing liquidation decision, we write simply $\mathbf{q} := \mathbf{q}(\mathbf{X})$. These sales induce $\boldsymbol{\delta} = \boldsymbol{\Lambda}\mathbf{q}$, with the component-wise price-impact cap $\boldsymbol{\Lambda}\mathbf{q} \le \bar{\delta}\mathbf{1}$, where $\bar{\delta} \ge 0$ is the maximum allowed stressed price decline for any asset. The repair is evaluated relative to the baseline system. Let $\mathbf{h}^{\text{agg,base}} := \sum_i \mathbf{h}_i^{\text{base}}$, and let $\mathbf{X}^{\text{base}}$ denote a minimum-loss feasible liquidation under the baseline holdings: $\mathbf{X}^{\text{base}} \in \arg\min_{\mathbf{X} \in \mathcal{X}(\mathbf{H}^{\text{base}})} L(\mathbf{h}^{\text{agg,base}}, \mathbf{q}(\mathbf{X}))$, where $\mathcal{X}(\mathbf{H}^{\text{base}})$ denotes the set of liquidation matrices satisfying the same liquidation and price-impact constraints under $\mathbf{H}^{\text{base}}$. We define $\mathbf{q}^{\text{base}} := \mathbf{q}(\mathbf{X}^{\text{base}})$, and the corresponding baseline loss as $L_0 := L(\mathbf{h}^{\text{agg,base}}, \mathbf{q}^{\text{base}})$. For a required reduction $\gamma \in [0,1]$, let $\widetilde{\mathbf{h}}^{\text{agg}} := \sum_i \widetilde{\mathbf{h}}_i$ and $L := L(\widetilde{\mathbf{h}}^{\text{agg}}, \mathbf{q})$. The repaired system must satisfy $L \le (1-\gamma)L_0$.

Individual participation and fairness. The baseline loss of fund F_i is $L_i^{\text{base}} = L_i(\mathbf{h}_i^{\text{base}}, \mathbf{q}^{\text{base}})$, while its post-rebalancing loss is $L_i = L_i(\widetilde{\mathbf{h}}_i, \mathbf{q})$. We define its achieved loss reduction as $G_i := L_i^{\text{base}} - L_i$ and require $G_i \ge \theta L_i^{\text{base}}$, where $\theta \ge 0$ is the minimum individual-improvement rate. Thus, $\theta = 0$ imposes a no-harm condition, while $\theta > 0$ guarantees a positive improvement whenever $L_i^{\text{base}} > 0$. The total improvement satisfies $G := \sum_i G_i = L_0 - L$.

Individual protection alone does not specify how this total improvement should be distributed. Building on the Shapley value [49], we introduce a measure of each fund's structural responsibility for portfolio crowding.

Definition 2 (Shapley-aligned fairness). *For a coalition $\mathcal{S} \subseteq \mathcal{F}$, let $\mathbf{h}_{\mathcal{S}}^{\text{base}} := \sum_{F_i \in \mathcal{S}} \mathbf{h}_i^{\text{base}}$ and define its crowding value as $v(\mathcal{S}) := \frac{1}{2}(\mathbf{h}_{\mathcal{S}}^{\text{base}})^\top \boldsymbol{\Lambda} \mathbf{h}_{\mathcal{S}}^{\text{base}}$. Let ϕ_i be the Shapley value of fund F_i and $s_i := \phi_i / \sum_k \phi_k$ its responsibility share. Define $d_i := |G_i - s_i G|$, $D_G := \sum_i d_i$, and $A_{\text{Shapley}} := 1 - D_G/(2G)$. A repair is perfectly Shapley-aligned if $A_{\text{Shapley}} = 1$, equivalently $G_i/G = s_i$ for all i*[1].

Since direct Shapley computation requires summing marginal contributions over all coalitions, the following proposition provides an equivalent closed-form expression that avoids explicit coalition enumeration. Its proof is given in App. B.1.

Proposition 1 (Closed-form Shapley responsibility). *If $\boldsymbol{\Lambda}$ is symmetric, then $\phi_i = \frac{1}{2}(\mathbf{h}_i^{\text{base}})^\top \boldsymbol{\Lambda} \mathbf{h}^{\text{agg,base}}$.*

Adjustment cost. Let $\Delta h_{ij} := \widetilde{h}_{ij} - h_{ij}^{\text{base}}$ denote the signed adjustment of asset A_j by fund F_i, and let $c_{ij}^{\text{tr}} \geq 0$ be its per-unit adjustment cost. The total adjustment cost is $C(\widetilde{\mathbf{H}}) := \sum_{i,j} c_{ij}^{\text{tr}} |\Delta h_{ij}|$, while fund-level turnover is constrained by $\sum_j |\Delta h_{ij}| \leq \omega_{\max} B_i^{\text{base}}$, where $\omega_{\max} \in [0,1]$ is the maximum allowed turnover fraction.

Repair objective. Combining adjustment cost and fairness, we solve

$$\min_{\widetilde{\mathbf{H}}, \mathbf{X}} (1-\beta) \frac{C(\widetilde{\mathbf{H}})}{C_{\text{ref}}} + \beta \frac{D_G}{2G_{\text{ref}}}, \tag{2}$$

subject to the constraints summarized in Fig. 1. These constraints ensure portfolio and liquidation feasibility, limit market impact and turnover, and enforce the required systemic and fund-level loss reductions. In the repair objective, $\beta \in [0,1]$ controls the cost–fairness trade-off, $C_{\text{ref}} := \sum_{i,j} c_{ij}^{\text{tr}} h_{ij}^{\text{base}}$ is the baseline cost scale, and $G_{\text{ref}} := \max\{\gamma, \theta\} L_0$ is the minimum total loss reduction implied by the systemic and individual-improvement requirements, assuming $C_{\text{ref}}, G_{\text{ref}} > 0$. Appendix B.2 examines how β affects Shapley alignment across different portfolio-concentration settings. Appendix B.3 characterizes the fairness-saturation threshold and its implications for calibrating β.

5.2 MPC-Based Portfolio Rebalancing

The resulting problem is jointly bilinear because both $L(\widetilde{\mathbf{h}}^{\text{agg}}, \mathbf{q})$ and $L_i(\widetilde{\mathbf{h}}_i, \mathbf{q})$ couple post-rebalancing holdings with aggregate stressed sales. Nevertheless, fixing $\widetilde{\mathbf{H}}$ yields a linear liquidation problem in $\mathbf{X}$, while fixing $\mathbf{X}$ yields a linear rebalancing problem in $\widetilde{\mathbf{H}}$. We therefore retain the alternating structure

$$(\widetilde{\mathbf{H}}^{(k)}, \mathbf{X}^{(k)}) \longrightarrow \mathbf{X}^{(k+1)} \longrightarrow \widetilde{\mathbf{H}}^{(k+1)}, \tag{3}$$

[1] For this definition, we assume $\phi_i \geq 0$ for all i, $\sum_k \phi_k > 0$, and $G > 0$.

Constraints of the multi-fund portfolio-rebalancing problem

(C1) Preserve each fund's risky-asset budget:	$\sum_j \widetilde{h}_{ij} = B_i^{\text{base}}$
(C2) Respect asset-position limits:	$\ell_{ij} \le \widetilde{h}_{ij} \le u_{ij}$
(C3) Meet each fund's stressed redemption need:	$\sum_j x_{ij} = D_i$
(C4) Limit sales by available holdings:	$0 \le x_{ij} \le \widetilde{h}_{ij}$
(C5) Limit stressed price declines:	$\boldsymbol{\Lambda}\mathbf{q} \le \bar{\delta}\mathbf{1}$
(C6) Reduce system-wide loss by at least γ:	$L(\widetilde{\mathbf{h}}^{\text{agg}}, \mathbf{q}) \le (1-\gamma)L_0$
(C7) Reduce each fund's loss by at least θ:	$L_i(\widetilde{\mathbf{h}}_i, \mathbf{q}) \le (1-\theta)L_i^{\text{base}}$
(C8) Limit fund-level turnover:	$\sum_j \lvert\Delta h_{ij}\rvert \le \omega_{\max} B_i^{\text{base}}.$

Fig. 1. Constraints of the portfolio-rebalancing problem. Fund-specific constraints hold for $i = 1, \ldots, N$ and, where applicable, $j = 1, \ldots, M$.

and focus on solving each linear block efficiently under MPC.

First-order optimization. In plaintext, the resulting linear programs (LPs) can be solved efficiently using mature simplex-based implementations [34]. Under MPC, however, simplex is poorly matched to secure computation: pivot selection and ratio tests require data-dependent comparisons and secret indexing, pivoting involves divisions, and the number and sequence of iterations depend on private intermediate values. We instead solve each block using the primal–dual hybrid gradient (PDHG) method [11]. PDHG replaces pivoting by repeated primal and dual updates consisting primarily of matrix–vector products, additions, multiplications, and simple projections. These operations map naturally to MPC and require neither secret pivot selection nor secret division. Public step sizes and a fixed iteration budget further give a data-independent execution. Appendix C.1 gives the complete update equations and their specialization to the two blocks.

Feasibility and optimization. Starting from $(\widetilde{\mathbf{H}}^{(0)}, \mathbf{X}^{(0)})$, Phase I introduces a non-negative slack η on the systemic-loss and individual-improvement requirements and minimizes the resulting feasibility violation for a fixed public number of alternating outer rounds. Phase II then removes this slack, restores the hard constraints, and minimizes Eqn. (2) for a second fixed public number of outer rounds. The phase lengths are fixed before execution and do not depend on whether or when η reaches zero. Thus, the economic formulation and alternating structure are unchanged, while both the outer schedule and each inner PDHG solve use fixed public iteration budgets.

5.3 Federated Portfolio Rebalancing

The MPC-based construction of Sec. 5.2 keeps every first-order iteration inside secure computation, providing the strongest privacy boundary but at substantial computation and communication cost. As a complementary performance-oriented

deployment, we move fund-local optimization outside MPC and use MPC only for secure aggregation. No trusted coordinator is required, as the aggregation functionality is jointly realized by the participating funds.

