mathematics

MDPI

Article

EquiPartition: Learning Player-Adaptive Deviation Certificates for Complete Enclosure of Approximate Nash Equilibria in Strategic Financial Markets

Jinjia Chen [1], Kun Yang [2], Danhua Liu [3], Yingchao Wang [4,*] and Feng Xue [2,*]

[1] Binary Business School, Binary University of Management & Entrepreneurship, Puchong 47100, Malaysia
[2] Graduate School of Information, Production and Systems, Waseda University, Kitakyushu 808-0135, Japan
[3] Graduate School of Business and Finance, Waseda University, Tokyo 169-8050, Japan
[4] School of Intelligent Manufacturing and Electrical Engineering, Xinjiang College of Science & Technology, Korla 841000, China
* Correspondence: w_yingchao_xjust@163.com (Y.W.); xue.feng89@toki.waseda.jp (F.X.)

Abstract

Strategic financial markets can admit multiple approximate Nash equilibria with materially different liquidation and price-impact regimes. Returning one equilibrium cannot exclude a more stressed alternative, yet complete enclosure becomes expensive when the profitable deviator changes inside a search region. We propose EquiPartition, a learning-augmented branch-and-bound method based on player-adaptive deviation certificates. The certificate combines institution-specific unilateral-deviation atoms through state-dependent Bernstein weights, while an outward-rounded verifier alone authorizes exclusion. Inconclusive proposals undergo bounded repair or conservative subdivision, so prediction quality affects efficiency but not the complete ε-Nash outer-enclosure guarantee. On the controlled and multi-institution synthetic one-period fire-sale markets, the finite recorded-reference audit retained all recorded high-precision reference-equilibrium components and found no accepted exclusion intersecting them. Relative to Nash-B&B, Predict+Repair reduced proof-tree nodes by 83.3% and 83.9%, unresolved volume by 75.8% and 69.0%, and mean capped runtime by 40.5% and 20.2%; it completed 113 of 120 multi-institution instances versus 103 for Nash-B&B. The same finite recorded-reference audit retained all recorded crisis components under the financial stress shifts. Runtime decomposition showed that amortized proposals avoided repeated neutral-start optimization and offset repair and verification costs. EquiPartition reduces proof fragmentation while preserving verifier-based completeness.

Keywords: approximate Nash equilibrium; equilibrium enclosure; deviation certificate; strategic financial markets; fire-sale equilibrium; systemic risk; branch and bound; learning-augmented optimization

MSC: 90C33; 90C57; 91G45; 68T05

Check for updates

Academic Editor: Banafsheh Behzad

Received: 6 September 2026
Revised: 20 October 2026
Accepted: 31 October 2026
Published: 4 November 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

1. Introduction

Strategic financial markets connect institutions through shared assets, price impact, funding needs, and network obligations. Classical clearing models explain shock propagation through liabilities, while price-impact and deleveraging models show how one institution's liquidation changes market values and other institutions' incentives [1–3].

Fire-sale games make this feedback strategic and may admit several equilibria under identical market conditions [4]. These equilibria can range from limited liquidation to severe price declines and broad balance-sheet losses. A solver returning one equilibrium therefore cannot rule out a more stressed equilibrium.

Most financial-network algorithms compute a selected clearing state, fixed point, or intervention outcome. Work on minimal clearing, derivatives, and generalized payment functions identifies computable states under richer rules [5–7]. Debt forgiveness, debt transfers, and claims trading optimize interventions or network changes [8–10]. Distributed dynamics and integer-programming formulations also reach or select equilibria under structural assumptions [11–13]. These methods answer important selection and intervention questions, but finitely many returned states cannot establish that no other approximate equilibrium exists in an unexplored part of a continuous strategy space.

Complete equilibrium methods address coverage by searching the strategy domain. Branch and prune and spatial branch and bound maintain boxes in the joint strategy space and discard a box only after a rigorous test proves that it contains no target equilibrium [14–16]. Unresolved boxes are subdivided, and the terminal boxes form an outer enclosure. Global formulations extend computation to mixed-integer, quasi-linear, rational, and polynomial generalized Nash settings [17–20]. Their cost still depends on regional proof strength: each failed certificate adds children, bound calculations, and verifier calls.

A major difficulty is that the institution providing a profitable unilateral deviation may change within one box. Many regional tests fix one player and require a sufficiently profitable deviation throughout the box [15], whereas nonequilibrium only requires some player to improve at each state. In fire-sale markets, the responsible institution can change with liquidation, price impact, and balance-sheet exposure. If the switch follows a slanted, curved, or higher-dimensional boundary, no fixed-player proof may cover the box, forcing an axis-aligned solver to represent the switch with many smaller boxes.

Existing methods offer partial ingredients. Nikaido–Isoda constructions and regularized merit functionals aggregate player-wise gains [21,22], while variational-inequality cuts and moment–sum-of-squares hierarchies tighten global equilibrium representations [23,24]. Adjustable robust optimization represents state-dependent decisions through affine policies or finite candidate actions [25,26], with extensions to stochastic settings, discrete uncertainty, and approximation guarantees [27–29]. However, these constructions do not directly provide a box-wide Nash exclusion certificate that remains rigorously valid when responsibility changes across institutions.

Learning can reduce repeated computation in exact solvers. Transfer learning in spatiotemporal recognition likewise illustrates how reusable representations can be adapted across related instances [30]. Learning-based coordination also addresses constrained cooperative planning and communication-aware collaboration in multi-agent systems [31,32]. It has been used for secure multi-agent control [33], cut selection, branching, node comparison, and separator configuration [34–37]; structural analyses study generalization across branch-and-cut instances [38]. Recent systems guide K-adaptability, repeatedly solved bilevel optimization, and branch-and-bound decisions [39–41]. Structured network-management systems similarly use graph attention for joint alarm analysis and schema-grounded agents for multi-vendor configuration [42,43]. Learning-augmented algorithms and learning-guided reasoning also combine predictions with worst-case safeguards or formal checks [44,45]. The remaining challenge is to synthesize strong regional certificates without letting proposal, repair, and verification costs outweigh saved branching work.

To address these challenges, we introduce EquiPartition, a complete enclosure algorithm based on player-adaptive deviation certificates. Each certificate contains a small set of institution-specific unilateral-deviation atoms and state-dependent responsibility weights.

Different institutions can support the exclusion proof in different parts of the same box while the certificate remains one continuous regional object. A rigorous outward-rounded verifier removes a box only when the certified lower bound exceeds the equilibrium tolerance.

A learned proposal map predicts certificate parameters for each market–region pair, and a bounded repair step improves proposals that are not immediately certified. Unsuccessful proposals return to conservative subdivision. Learning therefore affects how quickly a useful certificate is found, while the verifier alone determines whether a region can be removed.

The analysis relates the certificate to maximum unilateral regret, proves complete outer enclosure under verifier-authorized deletion, and characterizes the spatial cost of fixed-player proofs near non-axis-aligned responsibility switches. The experiments evaluate recorded-reference coverage, certificate coverage, proof-tree size, proposal and repair cost, verification cost, unresolved-region quality, and end-to-end runtime on controlled switching geometries and synthetic one-period fire-sale games. The contributions are as follows:

- We formulate a rigorously verifiable player-adaptive regional deviation certificate that allows deviation responsibility and deviation actions to vary with the strategy state inside one continuous proof object.
- We establish the certificate's soundness and the resulting complete ε-equilibrium outer enclosure, and derive an axis-aligned representation-size lower bound for fixed-player certification near a non-axis-aligned responsibility-switching interface.
- We develop reference, online, and amortized certificate-synthesis procedures, with bounded prediction-initiated repair, to reduce the per-box cost of constructing certificates across parameterized strategic financial-market games and search regions.
- We integrate the certificate into a complete branch-and-bound solver and evaluate whether reduced spatial fragmentation yields end-to-end gains on controlled and strategic fire-sale instances after accounting for proposal, repair, rigorous verification, and fallback.

2. Related Work

2.1. Strategic Financial Equilibria and Complete Computation

Financial-network models describe institutions whose payments, liquidation choices, and balance-sheet adjustments affect one another. The Eisenberg–Noe model established a fixed-point foundation for systemic-risk analysis in liability networks [1]. Later work introduced endogenous price impact and strategic fire sales, where market values and liquidation incentives are determined jointly [2,4]. Strategic deleveraging under price-mediated contagion further leads to generalized-game formulations in which institutional constraints interact through common asset prices [3]. These models show that the same market can support states with different liquidation levels, price effects, and balance-sheet losses.

Algorithmic work has examined which financial equilibrium or intervention outcome can be computed under richer network rules. Seniority choices and minimal clearing change equilibrium existence and convergence properties [5]. Derivative contracts can make strong approximation and exact clearing difficult, while Tarski-based methods compute selected fixed points under broad monotone payment functions [6,7]. Other approaches optimize debt forgiveness, debt transfers, and claims trading to improve liquidity or institutional assets [8–10]. These studies provide rigorous tools for selected clearing states, equilibria, and interventions.

These methods leave the coverage problem unresolved. Computing a minimal, maximal, dynamically selected, or intervention-induced equilibrium does not certify that no

other approximate equilibrium exists elsewhere in a continuous strategy region. An additional stressed equilibrium can therefore remain undetected. EquiPartition addresses this limitation by maintaining a regional outer enclosure and requiring a rigorously verified deviation certificate for every discarded box. Regions that cannot be certified remain in the enclosure or are subdivided, so the search does not omit other equilibria without proof.

2.2. Complete Equilibrium Enclosure and Adaptive Certification

Equilibrium computation includes methods that seek one solution, a selected solution, or the complete solution set. Existence and representation have been studied for classical, noncompact, and multiplayer tensor games [46–48]. Distributed dynamics and integer-programming formulations seek or optimize equilibria under structural assumptions [11–13]. Global formulations extend computation to rational, mixed-integer, quasi-linear, and algebraically structured generalized Nash problems [17–20]. Complete methods use branch and prune or spatial branch and bound and discard a strategy region only after a valid exclusion test succeeds [14–16].

Other approaches improve how state dependence or unilateral gains are represented. Nikaido–Isoda functionals, variational-inequality cuts, and moment–sum-of-squares relaxations aggregate improvement information or tighten global equilibrium representations [21–24]. Adjustable robust optimization uses affine rules, finite adaptability, and K-adaptability to encode decisions that vary with uncertainty [25–27]. Recent work develops discrete, piecewise-affine, and exact adjustable policies [49–52]. Approximation guarantees and learning-guided K-adaptability further relate finite policy complexity to solution quality and search effort [29,39].

These tools still do not resolve responsibility switching in complete regional certification. Fixed-player tests fail when the profitable deviator changes inside a box, while hard selection or explicit partitioning represents the switching boundary with many subboxes. Adjustable policies are also not designed to certify a box-wide lower bound on maximum unilateral regret. EquiPartition addresses these limitations by combining legal unilateral deviations with continuous state-dependent responsibility weights and verifying the complete weighted certificate over the whole box. Responsibility can therefore move across institutions without first partitioning the region along the switching boundary.

2.3. Learning-Augmented Exact Solvers and Verified Proposals

Learning has been integrated into exact and hybrid solvers to avoid recomputing expensive decisions across related instances. Representative work learns branching policies, node comparisons, and cutting-plane selection [34,35,37]. Structural analyses study the learnability of branch-and-cut parameters, while enumeration oracles and separator configuration connect predictions to valid solver operations [36,38,53]. Learned diving and large-neighborhood search use expert solutions or solver trajectories to propose primal decisions and neighborhoods [54–56]. Modular recognition-and-tracking pipelines similarly divide a complex inference task into specialized learned stages [57]. Prediction has also been coupled to downstream optimization in secure trajectory-aware task offloading [58]. Attention-based sequence forecasting provides another example of structured predictors reused across repeated inference settings [59].

A related line combines learned proposals with exact or symbolic checks. Forward-prediction and reverse-design pipelines generate candidate decisions for repeated engineering-design tasks [60]. Visual–inertial SLAM similarly fuses complementary state cues within a structured estimation pipeline [61]. Neural bilevel optimization learns value-function surrogates that are reinserted into an optimization model, while learning-augmented algorithms connect prediction quality to worst-case protection [40,44].

Metaheuristic optimization of attention-residual generative models provides another coupling between learned representations and explicit search [62]. Neural model checking and exact verification of learned barrier certificates use learned objects as candidate proofs and retain an exact checker as the final source of validity [63,64]. Learned interpolants follow the same proposal-and-check principle [45], and solver-level studies emphasize end-to-end search behavior [41].

Residual squeeze–excitation networks provide a complementary example of channel-wise feature recalibration in learned computational pipelines [65].