For the liquidation block, fund F_i updates its own liquidation $\mathbf{x}_i$ locally. At iteration r, the funds secret-share their local liquidation contributions and MPC reconstructs only their aggregate, $\mathbf{q}^{(r)} = \sum_{i=1}^{N} \mathbf{x}_i^{(r)}$. Analogously, in the rebalancing block MPC aggregates the local repaired holdings as $\widetilde{\mathbf{h}}^{\mathrm{agg},(r)} = \sum_i \widetilde{\mathbf{h}}_i^{(r)}$. The same mechanism is used for the additional contributions needed by the globally coupled loss, fairness, and feasibility constraints: each fund computes its local contribution, MPC sums the secret shares, and only the resulting aggregate is revealed.

The revealed aggregates provide the global information required for each fund to perform its next first-order update locally and in parallel. Individual fund contributions are never reconstructed, but the aggregate coordination values are intentionally visible to all participants. The outer feasibility-restoration and alternating liquidation/rebalancing procedure is otherwise unchanged.

This substantially reduces the amount of computation performed under MPC, at the cost of releasing iterative aggregate information. Section 6 and App. D analyze the corresponding weaker privacy boundary, while App. D.2 gives the complete aggregation details.

6 Security and Privacy

We analyze security in the standard real/ideal simulation paradigm [19], distinguishing confidentiality provided by the MPC protocol from information intentionally revealed by the functionality. For the pure-MPC constructions, security is inherited from the underlying semi-honest full-threshold MPC protocol by standard composition. In particular, for every coalition $\mathcal{C} \subsetneq \mathcal{F}$, there exists a PPT simulator $\mathcal{S}$ such that

$$\mathrm{View}_{\mathcal{C}}^{\Pi}(\mathbf{u}; \mathbf{p}) \approx_c \mathcal{S}(\{\mathbf{u}_i\}_{i\in\mathcal{C}}, \mathbf{p}, \{\mathcal{O}_i\}_{i\in\mathcal{C}}, \mathcal{O}_{\mathrm{pub}}) \ ,$$

where $\mathbf{u}_i$ is fund F_i's private input, $\mathbf{p}$ denotes public parameters, $\mathcal{O}_i$ its prescribed fund-specific output, and $\mathcal{O}_{\mathrm{pub}}$ any deliberately released system-level output. This follows because our protocols compose standard secure arithmetic, comparison, and oblivious-selection primitives.

Consequently, unreleased holdings, aggregates, correlations, optimization states, and liquidation decisions remain secret-shared. Fixed public iteration counts additionally prevent convergence-dependent control-flow leakage. For portfolio repair, fund F_i learns only its repaired portfolio $\widetilde{\mathbf{h}}_i$ and contingency liquidation $\mathbf{x}_i$, together with any outputs explicitly designated as public. Any inference from these authorized outputs is inherent in the requested functionality rather than additional leakage from the MPC execution.

The federated construction has a weaker privacy boundary. Individual contributions are protected during secure aggregation, but the resulting aggregate

coordination quantities are intentionally reconstructed to the funds across iterations. These values therefore form an additional application-level transcript that may permit inference about other funds, particularly under collusion. Appendix D addresses this additional leakage. A detailed analysis of iterative transcript inference is outside our scope and has been studied in the distributed-optimization literature [40,39].

7 Evaluation

We implement the secure computations in MP-SPDZ [35] and evaluate them on a single machine equipped with a 12th Gen Intel Core i7-12850HX processor and 32 GB RAM, with parties executing as separate processes. We use a semi-honest full-threshold protocol based on additive secret sharing over rings. Secure arithmetic is performed over a 128-bit ring, with real-valued quantities represented in fixed-point arithmetic [22] using 20 fractional bits. The large ring provides sufficient headroom for intermediate computations, while avoiding overflow without sacrificing the numerical precision required by the optimization and correlation computations. Consistent with Sec. 3, we report wall-clock time and assume that correlated randomness is generated beforehand through distributed preprocessing or by an independent semi-honest dealer. We report preprocessing requirements in App. F.3. Communication is reported per fund party and, for the federated construction, includes both MPC traffic and the fund–MPC client channel. We make our implementation available at [10].

We construct the benchmark from public SEC Form N-PORT [55] filings for Q2 2026, using reported holdings, cash, and fund values rather than synthetic portfolios. Our representative instance contains $N = 10$ funds and $M = 200$ stocks. The selected universe covers at least 83.8% of each fund's portfolio, 95.5% of the assets are held by at least two funds, and mean pairwise holdings similarity is 0.476. Thus, the instance exhibits substantial portfolio overlap while retaining meaningful heterogeneity across funds.

For network experiments, we additionally emulate increasingly constrained WAN conditions. The local profile uses 0.1 ms round-trip time (RTT) latency and 30 Gbit/s throughput; WAN-10 uses 10 ms RTT and 10 Gbit/s, WAN-20 uses 20 ms RTT and 1 Gbit/s, and WAN-40 uses 40 ms RTT and 500 Mbit/s. The RTT range is consistent with latencies observed between European cloud regions [2], while the bandwidth limits correspond to capacities offered by commercial private-connectivity services [1,3].

7.1 Private portfolio analysis

System-level analysis is naturally suited to recurring monitoring as portfolios and market conditions evolve. MPC allows such diagnostics to be refreshed without disclosing the underlying portfolios. Our representative analysis completes in approximately 230 seconds even under the constrained WAN setting, making

Table 1. Runtime of private portfolio analysis for the representative $N = 10$, $M = 200$, $T = 60$ instance under different network conditions.

$N = 10,\ M = 200,\ T = 60$	Local [s]	WAN-10 [s]	WAN-20 [s]	WAN-40 [s]
Private portfolio analysis	7.50	64.61	121.36	228.95

Table 2. Solution quality on the representative $N = 10$, $M = 200$ N-PORT instance. Min. individual reduction denotes the smallest relative reduction in allocated fire-sale loss across all funds, $\min_i (L_i^{\text{base}} - L_i)/L_i^{\text{base}}$. Shapley alignment is $100 A_{\text{Shapley}}\%$. Repair cost is reported in USD millions, with the corresponding percentage of the aggregate initial portfolio value additionally displayed. Plaintext HiGHS solves each alternating LP block exactly, whereas the secure variants use fixed-iteration PDHG.

Method	System loss reduction [%]	Min. individual reduction [%]	Mean turnover [%]	Shapley alignment [%]	Repair cost [$m]
Plaintext	10.000	6.044	0.927	100.00	0.919 (0.000186%)
MPC	10.027	6.056	1.097	99.80	0.834 (0.000169%)
Federated	10.007	6.145	1.977	98.86	1.349 (0.000273%)

repeated recomputation practical and allowing it to serve as a screening step before invoking the substantially more expensive portfolio-repair protocol.

For $N = 10$, $M = 200$, and $T = 60$ public asset-return observations, the secure protocol computes portfolio-return dependence, the correlation network and maximum spanning tree, clustering, and holdings-based Shapley responsibility. Plaintext verification reproduces all threshold-edge decisions and all nine MST edges. The resulting network is nontrivial: at threshold $\tau = 0.9$, seven funds form a connected core, while the most central fund by return correlation differs from the fund with largest structural-responsibility share. Appendix E reports the actual funds (Section E.1), complete per-fund statistics (App. E.2), MST (Fig. 4), threshold sensitivity (App. E.3) and component costs (App. E.4).

7.2 Private portfolio rebalancing

For portfolio rebalancing, unless stated otherwise, we use

$$\rho_{\text{red}} = 5\%, \quad \gamma = 10\%, \quad \omega_{\max} = 15\%, \quad \bar{\delta} = 30\%, \quad \beta = 0.1, \quad \theta = 0.001\,.$$

Thus, funds face a 5% common redemption shock and the repair must reduce systemic fire-sale loss by at least 10%. Each fund is limited to 15% gross turnover. With $\beta = 0.1$, the normalized objective assigns a larger coefficient to adjustment cost than to Shapley-deviation, while retaining the individual-improvement requirement. The secure solvers use fixed public iteration budgets to avoid data-dependent execution. Further solver and numerical details are given in App. F.

Solution quality. Table 2 shows that the computational savings do not materially change the economic outcome. All three methods attain the required 10% systemic-loss reduction, while even the least-benefiting fund improves by about 6%. Mean

Table 3. Scalability of portfolio repair with the number of funds for $M = 200$ assets. Communication denotes data transmitted by one fund party; for the federated construction, this includes both MPC communication and the fund–MPC client channel.

$M = 200$	Communication [GB]				Runtime [min]			
	$N = 5$	$N = 10$	$N = 15$	$N = 20$	$N = 5$	$N = 10$	$N = 15$	$N = 20$
MPC	16.770	72.331	167.307	302.621	22.86	27.65	46.77	94.19
Federated	0.108	0.198	0.289	0.380	7.59	7.72	7.92	8.21

turnover remains below 2%, far from the 15% cap, and Shapley alignment remains 99.80% for MPC PDHG and 98.86% for Federated PDHG. Repair costs are also of the same order as the plaintext reference using HiGHS LP solver [34]. Detailed fund- and asset-level comparisons are given in App. F.7.

Scalability. The gap between our two secure constructions widens as more funds participate. MPC PDHG communication grows approximately quadratically in N (see App. F.4), from 16.77 to 302.62 GB per party, whereas Federated PDHG communication grows approximately linearly, from only 0.108 to 0.380 GB. MPC PDHG runtime grows superlinearly (see App. F.4), while the federated runtime remains nearly constant over this range. Consequently, the runtime advantage increases from about 3× at $N = 5$ to 11.5× at $N = 20$, and the communication reduction approaches 800×. This reflects the trade-off of the federated construction: individual contributions remain protected during secure aggregation, while aggregate coordination values are released.

Appendix F.5 additionally varies the asset dimension from $M = 50$ to $M = 500$, and App. F.6 evaluates the representative instance under WAN conditions. Both experiments show the same qualitative trend: the federated construction becomes increasingly advantageous as secure computation becomes more expensive.

8 Conclusion

We presented, to the best of our knowledge, the first privacy-preserving framework that combines system-level portfolio analysis with coordinated portfolio repair. Using MPC, institutions can jointly identify cross-portfolio dependencies and compute concrete rebalancing and contingency-liquidation actions without disclosing their underlying holdings. Our evaluation on real N-PORT portfolios shows that private system analysis is practical under realistic network conditions and that coordinated repair can be carried out within the time scales of scheduled risk management. A complementary federated construction substantially reduces runtime and communication at the cost of additional leakage from iterative aggregate coordination. These results show that secure computation can support not only private financial analysis, but also coordinated intervention to reduce systemic risk.