Existing learning-augmented solvers do not directly solve repeated regional certificate construction in complete equilibrium search. Methods that learn branches, cuts, or candidate solutions do not represent institution-dependent deviation responsibility inside a region, while generic verified proposals do not address the spatial fragmentation caused by responsibility switching. EquiPartition predicts a player-adaptive certificate for each market–region pair, applies bounded repair when direct verification is inconclusive, and lets the rigorous verifier alone authorize deletion. If certification fails, the solver returns to complete subdivision, reducing neutral-start per-box synthesis cost without changing the correctness condition of the enclosure.

3. Methodology

EquiPartition consists of five components that correspond directly to Figure 1: (1) the certificate proposer, (2) the player-adaptive certificate, (3) the rigorous verifier, (4) bounded repair, and (5) the complete outer search. The top-level flow in Figure 1 shows their execution order, while panels A–C detail the player-adaptive certificate, rigorous verification, and complete-search mechanisms, respectively. Section 3.1 defines the enclosure objective and the fixed-player limitation processed by these components, and Sections 3.2–3.5 specify their construction and integration.

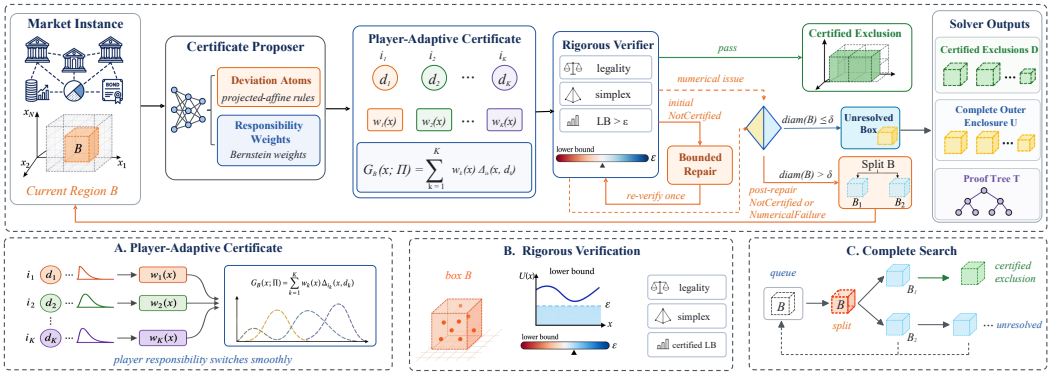


Figure 1. Overview of the five EquiPartition components and their execution flow. An inconclusive initial verification invokes at most one bounded repair; an unsuccessful post-repair verification returns the box to complete search. Panels (**A**–**C**) detail the certificate, verification, and search mechanisms.

3.1. Problem Formulation and Complete Enclosure Objective

Let $c \in \mathcal{C}$ denote a parameterized strategic financial-market instance and let $\mathcal{I} = \{1, \ldots, n\}$ index the financial institutions. Each institution acts as a player in the underlying Nash game. Institution i chooses a strategy x_i from a nonempty compact axis-aligned set $X_i \subset \mathbb{R}^{d_i}$, and the joint strategy space is

$$X = \prod_{i=1}^{n} X_i \subset \mathbb{R}^d, \qquad d = \sum_{i=1}^{n} d_i.$$

The instance may encode institutional exposures, market depth, exogenous shocks, and liquidation, funding, or inventory costs. Depending on the market model, x_i may

represent a liquidation rate, a deleveraging adjustment, a quote, or a liquidity-provision decision. We consider individually box-constrained games in which institutions interact through prices, payoffs, shared assets, or network states. For generalized Nash games with shared or decision-dependent constraints, regional exclusion additionally requires a rigorous certificate that each proposed unilateral deviation is feasible under the coupled constraints; feasibility and profitable deviation must then be verified jointly.

For a joint strategy $x = (x_i, x_{-i})$ and a feasible unilateral deviation $z_i \in X_i$, define

$$\begin{aligned} \Delta_i(c; x, z_i) &= u_i(z_i, x_{-i}; c) - u_i(x_i, x_{-i}; c), \\ R(c, x) &= \max_{i \in \mathcal{I}} \sup_{z_i \in X_i} \Delta_i(c; x, z_i), \\ \mathcal{E}_\varepsilon(c) &= \{x \in X : R(c, x) \le \varepsilon\}. \end{aligned} \tag{1}$$

Here Δ_i is the gain from a unilateral deviation, R is the maximum unilateral regret, and $\mathcal{E}_\varepsilon(c)$ is the ε-Nash equilibrium set.

The solver returns a finite collection of unresolved boxes $\mathcal{U}$ satisfying

$$\mathcal{E}_\varepsilon(c) \subseteq \bigcup_{B \in \mathcal{U}} B. \tag{2}$$

The union may contain nonequilibrium states, but it must not omit a target equilibrium. This one-sided coverage requirement distinguishes complete enclosure from methods that return only equilibrium candidates.

For an axis-aligned search box $B = \prod_{j=1}^{d} [\ell_j, r_j] \subseteq X$, pointwise nonequilibrium requires

$$\forall x \in B \; \exists i \in \mathcal{I} \; \exists z_i \in X_i : \Delta_i(c; x, z_i) > \varepsilon.$$

A fixed-player certificate instead requires the stronger condition

$$\exists i \in \mathcal{I} \; \forall x \in B \; \exists z_i \in X_i : \Delta_i(c; x, z_i) > \varepsilon. \tag{3}$$

The fixed-player condition can fail when the institution with the profitable deviation changes inside B, even if every state in the box is nonequilibrium. EquiPartition represents this state dependence inside the regional certificate.

3.2. Player-Adaptive Deviation Certificate

A player-adaptive certificate on B contains K deviation atoms. Atom $a_k = (i_k, d_k, w_k)$ associates an institution $i_k \in \mathcal{I}$ with a feasible deviation rule $d_k : B_{-i_k} \to X_{i_k}$ and a state-dependent responsibility weight $w_k : B \to [0, 1]$. With $\Pi = \{a_k\}_{k=1}^{K}$, the weights and certificate function satisfy

$$\begin{aligned} &w_k(x) \ge 0, \qquad \sum_{k=1}^{K} w_k(x) = 1 \quad \forall x \in B, \\ &G_B(x; \Pi) = \sum_{k=1}^{K} w_k(x) \Delta_{i_k}\big(c; x, d_k(x_{-i_k})\big). \end{aligned} \tag{4}$$

Every atom is a legal unilateral deviation, while different atoms may carry most of the proof responsibility in different parts of the same box. In a strategic financial market, an atom represents one institution and a feasible unilateral liquidation or balance-sheet adjustment under the current decisions of the other institutions.

For a legal certificate, let $\underline{G}_B(\Pi)$ be a rigorous lower bound satisfying $\underline{G}_B(\Pi) \leq \inf_{x \in B} G_B(x;\Pi)$. Each deviation gain in Equation (4) is bounded above by $R(c,x)$, and the nonnegative weights sum to one. Hence

$$G_B(x;\Pi) \leq R(c,x) \quad \forall x \in B, \qquad \underline{G}_B(\Pi) > \varepsilon \Longrightarrow B \cap \mathcal{E}_\varepsilon(c) = \varnothing. \tag{5}$$

Equation (5) is the exclusion rule used by the complete search. The proposer and repair procedure may construct candidate parameters, but only a rigorously certified lower bound above ε can remove a region.

The main representational benefit appears when deviation responsibility changes across a non-axis-aligned interface. An axis-aligned subbox is called a boundary box if it intersects that interface, and an axis-aligned representation resolves the interface at scale τ when every boundary box has diameter at most τ.

Theorem 1 (Axis-aligned representation size under player switching)**.** *Let a non-axis-aligned Lipschitz hypersurface $S \subset B$ with positive $(d-1)$-dimensional measure separate two regions of deviation responsibility. Suppose that no fixed-player certificate is valid on an axis-aligned subbox intersecting both sides of S, while a two-atom player-adaptive certificate verifies B with a uniform positive margin. Any finite axis-aligned fixed-player representation that resolves S at scale τ contains*

$$N_{\text{boundary}}(\tau) = \Omega\Big(\tau^{-(d-1)}\Big) \tag{6}$$

boundary boxes, whereas the player-adaptive representation uses two atoms on the original box B.

Proof. The boundary boxes cover S. A diameter-τ axis-aligned box covers at most a constant multiple of τ^{d-1} of a Lipschitz hypersurface, so covering a positive-measure part of S requires $\Omega(\tau^{-(d-1)})$ boxes. By assumption, a fixed-player certificate cannot remove a box crossing the two responsibility regions. The state-dependent weights of the two-atom certificate represent the same switch on B without external subdivision. The result concerns the axis-aligned representation of the switching interface, not every node or runtime contribution of the complete solver. □

3.3. Verifiable Certificate Parameterization

We realize the responsibility weights with tensor-product Bernstein polynomials. Map B to $[0,1]^d$ by

$$t_j = \frac{x_j - \ell_j}{r_j - \ell_j}, \qquad j = 1, \ldots, d.$$

For a multi-degree p and a multi-index $\alpha \preceq p$, define

$$\begin{aligned} b_{\alpha,p}(t) &= \prod_{j=1}^{d} \binom{p_j}{\alpha_j} t_j^{\alpha_j} (1-t_j)^{p_j - \alpha_j}, \\ w_k(t) &= \sum_{\alpha \preceq p} \rho_{k,\alpha} b_{\alpha,p}(t), \qquad \rho_{k,\alpha} \geq 0, \quad \sum_{k=1}^{K} \rho_{k,\alpha} = 1. \end{aligned} \tag{7}$$

The Bernstein basis is nonnegative and sums to one. The coefficient-wise simplex constraints therefore guarantee that the responsibility weights remain nonnegative and sum to one over the entire continuous box, without relying on checks at sampled points.

Each atom uses a projected-affine deviation rule,

$$d_k(x_{-i_k}) = \text{Proj}_{X_{i_k}}\big(A_k x_{-i_k} + b_k\big). \tag{8}$$

The projection keeps the deviation inside the corresponding institution's strategy set. If the affine image remains feasible throughout B_{-i_k}, the projection is inactive. If it crosses a strategy boundary, the verifier bounds the supported pieces separately; a candidate for which no rigorous bound can be obtained remains uncertified.

We write $\mathfrak{P}_K(B)$ for the finite-dimensional certificate family determined by the atom count, institution assignments, Bernstein coefficients, and bounded affine parameters. Reference synthesis, learned proposal, local repair, and rigorous verification all use this common family. The Bernstein and projected-affine forms implement the player-adaptive structure in Equation (4); they do not change the underlying definition of approximate Nash equilibrium.

3.4. Certificate Synthesis, Repair, and Verification

This subsection specifies the certificate interface for one search box. The selection, retention, and subdivision of boxes are handled by the complete search in Section 3.5.

Offline reference synthesis searches $\mathfrak{P}_K(B)$ for certificates with large rigorously certified margins and supplies the resulting parameters as training targets. At deployment, the proposer maps the market instance and current box to a candidate certificate,

$$\widehat{\Pi}_B = \pi_\theta(\phi(c, B)), \tag{9}$$

where $\phi(c, B)$ contains the market parameters, institution features, and normalized box bounds. The output specifies institution assignments, projected-affine deviation parameters, and Bernstein coefficients. Before verification, the Bernstein coefficients are projected onto the simplices in Equation (7).

The verifier $\mathcal{V}(c, B, \Pi, \varepsilon)$ evaluates the complete certificate over B and returns one of three statuses. `CertifiedExclusion` means that the rigorous lower bound exceeds ε; `NotCertified` means that the current bound is inconclusive; and `NumericalFailure` records a failure of the supported rigorous arithmetic. Only `CertifiedExclusion` authorizes removal. The other two statuses leave the enclosure unchanged.

On a polynomial verification cell $C \subseteq B$, write the certificate in Bernstein form as

$$G_B(x;\Pi)\big|_C = \sum_{\beta \preceq q} \gamma_{\beta,C} b_{\beta,q}(t_C(x)), \qquad \underline{G}_C(\Pi) = \min_{\beta \preceq q} \gamma_{\beta,C} \le \inf_{x \in C} G_B(x;\Pi). \tag{10}$$

On each supported polynomial piece of a verification cell C, the verifier writes the kth responsibility weight and deviation gain as Bernstein expansions with coefficient arrays $\rho_{k,\alpha,C}$ and $\eta_{k,\nu,C}$ and multi-degrees p and $r_{k,C}$. Their product is formed in the Bernstein basis through the coefficient convolution

$$\gamma^{(k)}_{\beta,C} = \sum_{\substack{\alpha+\nu=\beta \\ \alpha \preceq p,\, \nu \preceq r_{k,C}}} \omega^{(k,C)}_{\alpha,\nu,\beta} \rho_{k,\alpha,C} \eta_{k,\nu,C}$$

$$\omega^{(k,C)}_{\alpha,\nu,\beta} = \frac{\binom{p}{\alpha}\binom{r_{k,C}}{\nu}}{\binom{p+r_{k,C}}{\beta}}, \qquad \binom{p}{\alpha} = \prod_{j=1}^{d} \binom{p_j}{\alpha_j}.$$

The atom products are degree-elevated to a common multi-degree and summed to obtain the coefficients $\gamma_{\beta,C}$ in Equation (10). Projection boundaries and the positive-part shortfall term are handled as separate supported polynomial pieces, and exclusion requires every piece to have a lower bound above ε.