Acknowledgements

This work was supported by the Flemish Government through the Cybersecurity Research Program with grant number: VOEWICS02.

References

1. Azure ExpressRoute Overview: Connect over a private connection — learn.microsoft.com. https://learn.microsoft.com/en-us/azure/expressroute/expressroute-introduction?utm, [Accessed 10-09-2026]
2. Azure network round-trip latency statistics — learn.microsoft.com. https://learn.microsoft.com/da-dk/azure/networking/azure-network-latency?tabs={E}urope%2{C}{W}estern{E}urope, [Accessed 10-09-2026]
3. Direct Connect connection options - AWS Direct Connect — docs.aws.amazon.com. https://docs.aws.amazon.com/directconnect/latest/{U}ser{G}uide/connection_options.html?utm, [Accessed 10-09-2026]
4. Aaraj, N., Aly, A., Güneysu, T., Marcolla, C., Mono, J., Paludo, R., Santos-González, I., Scholz, M., Soria-Vazquez, E., Sucasas, V., et al.: FANNG-MPC: Framework for artificial neural networks and generic MPC. Cryptology ePrint Archive (2023)
5. Abbe, E.A., Khandani, A.E., Lo, A.W.: Privacy-Preserving Methods for Sharing Financial Risk Exposures. American Economic Review **102**(3), 65–70 (2012). https://doi.org/10.1257/aer.102.3.65
6. Agarwal, V., Jiang, W., Tang, Y., Yang, B.: Uncovering hedge fund skill from the portfolio holdings they hide. The Journal of Finance **68**(2), 739–783 (2013)
7. Agrawal, N., Bell, J., Gascón, A., Kusner, M.J.: Mpc-friendly commitments for publicly verifiable covert security. In: Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security. pp. 2685–2704 (2021)
8. Alford, D.: Supervisory colleges: The global financial crisis and improving international supervisory coordination. Emory Int'l L. Rev. **24**, 57 (2010)
9. Aly, A., Cuvelier, E., Mawet, S., Pereira, O., Van Vyve, M.: Securely solving simple combinatorial graph problems. In: International Conference on Financial Cryptography and Data Security. pp. 239–257. Springer (2013)
10. Anonymous: Anonymous GitHub Repository. https://anonymous.4open.science/r/Private-Portfolio-Analysis-and-Repair (2026), accessed: September 2026
11. Applegate, D., Díaz, M., Hinder, O., Lu, H., Lubin, M., O'Donoghue, B., Schudy, W.: Practical large-scale linear programming using primal-dual hybrid gradient. Advances in Neural Information Processing Systems **34**, 20243–20257 (2021)
12. Aragon, G.O., Hertzel, M., Shi, Z.: Why do hedge funds avoid disclosure? Evidence from confidential 13F filings. Journal of Financial and Quantitative Analysis **48**(5), 1499–1518 (2013)
13. Asharov, G., Balch, T., Polychroniadou, A.: Privacy-preserving portfolio pricing. In: Proceedings of the Second ACM International Conference on AI in Finance. pp. 1–8 (2021)
14. Barucca, P., Mahmood, T., Silvestri, L.: Common asset holdings and systemic vulnerability across multiple types of financial institution. Journal of Financial Stability **52**, 100810 (2021)
15. Beaver, D.: Efficient multiparty protocols using circuit randomization. In: Feigenbaum, J. (ed.) CRYPTO (1991)

16. Bednarz, A., Bean, N., Roughan, M.: Hiccups on the road to privacy-preserving linear programming. In: Proceedings of the 8th ACM Workshop on Privacy in the Electronic Society. pp. 117–120 (2009)
17. Bogdanov, D., Talviste, R., Willemson, J.: Deploying Secure Multi-Party Computation for Financial Data Analysis: (Short Paper). In: International Conference on Financial Cryptography and Data Security. pp. 57–64. Springer (2012)
18. Byun, J., Ko, H., Lee, J.: A privacy-preserving mean–variance optimal portfolio. Finance Research Letters **54**, 103794 (2023)
19. Canetti, R.: Security and composition of multiparty cryptographic protocols. Journal of Cryptology **13**(1), 143–202 (2000)
20. Capponi, A., Weber, M.: Systemic portfolio diversification. Operations Research **72**(1), 110–131 (2024)
21. Catrina, O., De Hoogh, S.: Secure multiparty linear programming using fixed-point arithmetic. In: European Symposium on Research in Computer Security. pp. 134–150. Springer (2010)
22. Catrina, O., Saxena, A.: Secure computation with fixed-point numbers. In: International Conference on Financial Cryptography and Data Security. pp. 35–50. Springer (2010)
23. Cont, R., Schaanning, E.: Fire sales, indirect contagion and systemic stress testing. Tech. Rep. 2/2017, Norges Bank (2017)
24. Damgård, I., Orlandi, C., Simkin, M.: Black-box transformations from passive to covert security with public verifiability. In: Annual International Cryptology Conference. pp. 647–676. Springer (2020)
25. Dreier, J., Kerschbaum, F.: Practical privacy-preserving multiparty linear programming based on problem transformation. In: 2011 IEEE Third International Conference on Privacy, Security, Risk and Trust and 2011 IEEE Third International Conference on Social Computing. pp. 916–924. IEEE (2011)
26. Escudero, D., Ghosh, S., Keller, M., Rachuri, R., Scholl, P.: Improved Primitives for MPC over Mixed Arithmetic-Binary Circuits. In: CRYPTO (2020)
27. European Securities and Markets Authority: Guidelines on Liquidity Stress Testing in UCITS and AIFs. Tech. Rep. ESMA34-39-897, European Securities and Markets Authority (2020), `https://www.esma.europa.eu/sites/default/files/library/esma34-39-897_guidelines_on_liquidity_stress_testing_in_ucits_and_aifs_en.pdf`
28. Faust, S., Hazay, C., Kretzler, D., Schlosser, B.: Financially backed covert security. In: IACR International Conference on Public-Key Cryptography. pp. 99–129. Springer (2022)
29. George, M., Kamara, S.: Adversarial level agreements for two-party protocols. In: Proceedings of the 2022 ACM on Asia Conference on Computer and Communications Security. pp. 816–830 (2022)
30. Greenwood, R., Landier, A., Thesmar, D.: Vulnerable banks. Journal of Financial Economics **115**(3), 471–485 (2015). https://doi.org/10.1016/j.jfineco.2014.11.006
31. Hastings, M., Falk, B.H., Tsoukalas, G.: Privacy-Preserving Network Analytics. Management Science **69**(9), 5482–5500 (2023). https://doi.org/10.1287/mnsc.2022.4582
32. Herfindahl, O.C.: Concentration in the Steel Industry. Ph.D. thesis, Columbia University (1950)
33. Hirschman, A.O.: National Power and the Structure of Foreign Trade. University of California Press, Berkeley (1945)
34. Huangfu, Q., Hall, J.J.: Parallelizing the dual revised simplex method. Mathematical programming computation **10**(1), 119–142 (2018)

35. Keller, M.: MP-SPDZ: A Versatile Framework for Multi-Party Computation. In: Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security (2020). https://doi.org/10.1145/3372297.3417872
36. Ko, H., Byun, J., Lee, J.: A Privacy-Preserving Robo-Advisory System with the Black–Litterman Portfolio Model: A New Framework and Insights into Investor Behavior. Journal of International Financial Markets, Institutions and Money **89**, 101873 (2023). https://doi.org/10.1016/j.intfin.2023.101873
37. Ko, H., Lee, J., Byun, J.: Advancing Financial Privacy: A Novel Integrative Approach for Privacy-Preserving Optimal Portfolio. Future Generation Computer Systems **174**, 107901 (2026). https://doi.org/10.1016/j.future.2025.107901
38. Laud, P.: Parallel oblivious array access for secure multiparty computation and privacy-preserving minimum spanning trees. Proceedings on Privacy Enhancing Technologies (2015)
39. Li, Q., Heusdens, R., Christensen, M.G.: Privacy-preserving distributed optimization via subspace perturbation: A general framework. IEEE Transactions on Signal Processing **68**, 5983–5996 (2020)
40. Li, Q., Lopuhaä-Zwakenberg, M., Yu, W., Heusdens, R.: On the privacy bound of distributed optimization and its application in federated learning. In: 2024 32nd European Signal Processing Conference (EUSIPCO). pp. 2232–2236. IEEE (2024)
41. Mantegna, R.N.: Hierarchical structure in financial markets. The European Physical Journal B **11**(1), 193–197 (1999). https://doi.org/10.1007/s100510050929
42. Marti, G., Nielsen, F., Bińkowski, M., Donnat, P.: A review of two decades of correlations, hierarchies, networks and clustering in financial markets. Progress in information geometry: Theory and applications pp. 245–274 (2021)
43. Nguyen-Nam, T., Tran-Truong, P.T., Nguyen-Ngoc, P., Son, H.X., Phan, T.H., Nguyen, T.T., Le, K.H., Nguyen, C.G.: Privacy-preserving computation in financial systems: A systematic survey and quantitative analysis. Information Systems p. 102777 (2026)
44. Parida, S., Teo, T.: The impact of more frequent portfolio disclosure on mutual fund performance. Journal of Banking & Finance **87**, 427–445 (2018)
45. Pichler, A., Poledna, S., Thurner, S.: Systemic risk-efficient asset allocations: Minimization of systemic risk as a network optimization problem. Journal of Financial Stability **52**, 100809 (2021)
46. Poledna, S., Martínez-Jaramillo, S., Caccioli, F., Thurner, S.: Quantification of systemic risk from overlapping portfolios in the financial system. Journal of Financial Stability **52**, 100808 (2021)
47. Rotaru, D., Wood, T.: MArBled Circuits: Mixing Arithmetic and Boolean Circuits with Active Security. In: INDOCRYPT (2019)
48. Scholl, P., Simkin, M., Siniscalchi, L.: Multiparty computation with covert security and public verifiability. Cryptology ePrint Archive (2021)
49. Shapley, L.S.: A value for n-person games. In: Kuhn, H.W., Tucker, A.W. (eds.) Contributions to the Theory of Games, Volume II, pp. 307–317. Princeton University Press, Princeton (1953)
50. Smart, N.P., Tanguy, T.: Taas: Commodity MPC via Triples-as-a-service. In: Proceedings of the 2019 ACM SIGSAC conference on cloud computing security workshop. pp. 105–116 (2019)
51. Sulas, A., Maringer, D., Paterlini, S.: Systemic risk from overlapping portfolios: A multi-objective optimization framework. International Review of Financial Analysis **97**, 103794 (2025)
52. Toft, T.: Solving linear programs using multiparty computation. In: International conference on financial cryptography and data security. pp. 90–107. Springer (2009)