Cell mapping, Bernstein conversion, coefficient products, degree elevation, summation, and threshold comparison use 128-bit MPFR arithmetic with directed outward rounding. For the internal cells $\mathcal{C}(B)$ visited in one verifier call, the box-level bound is

$$\underline{G}_B(\Pi) = \min_{C \in \mathcal{C}(B)} \underline{G}_C(\Pi) \leq \inf_{x \in B} G_B(x; \Pi).$$

If the initial cell is inconclusive, the fixed internal subdivision rule is applied within the common budget. Budget exhaustion returns `NotCertified`; unsupported expressions or failed rigorous arithmetic return `NumericalFailure`. Neither status authorizes exclusion.

When the initial candidate returns `NotCertified`, it initializes a repair procedure that remains inside $\mathfrak{P}_K(B)$ and performs at most J local updates to increase the certified margin. The repaired certificate is then verified once more. If the second verification is unsuccessful, or if rigorous arithmetic fails, the box returns directly to the conservative outer search. Thus, one box uses one proposal, at most one bounded repair stage, and at most two verifier calls before fallback.

3.5. Complete Enclosure Algorithm

Starting from the initial strategy box X, the complete solver maintains an active queue $\mathcal{Q}$, an unresolved set $\mathcal{U}$, a certified-exclusion set $\mathcal{D}$, and a proof tree $\mathcal{T}$. All certificate variants use the same maximum-diameter box selection, lexicographic tie breaking, and longest-edge bisection. The equilibrium tolerance ε defines the target equilibrium set, whereas the terminal resolution δ controls the size of boxes that remain unresolved.

Direct verification is attempted before repair, so repair is used only when the initial candidate is inconclusive. A numerical failure or an unsuccessful repair follows the same conservative retain-or-split path. The proposer changes the cost of finding a certificate, while rigorous verification remains the only condition for regional exclusion.

Theorem 2 (Completeness and finite termination)**.** *Suppose that every box added to $\mathcal{D}$ satisfies Equation (5), each split replaces its parent with children having the same union, and every terminal uncertified box is retained in $\mathcal{U}$. Then Algorithm 1 satisfies Equation (2), independently of the accuracy of π_θ. If X is bounded, $\delta > 0$, and longest-edge bisection is used, the algorithm terminates after processing finitely many boxes.*

Proof. At every iteration, each target equilibrium lies in an active box or in a box already retained in $\mathcal{U}$. The invariant holds initially because X contains $\mathcal{E}_\varepsilon(c)$. Equation (5) shows that removing a certified box preserves the invariant, splitting preserves the parent union, and a terminal uncertified box is transferred to $\mathcal{U}$. When the queue is empty, every target equilibrium therefore lies in the unresolved union. For bounded X, repeated longest-edge bisection reduces every root-to-leaf box below diameter δ after finitely many splits, which gives finite tree depth and finite termination. □

If a run is interrupted by the wall-clock limit, the box under evaluation is returned to the active set before output is recorded. The terminal unresolved boxes together with all active boxes then remain a conservative outer enclosure; only completion to resolution δ is missing.

Algorithm 1 EquiPartition

Input: *financial-market game c, strategy box X, tolerances ε and δ, proposal map π_θ, verifier $\mathcal{V}$, repair budget J*
Output: *unresolved boxes $\mathcal{U}$, certified exclusions $\mathcal{D}$, and proof tree $\mathcal{T}$*

$\mathcal{Q} \leftarrow \{X\}, \mathcal{U} \leftarrow \varnothing, \mathcal{D} \leftarrow \varnothing, \mathcal{T} \leftarrow$ InitializeTree(X)
while $\mathcal{Q} \neq \varnothing$ **do**
  $B \leftarrow$ PopFixed$(\mathcal{Q})$; $\Pi_B \leftarrow \pi_\theta(\phi(c, B))$
  $(s, \underline{g}, \xi) \leftarrow \mathcal{V}(c, B, \Pi_B, \varepsilon)$
  if $s =$ `NotCertified` **then**
    $\Pi_B \leftarrow$ Repair(c, B, Π_B, J); $(s, \underline{g}, \xi) \leftarrow \mathcal{V}(c, B, \Pi_B, \varepsilon)$
  end if
  if $s =$ `CertifiedExclusion` **then**
    $\mathcal{D} \leftarrow \mathcal{D} \cup \{B\}$; RecordCertificate$(\mathcal{T}, B, \Pi_B, \underline{g}, \xi)$
  else if $\operatorname{diam}(B) \leq \delta$ **then**
    $\mathcal{U} \leftarrow \mathcal{U} \cup \{B\}$; RecordUnresolved$(\mathcal{T}, B)$
  else
    $(B_1, B_2) \leftarrow$ LongestEdgeSplit(B); $\mathcal{Q} \leftarrow \mathcal{Q} \cup \{B_1, B_2\}$
    RecordSplit$(\mathcal{T}, B, B_1, B_2)$
  end if
end while
return $(\mathcal{U}, \mathcal{D}, \mathcal{T})$

4. Experimental Setup

4.1. Strategic Financial Models and Dataset Construction

We study a normalized one-period fire-sale market in which institutions interact through common asset holdings and endogenous price impact. A market contains n institutions and A assets. Institution i holds $h_{ia} \geq 0$ units of asset a and chooses a liquidation rate $x_i \in [0, 1]$. Aggregate sales are $q_a(x) = \sum_{j=1}^{n} h_{ja} x_j$, and the post-liquidation price is $p_a(q) = p_a^0(1 - \alpha_a q - \beta_a q^2)$. Let $\bar{p}_a(q)$ denote the average execution price from zero to q. The payoff is

$$u_i(x; c) = \sum_{a=1}^{A} h_{ia}[x_i \bar{p}_a(q_a(x)) + (1 - x_i) p_a(q_a(x))] - \frac{\kappa_i}{2} x_i^2 - \rho_i \left[s_i - \sum_{a=1}^{A} h_{ia} x_i \bar{p}_a(q_a(x)) \right]_+^2. \tag{11}$$

The first term combines liquidation proceeds with the marked value of the remaining portfolio. The coefficient κ_i controls direct liquidation cost, while s_i and ρ_i determine the funding requirement and the cost of any remaining shortfall. Liquidation by one institution changes the execution and residual values faced by other institutions holding the same assets. We restrict the impact parameters so that prices remain positive throughout the strategy domain. The resulting game has box-constrained actions and piecewise-polynomial payoffs, consistent with Section 3 [2,4].

The controlled dataset contains 240 markets with two or three institutions and one or two shared assets. Holdings, liquidation costs, funding requirements, and price-impact coefficients are generated heterogeneously. We retain instances for which the high-precision reference solver identifies at least two separated equilibrium components and a nonequilibrium region containing a non-axis-aligned switch in the institution providing the largest certified deviation. This produces diagnostic cases for the proposed certificate mechanism. The dataset is divided into 160 training, 40 validation, and 40 test markets.

We additionally evaluate 40 unfiltered controlled markets drawn directly from the same generator without conditioning on the number of equilibrium components or the geometry of deviation responsibility. The trained proposal model remains fixed. We also evaluate 119 controlled and 337 multi-institution single-dominant-deviator regions, defined by a recorded reference responsibility pattern that assigns one institution as dominant throughout the region.

The multi-institution dataset contains 720 markets with three to six institutions and two to four assets, divided into 480 training, 120 validation, and 120 independent and identically distributed (IID) test markets. Each portfolio combines common and institution-specific holdings, producing overlap levels between 0.2 and 0.9. After normalization, $\alpha_a \in [0.05, 0.20]$, $\beta_a \in [0, 0.05]$, $\kappa_i \in [0.01, 0.08]$, and the funding requirement is between 0.10 and 0.40 of the institution's initial marked portfolio value.

We construct paired stress variants of the held-out markets, with 40 controlled and 120 multi-institution markets per shift. Let $V_i^0 = \sum_a h_{ia} p_a^0$. For the high-leverage shift, s_i / V_i^0 is resampled from $[0.40, 0.55]$, while holdings and price-impact coefficients retain their paired IID values. For the low-liquidity shift, only α_a is resampled from $[0.20, 0.35]$.

The hidden-crisis shift combines $s_i / V_i^0 \in [0.32, 0.50]$, $\alpha_a \in [0.16, 0.30]$, and portfolio overlap in $[0.65, 0.95]$, while retaining the paired liquidation-cost and quadratic-impact parameters. All shifts preserve the action space, payoff equation, equilibrium tolerance, and verifier protocol. None of the shifted markets is used to update the proposal model.

For financial interpretation, each retained equilibrium state is assigned the stress score

$$S_{\text{stress}}(x) = 0.32L(x) + 0.32P(x) + 0.28E(x) + 0.08H(x), \tag{12}$$

where $L(x)$ is aggregate liquidation relative to initial holdings, $P(x)$ is the decline in the minimum asset price, and $E(x)$ is the system equity-loss share. The residual funding-pressure ratio is $r_i(x) = [s_i - \sum_a h_{ia} x_i \bar{p}_a(q_a(x))]_+ / V_i^0$, and $H(x) = n^{-1} \sum_i \mathbf{1}\{r_i(x) \geq 0.10\}$.

States with $S_{\text{stress}} < 0.30$ are labeled Mild Liquidation, those with $0.30 \leq S_{\text{stress}} < 0.50$ are labeled Partial Deleveraging, and those with $S_{\text{stress}} \geq 0.50$ are labeled Fire-Sale. Fire-Sale states define the crisis-equilibrium class used in the coverage analysis. These labels are assigned after equilibrium computation and do not affect certificate proposal, repair, verification, or regional exclusion.

All dataset splits are made at the market level. A stricter spatial branch-and-bound run provides the reference objects using $\varepsilon_{\text{ref}} = 2.5 \times 10^{-4}$ and $\delta_{\text{ref}} = 2^{-11}$ [15,16]. The reference run records a high-precision equilibrium enclosure and region-level certificate targets sampled across search depths, verifier outcomes, and responsibility patterns. These targets supervise proposal learning and support evaluation; every test-time exclusion remains subject to the online rigorous verifier.

An independent low-dimensional validation pipeline processes the 40 controlled test markets using active-set enumeration and interval root isolation. It operates independently of the learned proposal, player-adaptive certificates, and the reference search tree. Isolated candidates are matched to recorded reference components after satisfying the reference-tolerance regret test.

4.2. Baselines

We compare EquiPartition with complete-search, certificate-structure, and amortized-synthesis baselines. *Nash-B&B* is the reference complete spatial branch-and-bound solver without a learned player-adaptive certificate [15,16]. *Fixed-Institution* uses one institution and one constant deviation for the entire box. *Adaptive-Action* fixes the institution but optimizes a projected-affine deviation. *Hard-Partition* represents responsibility switching through outer subdivision and applies a Fixed-Institution certificate to each resulting subbox.

The remaining structural baselines use the same institution–deviation atoms but assign responsibility differently. Define $g_k(x) = \Delta_{i_k}(c; x, d_k(x_{-i_k}))$ and $\mathcal{S}_K = \{\lambda \in \mathbb{R}^K_+ : \sum_k \lambda_k = 1\}$. Their regional expressions are

$$\begin{aligned} G_B^{\mathrm{NI}}(x) &= \frac{1}{K}\sum_{k=1}^{K} g_k(x), \\ G_B^{\mathrm{AR}}(x) &= \sum_{k=1}^{K} \lambda_k g_k(x), \qquad \lambda \in \mathcal{S}_K, \\ s(x) &= \operatorname*{argmax}_{1\le k\le K}\{a_k^\top t(x) + b_k\}, \qquad G_B^{\mathrm{KA}}(x) = g_{s(x)}(x). \end{aligned} \tag{13}$$

Regional NI uses the equal-weight Nikaido–Isoda-type aggregate G_B^{NI} [21,22]. *Affine Recourse* optimizes a constant box-level mixture λ together with the projected-affine deviation rules [25]. *K-Adaptability* uses the same number of deviation rules and selects one through the affine score $s(x)$ [26,29,39].

For K-Adaptability, interval score bounds identify internal regions with a unique selected rule. Unresolved ties trigger internal subdivision and return `NotCertified` when the verifier budget is exhausted. EquiPartition instead uses continuous Bernstein responsibility weights. Because each $g_k(x) \le R(c, x)$, the convex aggregates and selected gains in Equation (13) use the same sound exclusion threshold as EquiPartition.

For a fixed atom count, the multi-rule methods use the same number and dimensions of projected-affine deviation rules. They also share the outer branching policy, per-box update budget, rigorous arithmetic, and internal subdivision budget specified in Section 4.4. HardPartition subdivisions are counted as outer proof-tree nodes, whereas selector refinements in K-Adaptability are counted as internal verifier cells.

Four EquiPartition variants separate certificate quality from synthesis cost. *EquiPartition-Online* optimizes a certificate independently at every visited box from a neutral initialization. *EquiPartition-Predict* directly verifies the amortized proposal. *EquiPartition-Predict+Repair* applies bounded repair after an inconclusive direct verification and is the deployed method. *EquiPartition-Oracle* uses multiple restarts and a larger search budget within the same certificate family to estimate attainable certificate quality.

Parent-WarmStart+Repair is a non-learning amortization baseline. The root box uses the neutral initialization. After a split, the final parent certificate is inherited by each child, with its Bernstein responsibility representation restricted to the child box before direct verification; an inconclusive candidate then receives the same bounded repair and rigorous verification as Predict+Repair.

4.3. Evaluation Metrics

Recorded-Reference Coverage is the fraction of recorded high-precision reference-equilibrium boxes intersected by the unresolved enclosure. *Reference-Set False Exclusion* counts certified-exclusion boxes that intersect the recorded reference set. These metrics audit the recorded reference objects. Formal complete-enclosure coverage is established by Equation (5) and Theorem 2. The stricter reference tolerance $\varepsilon_{\mathrm{ref}} < \varepsilon$ is used to generate the high-precision audit objects.

For a capped run, the reference audit uses the conservative partial enclosure formed by terminal unresolved boxes and all active boxes, including an interrupted current box returned to the queue. The *Completed* metric separately records whether the active queue was exhausted.

Three certification rates use different denominators. Full-Region Certification Rate is the fraction of selected switching regions certified without outer subdivision. Outer-Box

Exclusion Rate is the fraction of processed outer boxes removed before fallback. Candidate Certification Rate is the fraction of Oracle-certifiable candidates certified directly or after repair. Verifier-status frequencies are pooled over verifier calls and reported separately for the direct-proposal and post-repair stages; the three mutually exclusive outcomes are `CertifiedExclusion`, `NotCertified`, and `NumericalFailure`.

Unresolved Volume is the retained volume normalized by the initial strategy-box volume. *Proof-Tree Nodes* measure outer spatial fragmentation, while *Verifier Calls* measure the number of certificate-verification attempts. Proof-Tree Nodes count only boxes created by the outer search. Internal verifier cells are excluded from this structural metric for every method, while all arithmetic performed on those cells is included in verification time and End-to-End Runtime. In the default-cap accounting, mean and 95th-percentile internal-cell counts and cap-exhaustion frequency are pooled over verifier calls, whereas full-region certification, certified volume, and verification time retain the market-level aggregation.

End-to-End Runtime includes proposal, repair, rigorous verification, queue management, and spatial subdivision. Certificate-level diagnostics include direct and final candidate certification rates, certified volume, certified margin, and synthesis time. Dedicated profiling runs further separate proposal, refinement, verification, region-management, and residual overhead. Every selected profiling run contributes one complete component vector.

4.4. Experimental Protocol

All complete solvers use the same normalized strategy box, equilibrium tolerance $\varepsilon = 10^{-3}$, terminal resolution $\delta = 2^{-8}$, maximum-diameter box selection, lexicographic tie breaking, and longest-edge bisection. The same outward-rounded Bernstein–interval backend is used whenever a baseline admits the corresponding regional expression. All verifier-compatible methods use the same cap of 64 total internal cells per verifier call, including the root cell, and the same wall-clock limit. Cap exhaustion returns `NotCertified`.

Online synthesis, bounded repair, and the optimization-based structural baselines receive at most 20 projected updates per box. Numerical failures follow the conservative retain-or-split path and never produce exclusions. Controlled instances use two certificate atoms, while multi-institution instances use at most four. Responsibility weights use degree-two Bernstein polynomials, and deviation rules use the projected-affine family in Section 3.3.

Predict+Repair performs at most 20 certified-margin updates and stops when the verifier accepts the box. Online synthesis uses the same update rule and budget from a neutral initialization. Oracle synthesis uses eight restarts and 200 updates, corresponding to ten times the Online bound-evaluation budget, and is excluded from wall-clock rankings.

Parent-WarmStart+Repair uses the same 20-update repair limit and verifier settings but replaces the learned proposal with the final parent certificate restricted to the child box; the root box starts from the same neutral initialization as Online synthesis.

The three sensitivity studies reuse the same 24 controlled and 48 multi-institution market identities but use distinct aggregation units. The K–p sweep evaluates $K \in \{1, 2, 3, 4\}$ and Bernstein degree $p \in \{1, 2, 3\}$ with the same projected-affine deviation family. Online, Predict, and Predict+Repair each use five seed-indexed configuration runs, while Oracle contributes one reference run. The reported verification time is the configuration-sweep mean per region. These results are used only in Section 6.3 and are not mixed with the main test-set aggregates.

Repair-budget sensitivity reruns complete search separately for $J \in \{0, 5, 10, 20, 40\}$ at $K = 2$ and $p = 2$; its certificate-volume, tree, runtime, and completion quantities are on-policy summaries of the trajectory induced by each budget. Internal-cell-cap sensitivity holds the $K = 2$, $p = 2$ certificate inputs fixed within that sweep, varies the total internal-

cell cap over $\{1, 8, 16, 32, 64, 128\}$, and reports the verifier timer collected by the dedicated cap sweep. A cap of one evaluates only the root cell, and 64 is the default verifier protocol. Thus, the three sweeps share market identities but do not use an interchangeable timing or trajectory-level estimator.

Separate proposal models are trained for the controlled and multi-institution groups. Each uses a permutation-invariant three-layer institution encoder with 128 hidden units per layer, mean–maximum pooling, and two 256-unit global layers. The inputs summarize institution-level portfolio, funding, cost, overlap, and current-box information together with global market-size, price-impact, overlap, and search-depth descriptors. Separate heads predict institution assignments and continuous certificate parameters.

We sample 64 regions from each training-market trajectory, giving 10,240 controlled and 30,720 multi-institution training pairs. The fixed per-market quota gives every training market equal representation and prevents longer reference-search trajectories from dominating the proposal dataset. Within each trajectory, the sampled regions cover search depths, verifier outcomes, and responsibility patterns recorded by the reference run. Predicted and reference atoms are aligned by minimum-cost matching over institution assignments and continuous parameters. The training objective is

$$\mathcal{L} = \mathcal{L}_{\text{assign}} + \mathcal{L}_{\text{affine}} + 0.5\mathcal{L}_{\text{Bernstein}},$$

where assignment uses cross-entropy and continuous parameters use smooth-ℓ_1 loss.

The proposal models are trained with AdamW using a learning rate of 3×10^{-4}, weight decay 10^{-5}, batch size 256, and at most 120 epochs, with early stopping based on validation certified margin. Five independent proposal-model seeds are trained for Predict and Predict+Repair. During repair, institution assignments remain fixed, and continuous certificate parameters are updated by projected Adam for at most 20 steps.

Solver and profiling experiments were run on Ubuntu 22.04 with an AMD EPYC 7543P CPU, 256 GB of RAM, and an NVIDIA RTX A6000 GPU with 48 GB of memory. The GPU was used for proposal-model training. All reported online proposal, repair, verification, and branch-and-bound times were measured using CPU double-precision inference. Each market used one process pinned to one physical core, with multithreading and outer-node parallelism disabled.

The implementation used Python 3.11, PyTorch 2.4, and a C++17 Bernstein–interval verifier compiled with GCC 11.4 at `-O3`. Outward-rounded coefficient calculations and threshold comparisons used MPFR 4.2 at 128-bit precision. Runtime includes proposal inference, local repair, rigorous verification, queue operations, spatial subdivision, and proof-record serialization. It excludes offline reference generation, proposal-model training, and process startup. All methods use a two-hour cap.

The evaluation run count is method-specific. EquiPartition-Online, EquiPartition-Predict, and EquiPartition-Predict+Repair each use five seed-indexed runs per market; Predict and Predict+Repair use the corresponding proposal-model seed, whereas Online uses the same neutral-start protocol without a learned proposal. Nash-B&B, Fixed-Institution, Adaptive-Action, Hard-Partition, Regional NI, Affine Recourse, *K*-Adaptability, and Parent-WarmStart+Repair use one run per market. EquiPartition-Oracle contributes one reference run per market or configuration, with its eight restarts performed inside that run. The main solver-level and stress-shift results represent Online, Predict, and Predict+Repair by the actual seed run with the median capped runtime; each single-run method is represented by its sole run. A market is counted as completed when the applicable representative run finishes before the cap.

Runtime decomposition uses seed-level profiling availability. A market enters the runtime-decomposition analysis when at least one seed completes both the solver and every profiling timer. Among the eligible seeds, the actual run with the median profiling total

provides all component values. Solver timeouts, profiling timeouts, and incomplete profiles are censored. Consequently, the number of profiled markets can exceed the completion count based on the five-seed median run.

5. Experimental Results

We evaluate complete financial-equilibrium enclosure at the solver level, isolate institutional switching at the certificate level, examine the principal certificate components, and quantify the computational value of amortized certificate synthesis. The three five-seed variants—Online, Predict, and Predict+Repair—follow the market-level median-run aggregation protocol in Section 4; all listed single-run methods are evaluated once per market under the same verifier, branching rule, and time limit.

5.1. Complete Financial-Equilibrium Enclosure

Table 1 summarizes the complete-enclosure results on the controlled and multi-institution test markets.

Table 1. Complete approximate-equilibrium enclosure results. Ref. cov./FE denotes recorded-reference coverage (%)/false exclusions. Runtime is capped at 7200 s; timeout audits use the conservative partial enclosure. Online, Predict, and Predict+Repair report the seed run with median capped runtime across five runs; other methods use one run per market.