53. U.S. Securities and Exchange Commission: Net Asset Value. Investor.gov, https://www.investor.gov/introduction-investing/investing-basics/glossary/net-asset-value, accessed: 2026-09-11
54. U.S. Securities and Exchange Commission: Use of Derivatives by Registered Investment Companies and Business Development Companies. Investment Company Act Release No. 34084 (2020), https://www.sec.gov/rules/final/2020/ic-34084.pdf, accessed: 2026-09-11
55. U.S. Securities and Exchange Commission: Form N-PORT: Monthly Portfolio Holdings Report. 17 C.F.R. § 274.150 (2026), https://www.sec.gov/data-research/sec-markets-data/form-n-port-data-sets, accessed: 2026-09-10

A Operational Deployment Scenario

We consider a consortium of financial institutions that periodically assesses and, when necessary, reduces risks arising from their joint portfolio exposures. Each institution operates an MPC node within infrastructure under its control, and the parties communicate over authenticated channels. This periodic mode of operation is consistent with existing risk-management practice. For example, U.S. registered funds subject to the value-at-risk (VaR) requirements of Rule 18f-4 assess compliance with the applicable VaR limit at least once each business day and conduct stress testing at least weekly [54]. Mutual funds generally calculate their net asset value (NAV) each business day, typically after major U.S. exchanges close [53]. Similarly, the European Securities and Markets Authority (ESMA) recommends quarterly liquidity stress testing, with the frequency adapted to the fund's characteristics and liquidity profile [27].

These existing risk-management cadences provide natural operating windows for our protocols. Private portfolio analysis can, for example, be performed after end-of-day positions have been established, while more extensive systemic-risk assessments can be conducted on a weekly or periodic basis. Since the participants determine when an analysis is initiated, multi-hour secure computations can be scheduled between portfolio snapshots and the corresponding risk-management decision. If an agreed analysis or stress scenario identifies a material vulnerability, the portfolio-repair protocol can then compute coordinated adjustments intended to reduce the identified systemic risk before the next portfolio decision. The correlated randomness required by the MPC backend is independent of portfolio inputs and can therefore be generated in advance during preprocessing.

B Portfolio Rebalancing: Supplementary Analysis

B.1 Closed-Form Shapley Responsibility

We derive the closed-form Shapley responsibility used in Definition 2. The result applies to a general, possibly asymmetric, price-impact matrix and reduces to the expression used in the main text when $\boldsymbol{\Lambda}$ is symmetric.

Proposition 2 (Closed-form Shapley responsibility). *For the crowding game* $v(\mathcal{S}) = \frac{1}{2}(\mathbf{h}_{\mathcal{S}}^{\text{base}})^\top \boldsymbol{\Lambda} \mathbf{h}_{\mathcal{S}}^{\text{base}}$, *the Shapley responsibility of fund* F_i *is*

$$\phi_i = \frac{1}{4}(\mathbf{h}_i^{\text{base}})^\top \left(\boldsymbol{\Lambda} + \boldsymbol{\Lambda}^\top\right) \mathbf{h}^{\text{agg,base}}.$$

In particular, if $\boldsymbol{\Lambda}$ *is symmetric, then*

$$\phi_i = \frac{1}{2}(\mathbf{h}_i^{\text{base}})^\top \boldsymbol{\Lambda} \mathbf{h}^{\text{agg,base}}.$$

Proof. Let $\Pi(\mathcal{F})$ denote the set of all permutations of the N funds, and let $P_i(\pi)$ denote the set of funds preceding F_i in permutation π. By the definition of the Shapley value,

$$\phi_i = \frac{1}{N!} \sum_{\pi \in \Pi(\mathcal{F})} \left[v(P_i(\pi) \cup \{F_i\}) - v(P_i(\pi)) \right].$$

Thus, ϕ_i is the expected marginal contribution of F_i under a uniformly random ordering of the funds.

For any coalition $\mathcal{S} \subseteq \mathcal{F}$ not containing F_i, its marginal contribution is

$$\begin{aligned} v(\mathcal{S} \cup \{F_i\}) - v(\mathcal{S}) = {} & \frac{1}{2}(\mathbf{h}_i^{\text{base}})^\top \boldsymbol{\Lambda} \mathbf{h}_i^{\text{base}} \\ & + \frac{1}{2}(\mathbf{h}_i^{\text{base}})^\top \boldsymbol{\Lambda} \mathbf{h}_{\mathcal{S}}^{\text{base}} + \frac{1}{2}(\mathbf{h}_{\mathcal{S}}^{\text{base}})^\top \boldsymbol{\Lambda} \mathbf{h}_i^{\text{base}}. \end{aligned}$$

For every $F_\ell \neq F_i$, the probability that F_ℓ precedes F_i in a uniformly random permutation is $1/2$. Taking expectations therefore yields

$$\begin{aligned} \phi_i = {} & \frac{1}{2}(\mathbf{h}_i^{\text{base}})^\top \boldsymbol{\Lambda} \mathbf{h}_i^{\text{base}} \\ & + \frac{1}{4} \sum_{\ell \neq i} \left[(\mathbf{h}_i^{\text{base}})^\top \boldsymbol{\Lambda} \mathbf{h}_\ell^{\text{base}} + (\mathbf{h}_\ell^{\text{base}})^\top \boldsymbol{\Lambda} \mathbf{h}_i^{\text{base}} \right]. \end{aligned}$$

Using $\mathbf{h}^{\text{agg,base}} = \sum_{\ell=1}^{N} \mathbf{h}_\ell^{\text{base}}$, this simplifies to

$$\phi_i = \frac{1}{4}(\mathbf{h}_i^{\text{base}})^\top \left(\boldsymbol{\Lambda} + \boldsymbol{\Lambda}^\top\right) \mathbf{h}^{\text{agg,base}}.$$

If $\boldsymbol{\Lambda}$ is symmetric, then $\boldsymbol{\Lambda} + \boldsymbol{\Lambda}^\top = 2\boldsymbol{\Lambda}$, giving

$$\phi_i = \frac{1}{2}(\mathbf{h}_i^{\text{base}})^\top \boldsymbol{\Lambda} \mathbf{h}^{\text{agg,base}}.$$

For an asymmetric $\boldsymbol{\Lambda}$, the crowding game and its Shapley allocation therefore depend only on the symmetric component $\frac{1}{2}(\boldsymbol{\Lambda} + \boldsymbol{\Lambda}^\top)$. The closed form avoids explicit enumeration of coalitions or permutations.

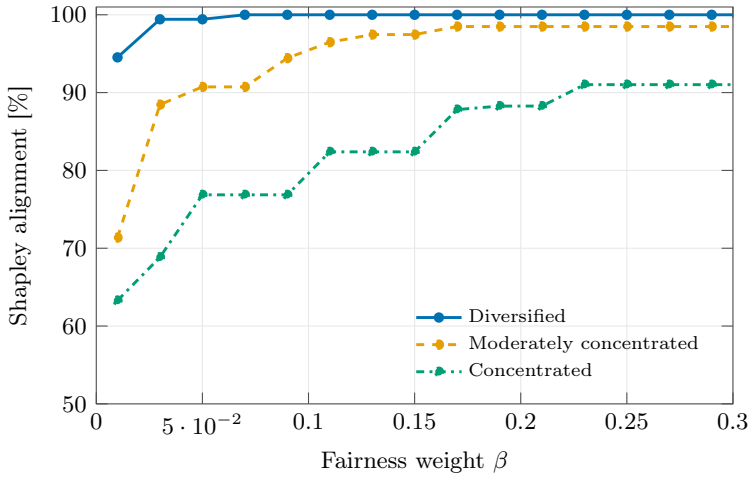

Fig. 2. Shapley alignment as a function of the fairness weight β under three portfolio structures. All economic constraints and the stress scenario are held fixed within each setting.

B.2 Sensitivity to the Fairness Weight β

We isolate the effect of the fairness weight β using plaintext alternating optimization, independently of the secure implementation. We use $N = 10$ funds and $M = 30$ assets and keep the market-impact model, stress scenario, repair target, turnover limit, and individual-participation constraints fixed throughout. Only the portfolio structure is varied. We consider three settings of increasing concentration. In the *diversified* setting, funds have moderate exposure to a common high-impact asset core and broad rebalancing flexibility. In the *moderately concentrated* setting, exposure to this common core is heterogeneous and the set of admissible substitutes is reduced. In the *concentrated* setting, several funds hold large fractions of their portfolios in the same high-impact assets and have comparatively narrow rebalancing opportunities. For each setting and each value of β, we solve the same repair problem from scratch; hence, changes in the resulting allocation can be attributed to the cost–fairness weighting rather than to changes in the underlying stress scenario.

Figure 2 shows that the value of β required to obtain a given level of fairness depends materially on the underlying portfolio structure. We measure alignment by $100A_{\text{Shapley}}\%$, so that 100% denotes exact agreement between realized loss-reduction shares and Shapley responsibility. In the diversified setting, the minimum-cost repair is already comparatively well aligned and a small positive fairness weight is sufficient to drive the alignment essentially to 100%. As concentration increases, the same alignment becomes harder to achieve: the moderately concentrated system improves progressively with β and approaches a high plateau, whereas the concentrated system exhibits both a lower initial alignment and

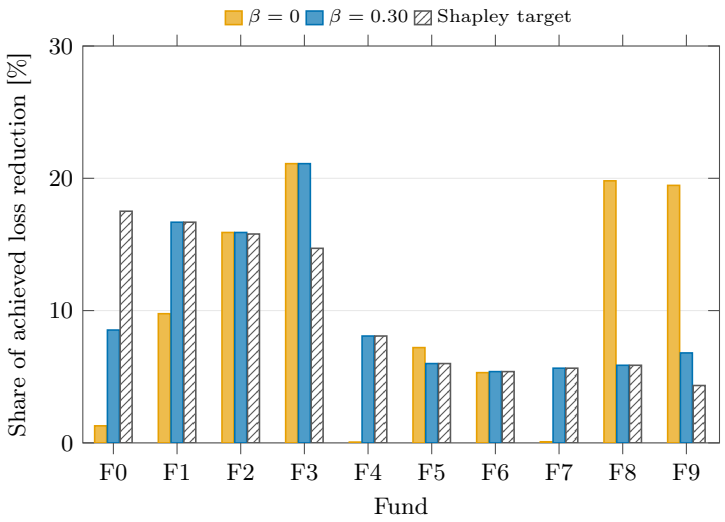

Fig. 3. Fund-level allocation of the achieved systemic-loss reduction in the concentrated setting. Bars report the realized shares G_i/G for $\beta = 0$ and $\beta = 0.30$, together with the Shapley targets s_i.

a slower improvement, reaching only about 91% over the shown range. Economically, concentration reduces the set of low-cost reallocations through which benefits can be assigned in proportion to responsibility. Consequently, a larger β is required before the optimizer is willing to depart from the minimum-cost allocation in order to improve fairness. The flattening of all three curves further indicates diminishing returns once the economically attainable Shapley allocation is approached.