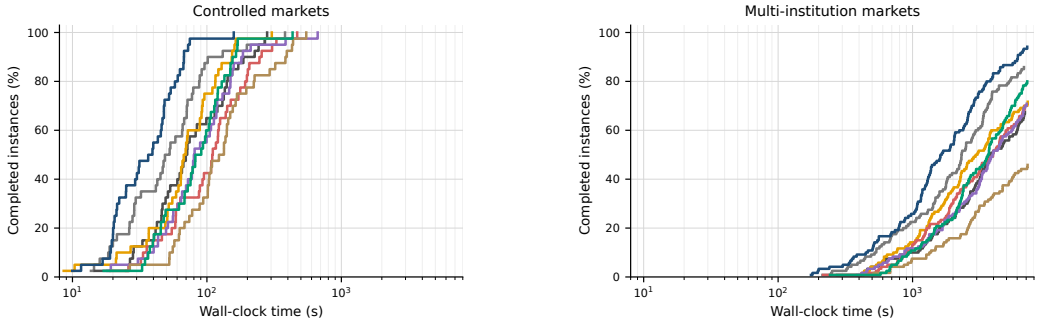
Controlled Financial Markets						
Method	**Ref. Cov./FE**	**Proof-Tree Nodes**	**Verifier Calls**	**Unres. Vol. (%)**	**Runtime (s)**	**Completed**
Nash-B&B	100/0	7780 ± 5560	5226 ± 3327	2.96 ± 1.34	69.3 ± 70.7	40/40
Fixed-Institution	100/0	7883 ± 4240	7944 ± 4542	2.29 ± 1.02	95.8 ± 69.4	40/40
Adaptive-Action	100/0	5782 ± 3192	5903 ± 3135	1.87 ± 0.83	81.5 ± 54.1	40/40
Hard-Partition	100/0	5945 ± 3389	6795 ± 3869	1.33 ± 0.51	129.5 ± 93.6	40/40
Regional NI	100/0	3976 ± 2143	4018 ± 2365	1.53 ± 0.67	113.9 ± 111.5	40/40
Affine Recourse	100/0	4523 ± 2294	4600 ± 2316	1.27 ± 0.68	103.4 ± 69.3	40/40
K-Adaptability	100/0	3966 ± 2614	5462 ± 4281	1.16 ± 0.56	161.5 ± 125.5	40/40
EquiPartition-Online	100/0	1260 ± 532	1556 ± 681	0.73 ± 0.30	97.7 ± 68.8	40/40
Parent-WarmStart+Repair	100/0	1467 ± 794	1837 ± 976	0.84 ± 0.34	69.8 ± 51.9	40/40
EquiPartition-Predict	100/0	2038 ± 881	2078 ± 867	1.06 ± 0.39	46.1 ± 35.5	40/40
EquiPartition-Predict+Repair	100/0	1297 ± 533	1699 ± 728	0.72 ± 0.26	41.3 ± 26.5	40/40
EquiPartition-Oracle	100/0	715 ± 377	969 ± 569	0.44 ± 0.24	297.9 ± 421.9	40/40
Multi-Institution Financial Markets						
Method	**Ref. Cov./FE**	**Proof-Tree Nodes**	**Verifier Calls**	**Unres. Vol. (%)**	**Runtime (s)**	**Completed**
Nash-B&B	100/0	53,322 ± 47,283	47,874 ± 48,053	15.05 ± 5.82	2989.0 ± 2265.2	103/120
Fixed-Institution	100/0	53,693 ± 39,188	84,352 ± 69,129	12.64 ± 4.63	4334.8 ± 2509.1	81/120
Adaptive-Action	100/0	35,533 ± 28,527	55,498 ± 50,189	10.44 ± 4.33	3804.6 ± 2595.5	86/120
Hard-Partition	100/0	29,731 ± 23,600	50,883 ± 47,603	8.39 ± 3.18	4133.5 ± 2542.3	84/120
Regional NI	100/0	27,114 ± 20,972	33,548 ± 28,299	8.84 ± 3.88	4273.2 ± 2473.1	85/120
Affine Recourse	100/0	26,996 ± 23,602	36,961 ± 37,663	8.16 ± 4.18	3880.2 ± 2511.2	90/120
K-Adaptability	100/0	32,742 ± 31,724	74,533 ± 80,730	7.37 ± 3.92	5284.4 ± 2420.2	55/120
EquiPartition-Online	100/0	8663 ± 6765	16,053 ± 14,734	4.56 ± 2.23	3924.3 ± 2315.7	96/120
Parent-WarmStart+Repair	100/0	11,892 ± 8721	17,366 ± 13,988	5.47 ± 2.35	3278.6 ± 2268.4	104/120
EquiPartition-Predict	100/0	15,763 ± 11,132	20,765 ± 15,928	6.70 ± 2.86	2882.5 ± 2274.3	106/120
EquiPartition-Predict+Repair	100/0	8585 ± 5948	14,235 ± 10,899	4.67 ± 1.99	2385.8 ± 2010.2	113/120
EquiPartition-Oracle	100/0	3731 ± 3106	7834 ± 7426	2.83 ± 1.49	4885.7 ± 2341.7	76/120

Across both market groups, recorded-reference coverage was 100%, and no accepted exclusion intersected the recorded reference set. Solver-level differences appear in enclosure tightness, proof-tree size, verification effort, completion rate, and capped wall-clock time.

On the controlled markets, EquiPartition-Predict+Repair used 1297 proof-tree nodes on average, left 0.72% of the initial strategy volume unresolved, and required 41.3 s. Relative to Nash-B&B, these values correspond to 83.3% fewer nodes, 75.8% lower unresolved volume, and 40.5% lower runtime. The strongest non-EquiPartition certificate baselines produced smaller trees than Nash-B&B, but the best of them still required 3,966 nodes and retained 1.16% unresolved volume. EquiPartition-Online obtained a slightly smaller tree than the deployed variant, whereas its per-box certificate optimization increased runtime to 97.7 s. Parent-WarmStart+Repair used 1467 nodes and 69.8 s, capturing part of the amortization benefit without a learned proposal but remaining above Predict+Repair in both tree size and runtime. The Oracle reference produced the tightest enclosure and the smallest tree, but its larger search budget raised mean runtime to 297.9 s.

The same pattern persists on the multi-institution markets. EquiPartition-Predict+Repair reduced the mean tree from 53,322 nodes for Nash-B&B to 8585 nodes and reduced unresolved volume from 15.05% to 4.67%. Its capped mean runtime was 2385.8 s, compared with 2989.0 s for Nash-B&B, and it completed 113 of 120 markets rather than 103. Against the best non-EquiPartition value for each structural metric, it used 68.2% fewer nodes than Affine Recourse and retained 36.7% less unresolved volume than K-Adaptability. Parent-WarmStart+Repair used 11,892 nodes, retained 5.47% unresolved volume, required 3278.6 s, and completed 104 of 120 markets. It improved on neutral-start Online in runtime and completion but did not match the learned Predict+Repair variant in tree size, runtime, or completion. The direct-prediction variant remained faster than several traditional baselines, but its resulting searches retained more nodes and unresolved volume than Predict+Repair. The Oracle reference again attained the smallest tree and tightest enclosure, while its higher synthesis cost and larger timeout rate prevent it from serving as an online method.

Figure 2 shows the solver-level completion profiles. On the multi-institution markets, EquiPartition-Predict+Repair reaches 94.2% completion, compared with 85.8% for Nash-B&B, 80.0% for EquiPartition-Online, and 45.8% for K-Adaptability. The controlled profiles reach full completion for all evaluated methods, while their horizontal separation reflects the different time required to construct equally complete enclosures.



(**a**) (**b**)

Figure 2. Completion profiles on (**a**) controlled and (**b**) multi-institution markets. Each curve gives the fraction of complete enclosures obtained within the wall-clock budget.

Figure 3 confirms that the aggregate improvement does not arise from uniform per-instance dominance. Among the 145 markets completed by both methods, EquiPartition was faster on 85 and slower on 60. Eight additional multi-institution markets were completed only by EquiPartition, three were completed only by the best baseline, and four reached the time limit on both sides. The retained reversals and censored cases provide a more complete view of solver behavior than a single aggregate runtime.

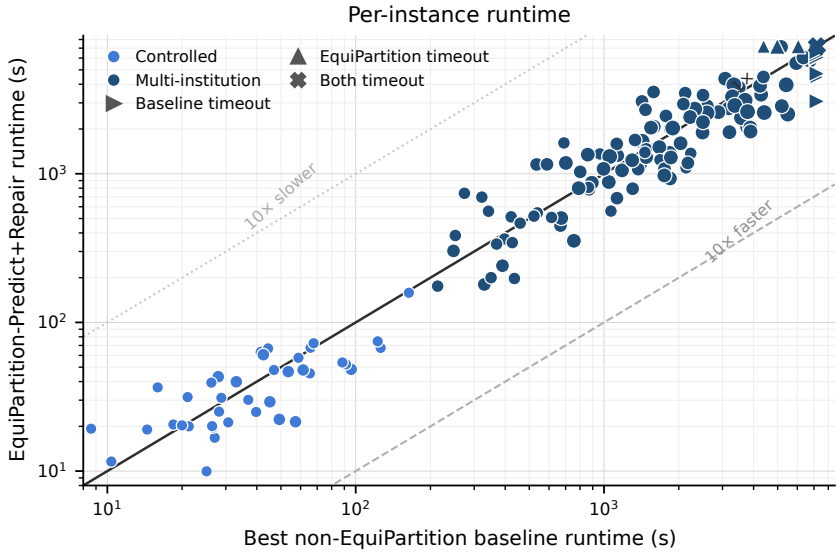


Figure 3. Per-instance runtime of EquiPartition-Predict+Repair versus the fastest non-EquiPartition baseline. Points below the diagonal favor EquiPartition; triangles denote one-sided timeouts, crosses denote two-sided timeouts, dashed lines mark tenfold differences, and marker size denotes institution count.

5.2. Institutional Switching and Regional Certification

Table 2 shows that player-adaptive certification covers switching regions with substantially fewer external partitions. On the controlled markets, EquiPartition-Online reached a full-region certification rate of 90.05%, certified 98.03% of the target volume, and required 1.37 partitions on average. The strongest competing full-region rate was 69.01% under Regional NI, while *K*-Adaptability attained the largest competing certified volume at 83.94%. EquiPartition improved these values by 21.04 and 14.09 percentage points, respectively, and used 73.9% fewer partitions than Affine Recourse, the lowest-partition competing method.

On the multi-institution markets, EquiPartition-Online obtained an 80.93% full-region certification rate and 94.29% certified volume with 1.78 partitions. Hard-Partition provided the strongest competing certification rate and volume at 58.77% and 76.61%, while Affine Recourse required 6.25 partitions. EquiPartition increased full-region certification by 22.16 percentage points, increased certified volume by 17.68 percentage points, and reduced the partition count by 71.6%. Its mean verification time was 0.157 s on controlled regions and 0.331 s on multi-institution regions, while the corresponding verifier-call counts of 1.86 and 2.00 were the lowest among all compared methods. The richer regional proof therefore converts a moderately higher per-certificate evaluation cost into a markedly smaller spatial representation.

Table 3 reports the selection-control results. On the unfiltered controlled markets, Predict+Repair completed all 40 instances, retained 0.74% unresolved volume, and required 31.5 s on average. Relative to Nash-B&B, it reduced proof-tree nodes by 66.0%, unresolved volume by 59.1%, and capped runtime by 16.9%. Online synthesis produced the smallest tree and unresolved volume in this control, but its 66.8 s runtime exceeded the amortized variant.

The single-dominant-deviator control shows a different method ordering from the switching regions in Table 2. On controlled regions, Adaptive-Action obtained the highest full-region certification rate and certified volume at 95.1% and 97.4%, while EquiPartition-Online reached 92.4% and 96.5%. On multi-institution regions, *K*-Adaptability reached 86.1% full certification and 95.7% certified volume, compared with 83.0% and 91.5% for EquiPartition-Online. The large advantage in Table 2 is therefore concentrated in regions where deviation responsibility switches rather than in regions governed by one dominant deviator.

Table 2. Regional certification on preselected institution-switching regions. Full cert. is the market-level full-region certification rate; other metrics are averaged within markets and then across markets.

Controlled Target Regions					
Method	**Full Cert. (%)**	**Cert. Vol. (%)**	**Partitions**	**Verif. Time (s)**	**Calls**
Fixed-Institution	36.45 ± 14.16	50.48 ± 12.66	8.10 ± 1.33	0.109 ± 0.017	4.44 ± 0.72
Adaptive-Action	53.11 ± 13.26	72.50 ± 8.70	7.05 ± 1.07	0.137 ± 0.020	4.09 ± 0.60
Hard-Partition	68.59 ± 10.84	83.73 ± 6.40	12.30 ± 2.23	0.269 ± 0.055	5.75 ± 0.97
Regional NI	69.01 ± 13.93	81.76 ± 9.01	5.63 ± 0.93	0.218 ± 0.040	3.42 ± 0.49
Affine Recourse	60.04 ± 15.26	76.47 ± 10.20	5.24 ± 0.73	0.192 ± 0.033	3.39 ± 0.61
K-Adaptability	65.65 ± 10.07	83.94 ± 5.67	7.84 ± 1.11	0.321 ± 0.058	4.21 ± 0.60
EquiPartition-Online	90.05 ± 6.66	98.03 ± 1.49	1.37 ± 0.23	0.157 ± 0.023	1.86 ± 0.28
Multi-Institution Target Regions					
Method	**Full Cert. (%)**	**Cert. Vol. (%)**	**Partitions**	**Verif. Time (s)**	**Calls**
Fixed-Institution	21.18 ± 11.16	33.03 ± 10.43	9.18 ± 1.27	0.225 ± 0.050	5.05 ± 0.73
Adaptive-Action	42.71 ± 15.15	61.48 ± 11.11	7.95 ± 1.20	0.285 ± 0.065	4.51 ± 0.74
Hard-Partition	58.77 ± 13.77	76.61 ± 8.51	14.66 ± 2.16	0.589 ± 0.142	6.89 ± 1.17
Regional NI	47.98 ± 15.68	63.58 ± 11.81	6.56 ± 0.94	0.445 ± 0.097	3.86 ± 0.68
Affine Recourse	39.66 ± 12.66	57.01 ± 10.45	6.25 ± 0.87	0.413 ± 0.093	3.87 ± 0.68
K-Adaptability	47.45 ± 15.36	67.08 ± 11.15	11.05 ± 2.65	0.773 ± 0.231	5.53 ± 1.28
EquiPartition-Online	80.93 ± 10.14	94.29 ± 3.51	1.78 ± 0.49	0.331 ± 0.070	2.00 ± 0.39

Table 3. Selection controls: (A) 40 unfiltered controlled markets and (B–C) regions with one recorded dominant deviator. Values are means ± sample standard deviations; Ref. cov./FE denotes recorded-reference coverage (%)/false exclusions.