Figure 3 makes the mechanism behind the aggregate alignment measure explicit. With $\beta = 0$, the optimizer minimizes adjustment cost subject to the hard repair requirements but has no incentive to distribute the resulting gain according to the Shapley shares. The allocation is therefore highly uneven: some funds receive substantially more than their responsibility-weighted target, while others receive essentially no share of the discretionary systemic improvement. In particular, the realized shares of some funds are zero or nearly zero at $\beta = 0$ even though their Shapley targets are strictly positive. When β is increased to 0.30, these funds receive positive loss reductions and the allocation moves substantially toward the Shapley reference; at the same time, funds that were strongly over-compensated under the cost-only solution receive smaller shares. The remaining deviations illustrate that Shapley alignment is not imposed as an equality constraint: it competes with adjustment cost and with the feasible portfolio and liquidation constraints. Thus, β controls a genuine economic trade-off rather than mechanically enforcing a predetermined allocation.

B.3 Fairness Saturation and Calibration of β

The plateaus in Fig. 2 suggest that the role of β can be characterized more precisely. For a fixed portfolio instance, increasing β initially trades adjustment cost for improved Shapley alignment, but beyond a critical value further increases cannot improve the best attainable fairness. We call the smallest such value the *fairness-saturation threshold*.

Let $z \in \mathcal{Z}$ denote a feasible repair and write the normalized objective as $J_\beta(z) = (1-\beta)c(z) + \beta f(z)$, where $c(z) := C(z)/C_{\text{ref}}$ is normalized adjustment cost and $f(z) := D_G(z)/(2G_{\text{ref}})$ is normalized Shapley deviation. Let $f^\star := \min_{z\in\mathcal{Z}} f(z)$ denote the smallest fairness deviation attainable under the hard economic constraints, and let $c^\star := \min\{c(z) : z \in \mathcal{Z},\ f(z) = f^\star\}$ be the minimum cost among all maximally fair repairs. Define

$$\lambda_{\text{sat}} := \sup_{\substack{z\in\mathcal{Z} \\ f(z)>f^\star}} \frac{[c^\star - c(z)]_+}{f(z) - f^\star}, \qquad \beta_{\text{sat}} := \frac{\lambda_{\text{sat}}}{1+\lambda_{\text{sat}}}, \tag{4}$$

where $[x]_+ := \max\{x, 0\}$. Thus, λ_{sat} is the largest adjustment-cost saving obtainable per unit of fairness sacrificed relative to the fairness-optimal endpoint.

Proposition (fairness saturation). Assume that $\mathcal{Z}$ is nonempty and compact, that c and f are continuous, and that $\lambda_{\text{sat}} < \infty$. Then, for every $\beta \in (\beta_{\text{sat}}, 1)$, every global minimizer z_β of J_β satisfies $f(z_\beta) = f^\star$ and $c(z_\beta) = c^\star$. Conversely, for $\beta < \beta_{\text{sat}}$, a repair attaining $(c^\star, f^\star)$ cannot be globally optimal. At $\beta = \beta_{\text{sat}}$, it may tie with a cheaper but less-fair repair.

Proof. For $\beta < 1$, multiplication of the objective by $(1-\beta)^{-1} > 0$ does not change its minimizers. Hence minimizing J_β is equivalent to minimizing $c(z) + \lambda f(z)$ with $\lambda := \beta/(1-\beta)$. Let $z^\star$ attain $(c^\star, f^\star)$. For any z with $f(z) > f^\star$, the difference between the two weighted objectives is $c(z) - c^\star + \lambda(f(z) - f^\star)$. If $c(z) \geq c^\star$, this quantity is strictly positive. If $c(z) < c^\star$, it is strictly positive whenever $\lambda > (c^\star - c(z))/(f(z) - f^\star)$. Therefore, if $\lambda > \lambda_{\text{sat}}$, every solution with $f(z) > f^\star$ is dominated by $z^\star$, while among solutions with $f = f^\star$ the objective selects the minimum cost $c^\star$. Conversely, if $\lambda < \lambda_{\text{sat}}$, the definition of the supremum implies the existence of a cheaper, less-fair repair whose cost reduction exceeds its λ-weighted fairness penalty. Since $\beta = \lambda/(1+\lambda)$ is strictly increasing in λ, the stated threshold follows. □

Empirical calibration. The proposition characterizes the global repair problem. Our plaintext implementation, however, uses alternating optimization and therefore does not provide a global-optimality certificate. We consequently estimate the threshold from the set $\widehat{\mathcal{Z}}$ of distinct repairs obtained in the β sweep. We first set $\widehat{f}^\star := \min_{z\in\widehat{\mathcal{Z}}} f(z)$ and choose the least-cost observed repair attaining this value, $\widehat{c}^\star := \min\{c(z) : z \in \widehat{\mathcal{Z}}, f(z) = \widehat{f}^\star\}$. We then replace the supremum in Eqn. (4) by the maximum over $\widehat{\mathcal{Z}}$, yielding $\widehat{\lambda}_{\text{sat}}$ and $\widehat{\beta}_{\text{sat}} = \widehat{\lambda}_{\text{sat}}/(1+\widehat{\lambda}_{\text{sat}})$. This procedure uses the full observed cost–fairness frontier rather than identifying the plateau visually.

In these experiments the achieved system-wide loss reduction remains essentially equal to its 10% target. Since $A_{\text{Shapley}} = 1 - D_G/(2G)$, we consequently have $f(z) \simeq 1 - A_{\text{Shapley}}(z)$, so minimizing the fairness term is equivalent, up to the negligible variation in G/G_{ref}, to maximizing the alignment shown in Fig. 2.

Table 4. Estimated fairness-saturation thresholds for the three portfolio structures. The final column reports the maximum observed Shapley alignment and the first sampled value of β attaining its plateau.

Setting	$\widehat{\beta}_{\text{sat}}$	Plateau alignment / first sampled β
Diversified	0.061	100.00% / 0.07
Moderately concentrated	0.164	98.49% / 0.17
Concentrated	0.226	91.02% / 0.23

The estimated saturation threshold rises markedly with portfolio concentration, from approximately 0.061 in the diversified system to 0.164 and 0.226 in the moderately concentrated and concentrated systems, respectively. Thus concentration affects not only the maximum attainable alignment but also the relative weight that must be placed on fairness before the optimizer selects the fairest observed repair. In each case, $\widehat{\beta}_{\text{sat}}$ lies just below the first sampled value on the corresponding plateau, providing a quantitative explanation for the flattening in Fig. 2. Beyond this threshold, increasing β changes the relative objective weights but yields no further fairness improvement on the observed cost–fairness frontier. Because the joint problem is bilinear and the alternating solver is not globally certified, the reported values are empirical estimates of the global threshold characterized in Eqn. (4).

C Privacy-Preserving Portfolio Repair

In the following we give more details about the PDHG block solver employed in the MPC-based solution and the federated approach.

C.1 PDHG Block Solver

Each step of the outer alternation in Eqn. (3) is an LP once the other decision block is fixed. To describe the common solver, collect the variables of the current block in $\mathbf{z}$ and write

$$\min_{\mathbf{z}\in\mathcal{B}} \mathbf{c}^\top \mathbf{z} \quad \text{s.t.} \quad \mathbf{A}\mathbf{z} \le \mathbf{b}, \qquad \mathbf{E}\mathbf{z} = \mathbf{e}\,, \tag{5}$$

where $\mathcal{B}$ contains simple variable bounds. For the liquidation block, $\mathbf{z}$ contains $\mathbf{X}$ and its auxiliary variables; for the rebalancing block, it contains $\widetilde{\mathbf{H}}$ and the corresponding auxiliaries.

Let $\mathbf{y} \geq 0$ and $\mathbf{w}$ denote the dual variables associated with the inequality and equality constraints. We apply the primal–dual hybrid gradient (PDHG) method [11]. One iteration is

$$\mathbf{y}^{(r+1)} = \Big[\mathbf{y}^{(r)} + \sigma\big(\mathbf{A}\bar{\mathbf{z}}^{(r)} - \mathbf{b}\big)\Big]_+, \tag{6}$$

$$\mathbf{w}^{(r+1)} = \mathbf{w}^{(r)} + \sigma\big(\mathbf{E}\bar{\mathbf{z}}^{(r)} - \mathbf{e}\big), \tag{7}$$

$$\mathbf{z}^{(r+1)} = \mathit{\Pi}_{\mathcal{B}}\Big(\mathbf{z}^{(r)} - \tau_{\mathrm{p}}\big(\mathbf{c} + \mathbf{A}^{\top}\mathbf{y}^{(r+1)} + \mathbf{E}^{\top}\mathbf{w}^{(r+1)}\big)\Big)\,, \tag{8}$$

followed by

$$\bar{\mathbf{z}}^{(r+1)} = \mathbf{z}^{(r+1)} + \omega\big(\mathbf{z}^{(r+1)} - \mathbf{z}^{(r)}\big)\,. \tag{9}$$

The step sizes $\tau_{\mathrm{p}}, \sigma$ and extrapolation parameter ω are public. Under MPC, the matrix–vector products and arithmetic operations are evaluated on secret shares, while the projections are implemented through secure comparisons. No secret pivot selection or secret division is required. We use a fixed public number of inner iterations, so neither execution length nor control flow depends on private convergence behavior.

The same solver is used in both stages of portfolio repair. During Phase I, a non-negative slack η relaxes the systemic-loss and individual-improvement requirements, and the alternating blocks minimize the remaining violation while all other constraints remain active. After the fixed Phase-I schedule completes, Phase II begins unconditionally, restores the hard constraints, and alternates the liquidation and rebalancing blocks while minimizing Eqn. (2).

C.2 Federated PDHG

The federated construction preserves the outer alternation and the PDHG structure above, but separates fund-local operations from quantities that depend jointly on several funds. For a fixed block, let $\mathbf{z}_i$ collect the variables held by fund F_i. The block can be written schematically as

$$\min_{\mathbf{z}_1,\dots,\mathbf{z}_N} \sum_{i=1}^{N} f_i(\mathbf{z}_i) \quad \text{s.t.} \quad \mathbf{z}_i \in \mathcal{Z}_i, \qquad \sum_{i=1}^{N} \mathbf{A}_i\mathbf{z}_i \leq \mathbf{b}\,,$$

with equality constraints handled analogously. Here $\mathcal{Z}_i$ contains constraints depending only on fund F_i, while the sum represents constraints coupled through system-wide quantities.

At PDHG iteration r, each fund evaluates its local contribution and secret-shares it to the MPC parties. The MPC computation forms

$$\mathbf{a}^{(r)} = \sum_{i=1}^{N} \mathbf{A}_i\bar{\mathbf{z}}_i^{(r)}\,,$$

and reconstructs the resulting aggregate to the participating funds. Consequently, the global residual $\mathbf{r}^{(r)} = \mathbf{a}^{(r)} - \mathbf{b}$ and the corresponding PDHG coordination

state can be updated consistently by all participants, e.g.,

$$\mathbf{y}^{(r+1)} = \left[\mathbf{y}^{(r)} + \sigma \mathbf{r}^{(r)}\right]_+ .$$

Fund F_i then performs its local primal update

$$\mathbf{z}_i^{(r+1)} = \Pi_{\mathcal{Z}_i}\left(\mathbf{z}_i^{(r)} - \tau_{\mathrm{p}}\big(\mathbf{c}_i + \mathbf{A}_i^\top \mathbf{y}^{(r+1)}\big)\right) ,$$

using its private state.

For the liquidation block, $\mathbf{z}_i$ contains the local liquidation $\mathbf{x}_i$ and the principal aggregate is $\mathbf{q}^{(r)} = \sum_{i=1}^N \mathbf{x}_i^{(r)}$. For the rebalancing block, $\mathbf{z}_i$ contains the local repaired holdings and associated auxiliary variables, with $\widetilde{\mathbf{h}}^{\mathrm{agg},(r)} = \sum_{i=1}^N \widetilde{\mathbf{h}}_i^{(r)}$. Additional quantities required by the loss, fairness, and feasibility conditions are aggregated in the same way. Individual fund contributions are not reconstructed; only the aggregate quantities required by the iterative solver are released.

Thus, the federated protocol moves fund-local projections and linear algebra out of MPC while retaining secure aggregation of cross-fund contributions. The price of this reduction in secure computation is that the sequence of released aggregates becomes part of the federated transcript; its leakage is analyzed in App. D.

D Security and Leakage Analysis

The simulation guarantee of Sec. 6 protects all values that are not explicitly reconstructed. Here we distinguish leakage inherent in authorized outputs from the additional transcript exposed by the federated construction.

D.1 Output leakage

Let U_i denote the private input of fund F_i, let $U_{\mathcal{C}}$ denote the inputs of coalition $\mathcal{C}$, and let P denote public information. If the coalition observes

$$O_{\mathcal{C}} := (\mathcal{O}_{\mathrm{pub}}, \{\mathcal{O}_i\}_{i\in\mathcal{C}}) ,$$

then for an honest fund $F_j \notin \mathcal{C}$ we define the information contained in the authorized outputs as

$$\mathcal{L}_{\mathrm{out}}^{j\to\mathcal{C}} := I(U_j; O_{\mathcal{C}} \mid U_{\mathcal{C}}, P) .$$

This quantity captures inference inherent in providing the requested result and is distinct from protocol leakage. For example, a released correlation matrix reveals pairwise dependence values, whereas releasing only a threshold network or an MST reveals a strictly coarser diagnostic. We therefore follow an output-minimization principle and release only the granularity required by the application.

D.2 Federated leakage

The federated construction intentionally adopts a weaker privacy boundary than pure MPC. Here, bounded collusion means $|\mathcal{C}| \leq N - 2$, ensuring that at least two funds remain non-colluding and thereby excluding the degenerate case in which a released aggregate reveals the sole remaining fund's contribution exactly by subtraction. Individual fund contributions remain protected during secure aggregation, but aggregate coordination signals required by the local updates are reconstructed to the participating funds across iterations. Such iterative signals may permit inference about private local data, as studied in the distributed-optimization literature [40]. Existing approaches mitigate this leakage by perturbing or otherwise modifying the exchanged optimization state while preserving useful convergence properties [39]. These techniques provide a natural extension to our federated construction. We leave the adaptation of these techniques to the portfolio-repair setting and the quantification of the resulting privacy–utility trade-off to future work.

Secure aggregation alone cannot prevent inference that follows from an aggregate once that aggregate is deliberately released. Outside the assumed bound, if a coalition observes $\mathbf{a} = \sum_i \mathbf{v}_i$ and all but one fund collude, the remaining contribution is recovered exactly as $\mathbf{v}_j = \mathbf{a} - \sum_{i \in \mathcal{C}} \mathbf{v}_i$. We therefore use the following single-aggregate experiment as an illustrative identifiability baseline.

Single-aggregate identifiability baseline. We illustrate the effect of collusion on the representative $N = 10$, $M = 200$ portfolio instance using a single-aggregate identifiability stress test. The test assumes that a coalition can derive the per-asset aggregate holdings of the remaining non-colluding funds. For each number k of non-colluding funds, we enumerate every subset $\mathcal{H} \subseteq \mathcal{F}$ with $|\mathcal{H}| = k$, every target fund $F_j \in \mathcal{H}$, and all M assets. For each resulting target–asset observation, define

$$X_{ja} := \mathbf{1}\{h_{ja} > 0\}, \qquad Z_{\mathcal{H}a} := \mathbf{1}\left\{\sum_{i \in \mathcal{H}} h_{ia} > 0\right\}.$$

Thus, X indicates whether the target fund holds an asset, whereas Z indicates only whether the aggregate position of the non-colluding funds in that asset is non-zero. Importantly, the aggregate amount itself is discarded.

For each k, we compute the empirical mutual information $I_k(X; Z)$ from all such binary observations. Let

$$\hat{p} := \frac{1}{NM} \sum_{j \in \mathcal{F}} \sum_{a=1}^{M} \mathbf{1}\{h_{ja} > 0\}$$

denote the empirical probability that a fund holds an asset, and let

$$h_2(\hat{p}) := -\hat{p} \log_2 \hat{p} - (1 - \hat{p}) \log_2 (1 - \hat{p})$$

denote the corresponding binary entropy. We also report

$$\frac{I_k(X; Z)}{h_2(\hat{p})}.$$

Because every fund–asset pair appears equally often among the target observations for any fixed k, the marginal distribution of X, and hence $h_2(\hat{p})$, is independent of k. The normalized quantity therefore measures the fraction of the target's prior support uncertainty removed by observing the aggregate-presence indicator.

Finally, we report the fraction of portfolio coordinates that are determined exactly from the aggregate using only non-negativity. If $k = 1$, subtraction of the colluding funds' holdings reveals the complete target portfolio. If $k > 1$, an aggregate value of zero implies

$$\sum_{i \in \mathcal{H}} h_{ia} = 0 \quad \Longrightarrow \quad h_{ia} = 0 \quad \text{for every } i \in \mathcal{H}\,,$$

so the corresponding target coordinate is known exactly.

Table 5. Single-aggregate identifiability baseline on the $N = 10$, $M = 200$ benchmark. The coalition is assumed to derive the per-asset aggregate of the remaining non-colluding funds. Only the binary presence of the aggregate is used to compute support leakage.

Non-colluding funds k	Colluding funds	Support leakage $I_k(X;Z)$ [bits]	Normalized leakage $I_k(X;Z)/h_2(\hat{p})$ [%]	Coordinates det. exactly [%]
1	9	0.9345	100.0	100.00
2	8	0.3849	41.2	42.89
3	7	0.2160	23.1	28.11
4	6	0.1269	13.6	17.98
5	5	0.0740	7.9	11.05
6	4	0.0414	4.4	6.38
7	3	0.0213	2.3	3.34
8	2	0.0093	1.0	1.48
9	1	0.0028	0.3	0.45
10	0	0.0000	0.0	0.00

With nine colluding funds, the remaining aggregate identifies the target portfolio exactly. Even with two non-colluding funds, the deliberately reduced binary observable removes 41.2% of the target's empirical support uncertainty and determines 42.89% of coordinates exactly. This experiment is not a privacy bound for the complete federated protocol. Iterative coordination signals may enable additional inference. Existing privacy-preserving distributed-optimization techniques can in principle be used to reduce this leakage by perturbing the exchanged optimization state [39], at the cost of an additional privacy–utility trade-off. Our experiment instead isolates the baseline information exposed by aggregate release itself.

E Additional Evaluation of Private Portfolio Analysis

E.1 Benchmark Funds

Table 6 reports the ten funds used in the representative experiment. The benchmark deliberately contains funds with different sizes, concentration levels, and liquidity buffers. N-PORT [55] is used to construct a reproducible benchmark from real portfolios; the protocols themselves do not depend on the participating institutions being registered funds and apply equally when the holdings are supplied privately by other asset managers.

Table 6. Funds in the representative N-PORT benchmark. V_i^{base} denotes the reported fund value used as the redemption base, $B_{i,200}^{\text{base}}$ the risky exposure represented in the selected 200-asset universe, c_i the cash position used by the model, and $D_i/B_{i,200}^{\text{base}}$ the stressed liquidation requirement under $\rho_{\text{red}} = 5\%$. Monetary quantities are in USD billions. Because $B_{i,200}^{\text{base}}$ includes only the selected asset universe, it may be smaller than the fund's total risky exposure B_i^{base}. Hence $c_i + B_{i,200}^{\text{base}} \le V_i^{\text{base}} = c_i + B_i^{\text{base}}$.