Panel A: Unfiltered Random Controlled Markets							
Method	**Ref. Cov./FE**	**Proof-Tree Nodes**	**Verifier Calls**	**Unres. Vol. (%)**	**Runtime (s)**	**Completed**	**Markets**
Nash-B&B	100/0	3426 ± 2834	2345 ± 1822	1.81 ± 1.07	37.9 ± 42.1	40/40	40
Fixed-Institution	100/0	3148 ± 2291	3279 ± 2410	1.42 ± 0.88	49.6 ± 48.3	40/40	40
EquiPartition-Online	100/0	1087 ± 716	1321 ± 853	0.71 ± 0.45	66.8 ± 55.4	40/40	40
EquiPartition-Predict+Repair	100/0	1164 ± 762	1458 ± 957	0.74 ± 0.47	31.5 ± 28.6	40/40	40
Panel B: Controlled Single-Dominant-Deviator Regions (40 Markets, 119 Regions)							
Method		**Full Cert. (%)**	**Cert. Vol. (%)**	**Partitions**	**Verif. Time (s)**	**Calls**	
Fixed-Institution		93.2 ± 6.84	96.8 ± 2.91	1.88 ± 0.37	0.078 ± 0.014	1.21 ± 0.19	
Adaptive-Action		95.1 ± 5.92	97.4 ± 2.23	1.21 ± 0.28	0.121 ± 0.026	1.56 ± 0.27	
Hard-Partition		86.4 ± 9.41	94.2 ± 4.18	2.98 ± 0.81	0.223 ± 0.057	2.96 ± 0.71	
Regional NI		82.6 ± 10.18	91.8 ± 4.72	2.37 ± 0.54	0.184 ± 0.042	2.69 ± 0.56	
Affine Recourse		88.7 ± 8.73	94.9 ± 3.86	1.47 ± 0.33	0.147 ± 0.035	2.14 ± 0.43	
K-Adaptability		84.9 ± 9.06	93.6 ± 4.01	3.42 ± 0.77	0.251 ± 0.068	2.81 ± 0.69	
EquiPartition-Online		92.4 ± 7.38	96.5 ± 2.64	2.12 ± 0.42	0.162 ± 0.031	1.60 ± 0.34	
Panel C: Multi-Institution Single-Dominant-Deviator Regions (120 Markets, 337 Regions)							
Method		**Full Cert. (%)**	**Cert. Vol. (%)**	**Partitions**	**Verif. Time (s)**	**Calls**	
Fixed-Institution		78.9 ± 12.03	88.9 ± 6.33	2.54 ± 0.61	0.174 ± 0.046	2.39 ± 0.52	
Adaptive-Action		84.3 ± 10.71	93.5 ± 5.12	3.30 ± 0.73	0.298 ± 0.083	2.84 ± 0.68	
Hard-Partition		73.5 ± 13.48	87.2 ± 7.04	5.72 ± 1.24	0.561 ± 0.162	4.44 ± 1.08	
Regional NI		81.2 ± 11.92	94.8 ± 5.89	2.16 ± 0.58	0.372 ± 0.097	2.49 ± 0.59	
Affine Recourse		76.6 ± 14.26	86.9 ± 7.42	4.31 ± 0.96	0.315 ± 0.081	2.27 ± 0.48	
K-Adaptability		86.1 ± 9.88	95.7 ± 4.73	3.05 ± 0.69	0.487 ± 0.134	2.66 ± 0.64	
EquiPartition-Online		83.0 ± 11.37	91.5 ± 6.11	2.05 ± 0.65	0.341 ± 0.091	2.71 ± 0.61	

Figure 4 visualizes this difference on a representative two-institution market. The maximum-regret responsibility changes along a curved boundary that crosses the target box. The fixed-institution construction represents that box with 25 axis-aligned cells, whereas the player-adaptive weight changes continuously across the interface and covers the target region with one certificate. Along the reported cross-section, the certificate value remains between 0.00315 and 0.00727, above the equilibrium tolerance of 0.001. The first institution's selected deviation falls below the threshold after the responsibility switch, while the combined certificate preserves a positive regional margin. The visualization supports the mechanism measured in Table 2: the principal gain is a reduction in external spatial fragmentation, not a reduction in the cost of every verifier call.

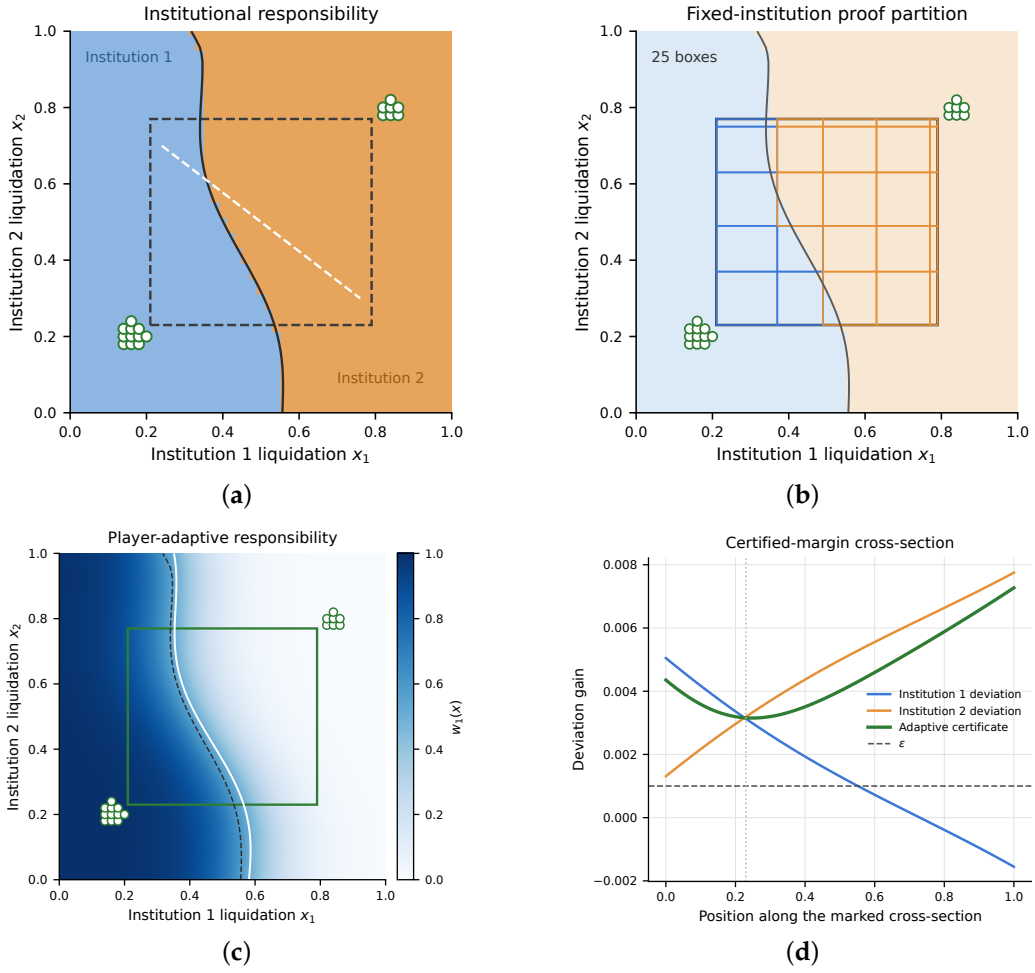


Figure 4. Institutional switching in a representative controlled market: (**a**) a non-axis-aligned dominant-deviator boundary and reference points; (**b**) the fixed-institution partition; (**c**) the adaptive responsibility weight; and (**d**) a cross-section with the certificate above $\varepsilon = 10^{-3}$.

5.3. Ablation Study

The full configuration is EquiPartition-Predict+Repair. *Static Deviation* retains player-adaptive responsibility weights but makes the deviation actions state-independent, *Hard Assignment* replaces continuous responsibility weights with a hard institution assignment, and *Without Repair* bypasses the bounded repair step. Fixed-Institution is the corresponding fixed-player regional proof. These variants leave the rigorous verifier and outer retain-or-split logic unchanged. Table 4 reports all five configurations under this common protocol.

Table 4. Core ablations. Outer-box exclusion rate is the fraction of processed boxes removed before retain-or-split fallback; the finite recorded-reference audit gives 100% coverage and zero intersecting exclusions for every variant.

Controlled Financial Markets				
Variant	**Box Excl. (%)**	**Cert. Vol. (%)**	**Nodes**	**Runtime (s)**
Full EquiPartition	41.20 ± 1.74	68.88 ± 8.23	1297 ± 533	41.3 ± 26.5
Fixed-Institution	33.19 ± 3.78	58.16 ± 10.13	7883 ± 4240	95.8 ± 69.4
Without Repair	38.98 ± 2.52	62.83 ± 7.00	2038 ± 881	46.1 ± 35.5
Static Deviation	23.48 ± 4.99	54.27 ± 8.08	2156 ± 826	51.3 ± 34.8
Hard Assignment	27.15 ± 4.50	57.92 ± 7.13	1887 ± 821	48.8 ± 32.6
Multi-Institution Financial Markets				
Variant	**Box Excl. (%)**	**Cert. Vol. (%)**	**Nodes**	**Runtime (s)**
Full EquiPartition	35.44 ± 3.90	62.70 ± 8.47	8585 ± 5948	2385.8 ± 2010.2
Fixed-Institution	25.03 ± 4.88	49.85 ± 15.75	53,693 ± 39,188	4334.8 ± 2509.1
Without Repair	32.30 ± 4.30	57.54 ± 9.40	15,763 ± 11,132	2882.5 ± 2274.3
Static Deviation	18.08 ± 5.43	49.73 ± 9.34	15,248 ± 11,027	2864.7 ± 2180.8
Hard Assignment	21.95 ± 5.08	51.62 ± 9.72	12,792 ± 8854	2749.4 ± 2163.3

All ablations retain the same verifier-authorized deletion rule, so the completeness argument is unchanged. The finite recorded-reference audit also retained every recorded component and found no intersecting exclusion; their empirical differences concern proof efficiency. On the controlled markets, the full method achieved a 41.20% outer-box exclusion rate and 68.88% certified volume. Fixing the institution reduced these values to 33.19% and 58.16%, while increasing the proof tree from 1297 to 7883 nodes and runtime from 41.3 to 95.8 s. Static Deviations and Hard Assignments also reduced certified volume by 14.61 and 10.96 percentage points; their trees contained 66.2% and 45.5% more nodes than the full method, respectively. Removing repair produced a smaller change in certificate rate, but still increased nodes by 57.2% relative to the full configuration and increased runtime from 41.3 to 46.1 s.

The multi-institution ablations show the same ordering in the reported means. The full method obtained a 35.44% outer-box exclusion rate and 62.70% certified volume. Fixed-Institution reduced these values by 10.41 and 12.85 percentage points and increased the tree from 8585 to 53,693 nodes. Static Deviation and Hard Assignment lowered certified volume to 49.73% and 51.62%, yielding 15,248 and 12,792 nodes. Without Repair retained 57.54% certified volume and 15,763 nodes, with a capped runtime of 2882.5 s rather than 2385.8 s. The ablations therefore associate the largest structural loss with removing player adaptivity, while state-dependent deviations, continuous responsibility weights, and bounded repair each contribute additional empirical gains under the common solver protocol.

5.4. Amortized Certificate Synthesis and End-to-End Cost

Table 5 compares neutral-start, parent-inherited, and learned certificate initialization, while Figure 5 visualizes the learned and neutral-start quality–cost diagnostics. The region-level columns are on-policy diagnostics, whereas Predict and Predict+Repair are directly paired because they share the same proposals and candidate denominator. On controlled markets, direct prediction certified 46.59% of candidate regions in 1.43 ms. Bounded repair raised this paired rate to 74.03%, a gain of 27.44 percentage points, while increasing synthesis time to 2.69 ms. Neutral-start Online synthesis required 44.39 ms on its trajectory. At the matched market level, Predict+Repair used 1297 nodes, close to the 1260 nodes

of Online synthesis, and reduced mean runtime from 97.7 to 41.3 s. Relative to direct prediction, repair reduced the proof tree by 36.4% and runtime by 10.5%.

Table 5. Amortized certificate synthesis and solver-level cost. Candidate certification and synthesis time are on-policy region-level diagnostics; nodes and runtime are market-level. Predict and Predict+Repair share the direct proposals, and Parent-WarmStart+Repair uses no learned proposal.

Controlled Financial Markets					
Method	**Direct Cert. (%)**	**Final Cert. (%)**	**Synth. Time (ms)**	**Nodes**	**Runtime (s)**
EquiPartition-Online	66.28 ± 7.65	66.28 ± 7.65	44.39 ± 5.37	1260 ± 532	97.7 ± 68.8
Parent-WarmStart+Repair	55.48 ± 8.63	71.16 ± 7.21	7.58 ± 1.17	1467 ± 794	69.8 ± 51.9
EquiPartition-Predict	46.59 ± 8.24	46.59 ± 8.24	1.43 ± 0.07	2038 ± 881	46.1 ± 35.5
EquiPartition-Predict+Repair	46.59 ± 8.24	74.03 ± 6.59	2.69 ± 0.26	1297 ± 533	41.3 ± 26.5
EquiPartition-Oracle	100.00 ± 0.00	100.00 ± 0.00	73.70 ± 9.67	715 ± 377	297.9 ± 421.9
Multi-Institution Financial Markets					
Method	**Direct Cert. (%)**	**Final Cert. (%)**	**Synth. Time (ms)**	**Nodes**	**Runtime (s)**
EquiPartition-Online	46.87 ± 9.89	46.87 ± 9.89	318.90 ± 49.66	8663 ± 6765	3924.3 ± 2315.7
Parent-WarmStart+Repair	39.72 ± 9.27	50.06 ± 9.83	61.42 ± 9.64	11,892 ± 8721	3278.6 ± 2268.4
EquiPartition-Predict	25.11 ± 7.85	25.11 ± 7.85	5.07 ± 0.43	15,763 ± 11,132	2882.5 ± 2274.3
EquiPartition-Predict+Repair	25.11 ± 7.85	57.62 ± 9.60	18.75 ± 2.28	8585 ± 5948	2385.8 ± 2010.2
EquiPartition-Oracle	100.00 ± 0.00	100.00 ± 0.00	508.94 ± 82.66	3731 ± 3106	4885.7 ± 2341.7

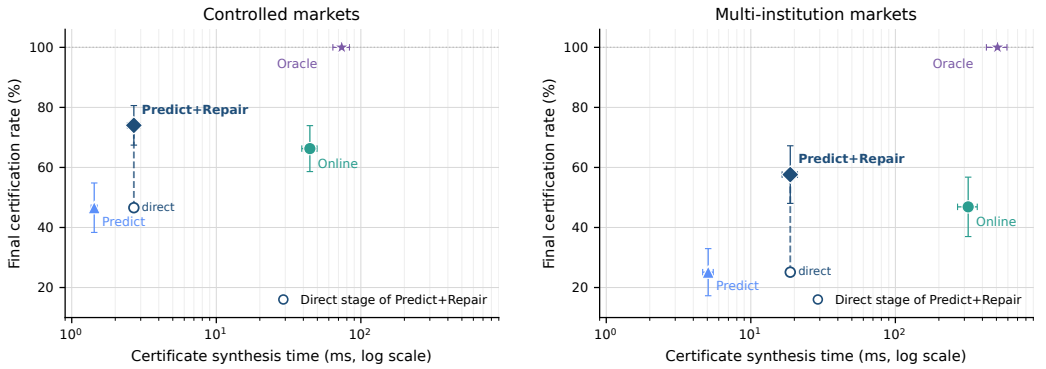


Figure 5. Certificate quality–cost diagnostics on (**a**) controlled and (**b**) multi-institution markets. Filled markers show final certification, the open marker shows the paired pre-repair rate, and error bars show variation across markets.

The multi-institution results show the same paired repair effect. Direct prediction certified 25.11% of candidate regions, and repair increased the rate on the shared panel to 57.62%. Predict+Repair required 18.75 ms of synthesis time, compared with 318.90 ms for neutral-start Online synthesis on its trajectory. At the matched market level, Predict+Repair used 8585 proof-tree nodes and 2385.8 s, compared with 8663 nodes and 3924.3 s for Online synthesis. It also reduced the tree by 45.5% and runtime by 17.2% relative to direct prediction, and completed 113 of 120 markets compared with 96 for Online synthesis.

The non-learning parent warm start isolates the effect of amortized initialization without proposal learning. On controlled markets, it reached 71.16% final certification in 7.58 ms, with 1467 nodes and 69.8 s runtime; Predict+Repair reached 74.03% with 1297 nodes and 41.3 s. On multi-institution markets, Parent-WarmStart+Repair reached 50.06% final certification, used 11,892 nodes, required 3278.6 s, and completed 104 of 120 markets, whereas Predict+Repair reached 57.62%, used 8585 nodes, required 2385.8 s, and completed 113 markets. Parent inheritance therefore captures part of the amortization benefit, while the learned proposal provides additional certificate quality and end-to-end savings under the same repair and verification protocol.

Repair was invoked on 52.81% of controlled candidate regions and 73.60% of multi-institution regions, with 6.99 and 10.88 update steps on average among attempted repairs. This concentrates additional optimization on proposals that do not pass direct verification. The Oracle reference achieved full candidate certification and the smallest proof trees, while its larger synthesis budget produced higher solver cost. Together with the parent-warm-start control, the comparison separates the benefit of generic certificate reuse from the additional benefit of the learned proposal.

Table 6 summarizes the on-policy repair-budget sweep. The direct certification rate is unchanged because the common pre-repair proposal panel is reused, whereas final certification, certified volume, proof-tree size, and runtime follow the trajectory induced by each J. On controlled markets, increasing J from 0 to 20 raised final certification from 46.72% to 74.21% and reduced the mean tree from 2067 to 1294 nodes; the additional increase to $J = 40$ was 1.15 percentage points, while runtime rose from 41.4 to 47.1 s. On multi-institution markets, $J = 20$ raised final certification from 25.34% to 57.84%, reduced nodes from 15,881 to 8537, and produced the lowest capped runtime, 2298.5 s. Its on-policy certified-volume values of 73.04% and 68.71% are therefore separate trajectory summaries from the configuration-sweep K–p reference values of 73.36% and 68.86%. The results show diminishing certification and tree-size returns beyond the 20-update protocol.

Table 6. Repair-budget sensitivity at $K = 2$, $p = 2$ on the fixed held-out subsets. Direct certification uses the common pre-repair panel; remaining certificate and solver metrics follow each budget-specific search trajectory.

					Controlled Markets (24)				
J	Direct Cert. (%)	Final Cert. (%)	Cert. Vol. (%)	Repair Steps	Synth. Time (ms)	Nodes	Runtime (s)	Completed	Markets
0	46.72 ± 8.41	46.72 ± 8.41	61.91 ± 7.38	0.00 ± 0.00	1.42 ± 0.08	2067 ± 918	45.8 ± 34.7	24/24	24
5	46.72 ± 8.41	63.84 ± 7.46	66.38 ± 7.10	4.17 ± 0.62	2.06 ± 0.21	1664 ± 731	40.2 ± 29.8	24/24	24
10	46.72 ± 8.41	70.59 ± 6.89	70.77 ± 6.74	6.21 ± 1.08	2.34 ± 0.28	1439 ± 611	38.7 ± 27.4	24/24	24
20	46.72 ± 8.41	74.21 ± 6.54	73.04 ± 6.49	7.06 ± 1.49	2.78 ± 0.36	1294 ± 551	41.4 ± 28.3	24/24	24
40	46.72 ± 8.41	75.36 ± 6.31	73.82 ± 6.33	8.64 ± 2.17	3.71 ± 0.55	1277 ± 536	47.1 ± 32.9	24/24	24
					Multi-Institution Markets (48)				
J	Direct Cert. (%)	Final Cert. (%)	Cert. Vol. (%)	Repair Steps	Synth. Time (ms)	Nodes	Runtime (s)	Completed	Markets
0	25.34 ± 8.03	25.34 ± 8.03	32.57 ± 9.72	0.00 ± 0.00	5.16 ± 0.47	15,881 ± 11,240	2867.4 ± 2268.1	42/48	48
5	25.34 ± 8.03	43.92 ± 9.41	49.86 ± 9.15	4.44 ± 0.69	10.62 ± 1.38	12,134 ± 8647	2698.6 ± 2112.4	43/48	48
10	25.34 ± 8.03	52.47 ± 9.66	60.63 ± 8.84	7.38 ± 1.25	15.21 ± 1.96	9876 ± 7321	2472.3 ± 1987.2	45/48	48
20	25.34 ± 8.03	57.84 ± 9.38	68.71 ± 8.39	10.62 ± 1.93	19.64 ± 2.61	8537 ± 6458	2298.5 ± 1851.6	46/48	48
40	25.34 ± 8.03	59.18 ± 9.21	70.05 ± 8.26	13.31 ± 2.78	29.83 ± 4.42	8299 ± 6314	2479.1 ± 1974.8	45/48	48

6. Additional Analysis and Discussion

The main results establish the solver-level effect of player-adaptive certificates. This section examines the behavior of the complete enclosure under stronger financial stress, independent reference and verifier diagnostics, the trade-off between certificate expressiveness and verification cost, the runtime source of the amortization gain, and the financial profiles represented by the retained equilibria.

6.1. Robustness Under Financial Stress Shifts

Across the IID condition and the high-leverage, low-liquidity, and hidden-crisis shifts, the finite recorded-reference audit for both EquiPartition variants retained every recorded reference-equilibrium component and every recorded crisis component, with no accepted exclusion intersecting the reference set. Crisis components are the recorded reference states classified as Fire-Sale by Equation (12); the classification is used only for reporting. The hidden-crisis condition increased the mean number of recorded reference components

from 5.12 to 6.63 and the mean number of recorded crisis components from 1.86 to 3.13 in the multi-institution markets, while the audit coverage remained unchanged.

The proposal model is held fixed across all three stress shifts. high leverage changes funding ratios, low liquidity changes price-impact coefficients, and hidden crisis jointly changes funding ratios, price impact, and portfolio overlap. This is a parameter-level distribution-shift evaluation: Table 7 measures transfer to unseen leverage, liquidity, and overlap regimes while retaining the same action space, payoff equation, equilibrium tolerance, and verifier protocol.

Table 7. Completion and capped runtime under IID and financial stress shifts. Online and Predict+Repair report the seed run with median capped runtime across five runs; recorded-reference and crisis-component coverage are 100% with zero intersecting exclusions. Controlled completion is 100% for both variants.

	Controlled Runtime (s)		Multi-Institution Completion (%)		Multi-Institution Runtime (s)	
Shift	**Online**	**Predict+Repair**	**Online**	**Predict+Repair**	**Online**	**Predict+Repair**
IID	97.7	41.3	80.0	94.2	3924.3	2385.8
High Leverage	97.5	78.9	80.0	90.0	4615.3	3414.7
Low Liquidity	134.0	92.2	72.5	86.7	4625.3	3697.3
Hidden Crisis	149.5	124.7	60.0	75.8	5267.4	4318.7

Table 7 shows that Predict+Repair completed all controlled markets under every condition and retained a lower capped runtime than Online synthesis. The reductions were 57.8% under IID sampling, 19.1% under high leverage, 31.2% under low liquidity, and 16.6% under the hidden-crisis shift. On the multi-institution markets, Predict+Repair completed 94.2%, 90.0%, 86.7%, and 75.8% of the instances across the four conditions, compared with 80.0%, 80.0%, 72.5%, and 60.0% for Online synthesis. Its capped runtime was lower by 39.2%, 26.0%, 20.1%, and 18.0%, respectively. The amortized proposal therefore remained useful relative to neutral-start Online synthesis as the financial conditions increased the search burden, while strict verification maintained the same outer-enclosure criterion.

6.2. Reference and Verifier Diagnostics

Table 8 reports the independent component cross-check and the pooled verifier outcomes. The independent pipeline matched 156 of the 157 recorded reference components, found no independently validated component outside the recorded reference set, and reproduced the full component set on 39 of 40 markets. One recorded component was not recovered by the independent pipeline.

Table 8. Independent reference cross-check (A) and pooled verifier outcomes (B) on the controlled test markets.

Panel A: Independent Reference Cross-Check							
Markets	**Reference Components**	**Independent Components**	**Matched**	**Independent-Only**	**Reference-Only**	**Exact-Match Markets**	**Exact-Match Rate**
40	157	156	156	0	1	39/40	97.5%

Panel B: Verifier-Status Frequencies					
Market Group	**Stage**	**Calls**	**Certified Exclusion**	**Not Certified**	**Numerical Failure**
Controlled	Direct	44,732	20,801 (46.50%)	23,668 (52.91%)	263 (0.59%)
Controlled	Post-repair	23,668	12,214 (51.61%)	11,331 (47.87%)	123 (0.52%)
Multi-institution	Direct	984,231	246,783 (25.07%)	724,182 (73.58%)	13,266 (1.35%)
Multi-institution	Post-repair	724,182	317,839 (43.89%)	394,424 (54.46%)	11,919 (1.65%)

At the direct stage, `NumericalFailure` occurred in 0.59% of controlled calls and 1.35% of multi-institution calls; the post-repair frequencies were 0.52% and 1.65%. `NotCertified` accounted for most unsuccessful direct calls: 52.91% on controlled markets and 73.58% on multi-institution markets. Repair converted 51.61% and 43.89% of the corresponding second-stage calls into certified exclusions.