	Fund	SEC series	V_i^{base}	$B_{i,200}^{\text{base}}$	c_i	$D_i/B_{i,200}^{\text{base}}$ [%]
P_0	Growth Fund	S000008791	49.59	36.28	2.11	1.02
P_1	College Retirement Equities Fund – Growth Account	S000005082	37.62	35.85	0.07	5.06
P_2	Fidelity Series Large Cap Stock Fund	S000039214	24.80	19.34	1.15	0.49
P_3	Franklin Dynatech Fund	S000006754	24.77	22.11	0.28	4.36
P_4	Franklin Rising Dividends Fund	S000006858	25.78	23.15	0.28	4.34
P_5	JPMorgan Hedged Equity Fund	S000043249	20.06	16.97	1.70	0.00
P_6	JPMorgan Large Cap Growth Fund	S000003856	107.43	98.78	5.13	0.24
P_7	State Street Financial Select Sector SPDR ETF	S000006411	47.79	40.24	0.13	5.62
P_8	State Street Technology Select Sector SPDR ETF	S000006415	84.20	77.99	0.10	5.27
P_9	Vanguard PRIMECAP Fund	S000002568	72.40	54.32	0.70	5.38

The sample spans net asset values from approximately \$20 bn to \$107 bn and exhibits substantial heterogeneity in liquidity buffers. For example, P_5 has sufficient cash to absorb the modeled redemption shock without forced selling, whereas P_7–P_9 must liquidate more than 5% of their selected risky exposure. The selected universe nevertheless retains at least 83.8% coverage under our asset-selection criterion; 95.5% of selected assets are shared by at least two funds, and mean pairwise holdings cosine similarity is 0.476.

E.2 Private Network Analysis

For the representative instance we use $T = 60$ asset-return observations and the public correlation threshold $\tau = 0.9$. The secure computation produces the complete set of portfolio and network statistics shown in Tab. 7.

Table 7. Per-fund output of the private portfolio analysis for $N = 10$, $M = 200$, $T = 60$, and $\tau = 0.9$. $\deg_\tau$ and str_τ denote threshold-network degree and correlation strength; $\deg_{\mathrm{MST}}$ denotes MST degree.

Fund	$\bar{\rho}$	$\rho_{\max}$	Nearest	$\deg_\tau$	str_τ	$\deg_{\mathrm{MST}}$	Component size	Shapley share [%]
P_0	0.884	0.967	P_1	6	5.660	1	7	6.74
P_1	0.883	0.977	P_3	6	5.750	3	7	7.78
P_2	0.888	0.973	P_5	6	5.623	1	7	3.79
P_3	0.864	0.977	P_1	6	5.681	2	7	4.46
P_4	0.801	0.889	P_5	0	0	2	1	5.32
P_5	0.903	0.973	P_2	6	5.702	5	7	3.51
P_6	0.875	0.962	P_5	6	5.639	1	7	25.98
P_7	0.594	0.683	P_4	0	0	1	1	11.71
P_8	0.850	0.959	P_3	6	5.548	1	7	15.88
P_9	0.813	0.871	P_5	0	0	1	1	14.83

The results illustrate why the different statistics are complementary. At $\tau = 0.9$, seven funds form a dense connected component, whereas P_4, P_7, and P_9 are isolated in the threshold graph. P_1 has the largest thresholded correlation strength, but it is not the fund with greatest structural responsibility. Instead, P_6 accounts for approximately 26% of holdings-based Shapley crowding. Conversely, P_5 has the largest MST degree, connecting five other funds in the sparse representation despite having only a 3.5% Shapley share. Thus, return centrality, sparse-network position, and holdings-based systemic responsibility capture genuinely different properties.

E.3 Threshold Sensitivity

For this instance, the qualitative network structure is not unique to $\tau = 0.9$. Raising the threshold progressively removes weaker links, but a seven-fund core persists from $\tau = 0.9$ through $\tau = 0.95$. The threshold therefore controls the sparsity of the reported network without eliminating the strongly correlated core observed in this instance.

E.4 Computational Cost

The complete representative analysis finishes in 7.50 seconds. Table 9 separates the principal components.

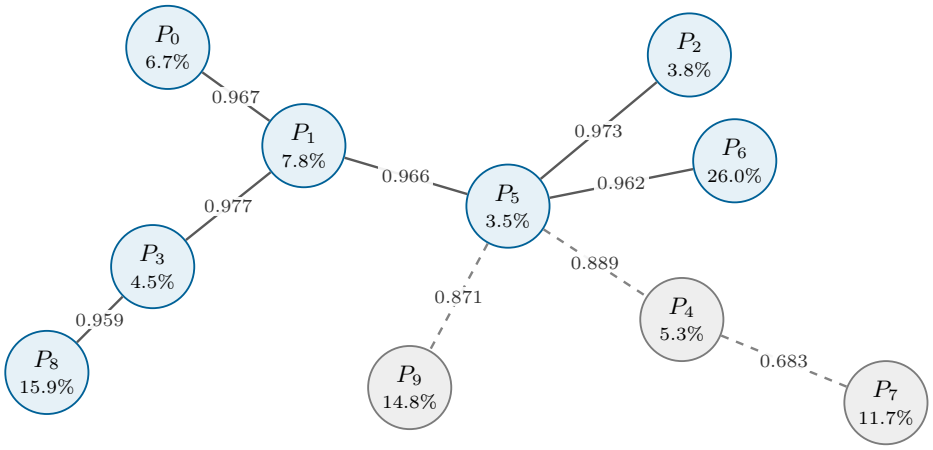

Fig. 4. Maximum spanning tree of portfolio-return correlations for the representative instance. Edge labels report Pearson correlation. Node percentages report holdings-based Shapley responsibility. Blue nodes form the seven-fund connected component of the threshold graph at $\tau = 0.9$; gray nodes are isolated at this threshold. Dashed MST edges have correlation below $\tau = 0.9$. Fund identifiers correspond to Table 6.

Table 8. Sensitivity of the portfolio-correlation network to the public edge threshold τ. Density is the fraction of the 45 possible undirected fund pairs retained.

τ	Edges	Density [%]	Components	Largest component
0.80	34	75.6	2	9
0.85	26	57.8	2	9
0.875	22	48.9	3	8
0.90	21	46.7	4	7
0.925	15	33.3	4	7
0.95	10	22.2	4	7

F Additional Portfolio-Repair Evaluation

F.1 Solver Configuration

The joint portfolio-repair problem is bilinear, but becomes an LP when either holdings or contingency liquidations are fixed. Plaintext HiGHS [34] therefore serves as an exact-block reference: it solves each LP arising in the alternating procedure exactly, rather than solving the joint problem globally.

The secure implementations instead use fixed public iteration budgets, which make the optimization trace independent of private convergence behavior. Federated PDHG uses 128 inner iterations per block solve and one Phase-I followed by nine Phase-II outer rounds. We selected 128 as the smallest tested inner budget that reliably satisfies the hard constraints while remaining close to the converged objective. Increasing the budget to 144 changes the final objective

Table 9. Runtime and communication cost breakdown for private analysis with $N = 10$, $M = 200$, and $T = 60$.

Component	Time [s]	Data [MB]
Portfolio-return construction	0.02	0.09
Division-free Prim MST	3.17	58.24
Full Pearson matrix	0.45	14.16
Network statistics, clusters, Shapley	2.89	30.88
Full protocol	7.50	230.67

by only 1.1%, at the cost of 10.6% additional secure-aggregation work. The objective likewise stabilizes after nine Phase-II rounds. MPC PDHG uses fixed stage-specific iteration budgets throughout.

All secure arithmetic uses fixed-point representation over a 128-bit ring. For the representative instance, the 20-bit fractional representation gives a quantum of approximately 9.54×10^{-7}. Final solutions are explicitly checked for budget preservation, sale feasibility, holding bounds, turnover, price-impact, systemic-loss, and individual-improvement constraints before release.

F.2 Dataset Diagnostics

The representative $N = 10, M = 200$ instance was selected to balance realistic portfolio overlap with sufficient heterogeneity for meaningful rebalancing. In addition to the statistics reported in the main text, the maximum pairwise holdings similarity is 0.897, so the set contains some strongly overlapping portfolios without making all funds effectively identical.

The stressed liquidity requirements are also well within the feasible portfolio-repair regime: the largest required liquidation is 5.6% of a fund's risky-asset budget. Several funds retain substantial cash buffers and therefore require little or no direct liquidation, introducing heterogeneity in how funds contribute to systemic risk reduction.

Market-impact coefficients span approximately $\Lambda_{aa} \in [0.0030, 0.655]$, while the median daily trading volume used in the liquidity calibration is approximately \$933 m. Transaction-cost coefficients lie between 1.61×10^{-4} and 3.12×10^{-4}. These magnitudes provide adequate numerical resolution under the 20-bit precision fixed-point representation while avoiding the extreme price-cap saturation of more illiquid candidate universes.

F.3 Preprocessing Requirements

Table 10 reports the dominant preprocessing requirements of the representative compiled circuit with $N = 10$ and $M = 200$.

Table 10. Preprocessing requirements for the 10×200 instance. L-edaBits and S-edaBits denote loose and strict extended daBits (edaBits) [26], respectively, with the subscript indicating their bit length. daBits are doubly-authenticated bits introduced by Rotaru and Wood [47], while triples denote Beaver multiplication triples [15].

Arith. triples	Bool. triples	daBits	L-edaBits$_{51}$	S-edaBits$_{20}$	S-edaBits$_{50}$
140,459,788	3,721,516	62,792,056	62,768,584	63,692,978	63,739,943

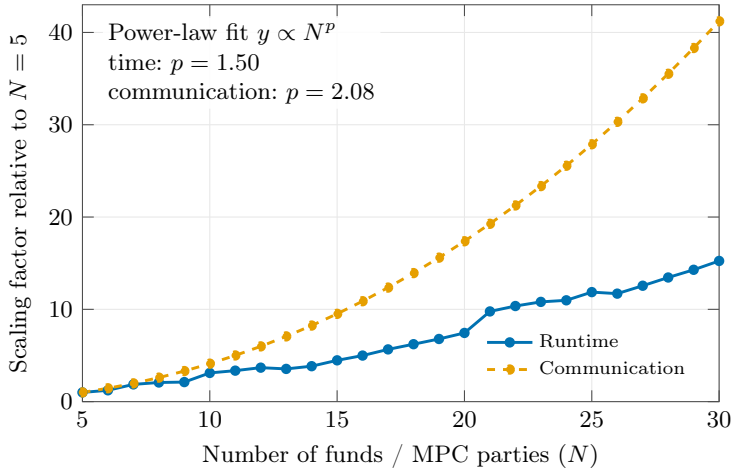

Fig. 5. Measured scaling of the pure-MPC portfolio-repair implementation with the number of funds / MPC parties N for a fixed asset universe of $M = 50$. Both metrics are reported relative to the $N = 5$ baseline. The results show superlinear growth in wall-clock time and approximately quadratic growth in communication.

F.4 Party Scalability

Figure 5 shows that the pure-MPC implementation becomes substantially more expensive as the number of participating funds increases, even when the asset universe is held fixed at $M = 50$. Relative to the $N = 5$ baseline, wall-clock time grows by roughly a factor of $15.2\times$ at $N = 30$, while total communication increases by about $41.2\times$. A power-law fit gives an exponent of $p \approx 1.50$ for runtime, indicating superlinear but subquadratic time growth, whereas communication is much closer to quadratic with $p \approx 2.08$. This suggests that communication, rather than local arithmetic, is the dominant factor limiting scalability as the number of MPC parties grows.

F.5 Asset Scalability

We complement the fund-scaling experiment in the main text by fixing $N = 10$ and varying the number of assets from $M = 50$ to $M = 500$.

Table 11. Scalability of portfolio repair with the number of assets for $N = 10$ funds. Communication denotes data transmitted by one fund party; for the federated construction, this includes both MPC communication and the fund–MPC client channel.

$N = 10$	Communication [GB]				Runtime [min]			
	$M = 50$	$M = 100$	$M = 200$	$M = 500$	$M = 50$	$M = 100$	$M = 200$	$M = 500$
MPC	19.298	36.968	72.331	178.433	15.91	24.76	27.65	47.13
Federated	0.052	0.101	0.198	0.490	7.62	7.65	7.72	8.86

Communication grows approximately linearly with M for both constructions, but at very different scales. Across all four sizes, Federated PDHG communicates roughly 364–372× less data per fund party. More importantly, increasing the asset universe tenfold from 50 to 500 raises MPC-PDHG runtime from 15.91 to 47.13 minutes, whereas Federated PDHG increases only from 7.62 to 8.86 minutes. The corresponding runtime advantage therefore grows from 2.1× to 5.3×.

The relatively weak dependence of federated runtime on M results from moving the repeated fund-local vector operations outside MPC. Secure coordination still depends on the asset dimension, as reflected in the approximately linear communication growth, but it no longer dominates the full local optimization workload.

F.6 Network Sensitivity

Table 12. Runtime of portfolio repair under different network conditions for the representative $N = 10$, $M = 200$ instance. Local denotes 0.1 ms RTT and 30 Gbit/s; WAN-10 denotes 10 ms RTT and 10 Gbit/s; WAN-20 denotes 20 ms RTT and 1 Gbit/s; and WAN-40 denotes 40 ms RTT and 500 Mbit/s.

$N = 10,\ M = 200$	Runtime [min]			
	Local	WAN-10	WAN-20	WAN-40
MPC	27.65	154.99	279.50	526.38
Federated	7.72	12.68	17.39	20.73

We evaluate the representative $N = 10, M = 200$ instance under four network profiles in Table 12. Network conditions have a markedly different effect on the two constructions. MPC PDHG increases from 27.65 minutes locally to 154.99 minutes under WAN-10 and 279.50 minutes under WAN-20. Federated PDHG rises much more slowly, from 7.72 to 12.68 and 17.39 minutes, respectively. Its runtime advantage consequently grows from 3.6× locally to 12.2× under WAN-10 and 16.1× under WAN-20.

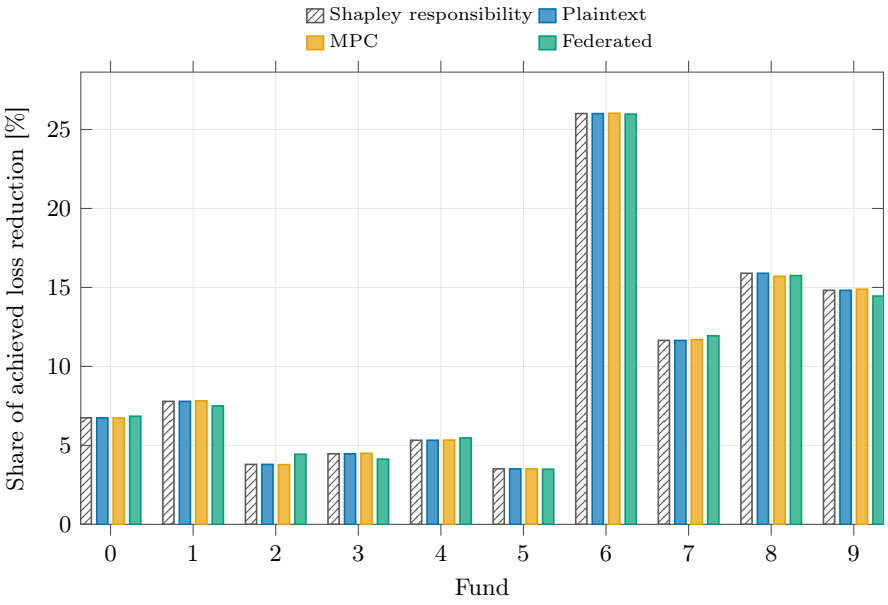

Fig. 6. Responsibility-aligned repair on the representative $N = 10$, $M = 200$ instance. Bars compare each fund's initial Shapley responsibility s_i with its realized share of achieved systemic-loss reduction, G_i/G, under plaintext HiGHS, pure-MPC PDHG, and federated PDHG.

This behavior is consistent with the interaction structure of the two protocols. The MPC-PDHG WAN runs spend more than 95% of their wall-clock time in communication, so round complexity rather than local arithmetic becomes the dominant bottleneck. Federated PDHG reduces the amount of repeated work carried out interactively while preserving secure aggregation and global coordination.

F.7 Detailed Solution Quality

Table 2 summarizes the system-level solution quality. Here we examine whether the same agreement is also visible at the fund and asset levels.

We report Shapley alignment as $100A_{\text{Shapley}}\%$. An alignment of 100% therefore indicates exact agreement between the realized loss-reduction shares G_i/G and the Shapley responsibility shares s_i.

The plaintext solution is effectively coincident with the Shapley target. For MPC PDHG, the realized reduction shares differ from the target by only 0.04 percentage points per fund on average and at most 0.20 percentage points. Federated PDHG remains close despite performing much of the optimization locally: its mean absolute deviation is approximately 0.23 percentage points and its maximum deviation 0.64 percentage points. Hence, the aggregate alignment values in Tab. 2 are not hiding large errors for individual funds.

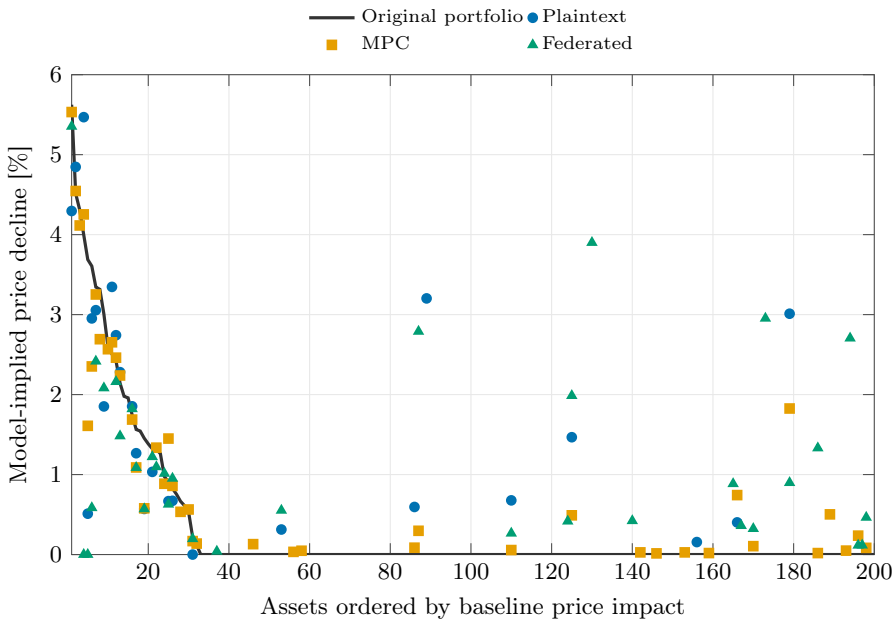

Fig. 7. Asset-level model-implied price impact before and after portfolio repair on the representative $N = 10$, $M = 200$ N-PORT instance. Assets are ordered by decreasing baseline impact. The black line shows the original price-impact profile; markers show repaired impacts only for assets whose impact changes by more than 0.01 percentage points.

The three methods also produce different, but economically comparable, asset-level repair paths.

Figure 7 orders assets by decreasing price impact under the original portfolio. The solid line denotes this original profile, while markers denote repaired impacts only where the absolute change exceeds 0.01 percentage points. This threshold suppresses numerically negligible changes and makes the structure of the repairs visible.

At this threshold, plaintext HiGHS modifies 25 assets, MPC PDHG 42, and Federated PDHG 35. The exact-block method therefore reaches a comparatively sparse repair, whereas the first-order methods spread smaller adjustments over more assets. Importantly, this difference in the path taken by the optimizer does not translate into a qualitatively different outcome: systemic-loss reduction remains near 10%, minimum individual improvement near 6%, and Shapley alignment above 98.8% for all reported solutions.

The repair-cost differences should likewise be interpreted in light of the outer bilinear optimization. Plaintext HiGHS minimizes each fixed-block LP exactly, but the complete problem contains products of aggregate holdings and liquidations and is therefore non-convex. Different block solvers and finite-iteration trajectories

can consequently reach different outer fixed points. The \$0.834 m cost of MPC PDHG, compared with \$0.919 m for plaintext HiGHS, is therefore not evidence that the approximate solver outperforms an exact global optimizer. Federated PDHG reaches a somewhat more expensive \$1.349 m repair, but remains well within the same economic scale while providing the substantially lower secure-computation cost reported above.