6.3. Certificate Complexity and Computational Trade-Offs

For the formal scaling analysis, let A denote the number of assets, p the common responsibility-weight degree, r an upper bound on the per-coordinate degree of one deviation-gain polynomial on a verification cell, H the number of internal cells processed by one verifier call, and $M(b)$ the cost of a b-bit MPFR operation. A dense tensor representation has

$$N_p = (p+1)^d, \qquad N_g = (r+1)^d, \qquad N_{pg} = (p+r+1)^d.$$

Here N_p is the coefficient count of one responsibility weight, N_g is the gain-coefficient count per atom, and N_{pg} is the coefficient count after multiplication. The full responsibility representation contains KN_p coefficients. A direct tensor implementation has the bounds

$$T_{\text{verify}} = O\big(HK\,N_g\,[A+N_p]M(b)\big), \qquad S_{\text{verify}} = O\big(KN_p + N_{pg}\big).$$

The projected-affine parameters contribute at most $O(Kd^2)$ additional scalars in a general box-constrained game and $O(Kd)$ in the scalar-action fire-sale experiments. In this model $d = n$, while A enters gain construction approximately linearly. At $p = 2$ and $K = 4$, the weight tensors contain $4 \cdot 3^d$ coefficients: 2916 at $d = 6$, 236,196 at $d = 10$, and 57,395,628 at $d = 15$, before gain-product tensors. The evaluated markets reach $d = 6$. A universal intractability cutoff is hardware- and protocol-dependent, but the increase from 2916 coefficients at $d = 6$ to 57,395,628 at $d = 15$ identifies the dense tensor representation as the practical scaling limit. Joint strategy dimension and polynomial degree therefore determine the principal dense-verifier bottleneck.

To assess sensitivity to the number of deviation atoms and the Bernstein degree, we vary $K \in \{1,2,3,4\}$ and $p \in \{1,2,3\}$ for all four EquiPartition variants on the fixed held-out subset described in Section 4.4. Figure 6 plots all 96 configuration summaries using certified region volume and the configuration-sweep mean verification time per region; Online, Predict, and Predict+Repair are five-seed summaries, while Oracle contributes one reference run. For a common degree p in joint strategy dimension d, the K responsibility weights contain $K(p+1)^d$ Bernstein coefficients before simplex constraints, so increased expressiveness also enlarges the verification representation. The controlled-market Pareto frontier ranges from 87.45% certified volume at 1.42 ms to 97.96% at 8.64 ms. The multi-institution frontier spans a wider regime, from 14.60% at 9.86 ms to 95.99% at 49.06 ms.

At the common $K = 2$, $p = 2$ reference configuration, Predict+Repair certified 73.36% of the controlled region volume and 68.86% of the multi-institution volume. Its configuration-sweep mean verification times were 4.55 and 27.94 ms per region. Relative to Predict at the same configuration, the certified-volume values increased by 23.96 and 36.06 percentage points. Predict+Repair also used 47.3% and 57.7% fewer proof-tree nodes and 44.0% and 43.6% less end-to-end runtime than Predict on the two evaluation groups. The sweep reveals an expressiveness–cost trade-off. For Predict+Repair, the $K = 4$, $p = 3$ configuration raised certified volume to 88.73% on controlled markets and 76.51% on multi-institution markets, while the displayed mean verification time per region increased from 4.55 to 9.44 ms (2.07×) and from 27.94 to 64.92 ms (2.32×), respectively. The shared reference configuration is selected by its favorable cross-dataset balance between certified volume and configuration-sweep verification cost.

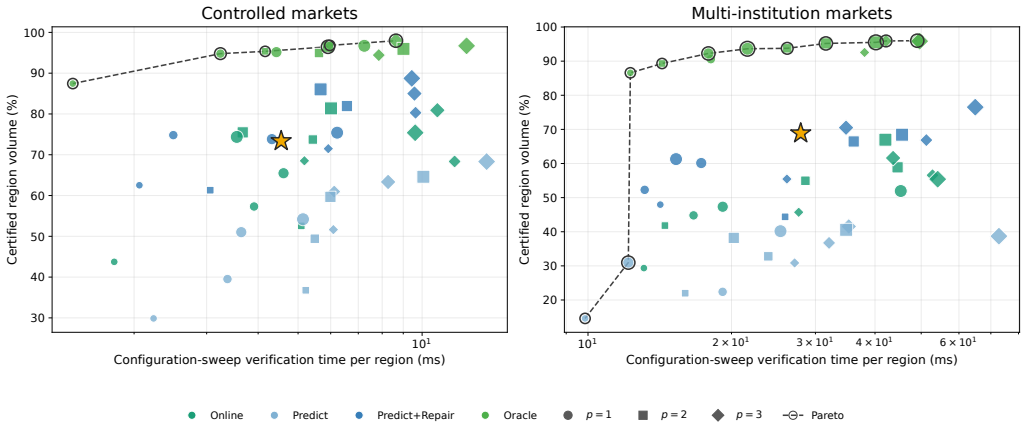
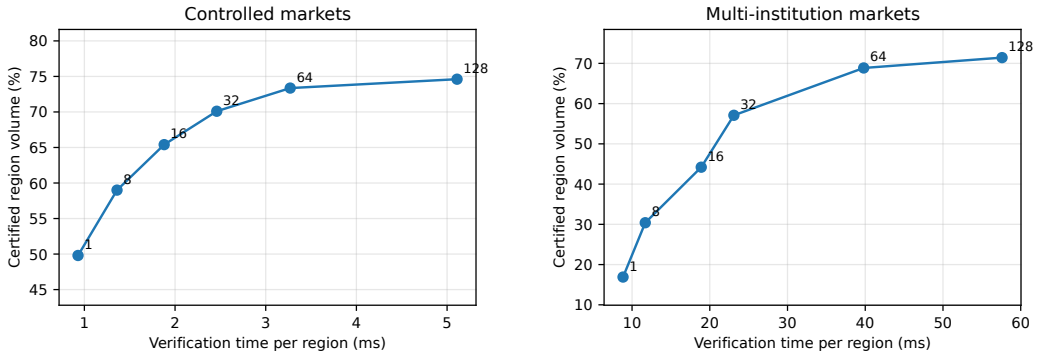


Figure 6. Certificate expressiveness–verification-cost trade-off on the fixed complexity subsets. Colors denote synthesis variants, shapes Bernstein degree p, and marker size atom count K; rings and dashed segments mark the Pareto frontier, and stars the $K = 2$, $p = 2$ reference.

Default Internal-Verifier Accounting

Table 9 reports the default-cap accounting. At the default cap, EquiPartition-Online processes 16.8 internal cells per controlled-region verifier call with a 4.2% exhaustion rate and 26.5 cells per multi-institution call with an 8.9% exhaustion rate. The corresponding full-region certification rates are 90.05% and 80.93%. Hard-Partition processes 6.2 and 13.8 cells on average but requires 0.269 and 0.589 s per region, whereas EquiPartition-Online requires 0.157 and 0.331 s; internal-cell count and per-cell expression cost therefore jointly determine verification time. K-Adaptability processes 27.3 and 44.2 cells, reaches the cap on 18.8% and 35.1% of calls, and requires 0.321 and 0.773 s.

Table 10 and Figure 7 show the internal-cell-cap trade-off under the cap-sweep mean verification-time estimator. Increasing the cell cap from one to 64 raised certified volume from 49.80% to 73.36% on controlled markets and from 16.90% to 68.86% on multi-institution markets, while cap-sweep time increased from 0.93 to 3.27 ms and from 8.84 to 39.80 ms, respectively. The certified-volume values at 64 match the configuration-sweep $K = 2$, $p = 2$ reference, whereas the cap-sweep timer is a separate estimator from the five-seed configuration-sweep times of 4.55 and 27.94 ms. Extending the cap from 64 to 128 added 1.25 and 2.57 percentage points of certified volume but increased cap-sweep time to 5.11 and 57.60 ms. The default 64-cell protocol lies near the flattening portion of both cap-sweep-time–certified-volume curves.



(**a**) (**b**)

Figure 7. Cap-sweep verification-time–certified-volume trade-off on (**a**) controlled and (**b**) multi-institution subsets. Labels show the total cell cap, including the root.

Table 9. Default 64-cell verifier accounting on the switching regions. Full certification, certified volume, and verification time are market-level; cell counts and exhaustion are pooled over verifier calls.

Controlled Target Regions						
Method	**Full Cert. (%)**	**Cert. Vol. (%)**	**Mean Cells**	**P95 Cells**	**Exhaustion (%)**	**Verif. Time (s)**
Fixed-Institution	36.45 ± 14.16	50.48 ± 12.66	10.4	64	7.9	0.109 ± 0.017
Adaptive-Action	53.11 ± 13.26	72.50 ± 8.70	13.7	58	4.6	0.137 ± 0.020
Hard-Partition	68.59 ± 10.84	83.73 ± 6.40	6.2	29	2.8	0.269 ± 0.055
Regional NI	69.01 ± 13.93	81.76 ± 9.01	15.6	64	6.1	0.218 ± 0.040
Affine Recourse	60.04 ± 15.26	76.47 ± 10.20	18.9	64	9.4	0.192 ± 0.033
K-Adaptability	65.65 ± 10.07	83.94 ± 5.67	27.3	64	18.8	0.321 ± 0.058
EquiPartition-Online	90.05 ± 6.66	98.03 ± 1.49	16.8	51	4.2	0.157 ± 0.023
Multi-Institution Target Regions						
Method	**Full Cert. (%)**	**Cert. Vol. (%)**	**Mean Cells**	**P95 Cells**	**Exhaustion (%)**	**Verif. Time (s)**
Fixed-Institution	21.18 ± 11.16	33.03 ± 10.43	24.7	64	19.6	0.225 ± 0.050
Adaptive-Action	42.71 ± 15.15	61.48 ± 11.11	31.9	64	15.2	0.285 ± 0.065
Hard-Partition	58.77 ± 13.77	76.61 ± 8.51	13.8	64	6.4	0.589 ± 0.142
Regional NI	47.98 ± 15.68	63.58 ± 11.81	28.1	64	17.9	0.445 ± 0.097
Affine Recourse	39.66 ± 12.66	57.01 ± 10.45	39.6	64	29.8	0.413 ± 0.093
K-Adaptability	47.45 ± 15.36	67.08 ± 11.15	44.2	64	35.1	0.773 ± 0.231
EquiPartition-Online	80.93 ± 10.14	94.29 ± 3.51	26.5	64	8.9	0.331 ± 0.070

Table 10. Internal-cell-cap sensitivity for Predict+Repair at $K = 2$, $p = 2$. The root cell counts toward the cap; 64 is the default. Verification time is measured by the dedicated cap sweep.

Controlled Markets (24)					
Cell Cap	**Full Cert. (%)**	**Cert. Vol. (%)**	**Internal Cells**	**Exhaustion (%)**	**Cap-Sweep Time (ms)**
1	38.9	49.80	1.0	63.1	0.93
8	48.1	59.00	5.9	46.3	1.36
16	55.3	65.40	10.8	35.8	1.88
32	61.6	70.10	19.4	24.1	2.46
64	66.7	73.36	32.6	14.6	3.27
128	68.2	74.61	47.3	8.3	5.11
Multi-Institution Markets (48)					
Cell Cap	**Full Cert. (%)**	**Cert. Vol. (%)**	**Internal Cells**	**Exhaustion (%)**	**Cap-Sweep Time (ms)**
1	13.1	16.90	1.0	82.3	8.84
8	22.6	30.40	6.8	69.1	11.70
16	31.4	44.20	12.7	57.4	18.90
32	43.2	57.10	24.1	41.2	23.10
64	52.2	68.86	44.2	25.7	39.80
128	55.6	71.43	73.8	14.1	57.60

6.4. Runtime Decomposition and Sources of Runtime Savings

Figure 8 decomposes completed profiling runs for EquiPartition-Online and EquiPartition-Predict+Repair under the IID, low-liquidity, and hidden-crisis conditions. We display these three representative conditions because Table 7 already reports the four-condition comparison. Each stack reports the mean of five mutually exclusive timers. A market contributes when at least one of its five seeds yields both a completed solver run and a complete profile; the eligible seed with the median profiling total supplies that market's component vector. The resulting n is therefore a profiling-availability count and need not equal the five-seed-median completion count in Tables 1 and 7. The stack height also need not equal the all-instance capped runtime. The controlled panel contains 40 profiled markets for every bar. In the multi-institution panel, the profiled counts for Online/Predict+Repair are 108/119, 104/112, and 100/104 across the three displayed conditions.

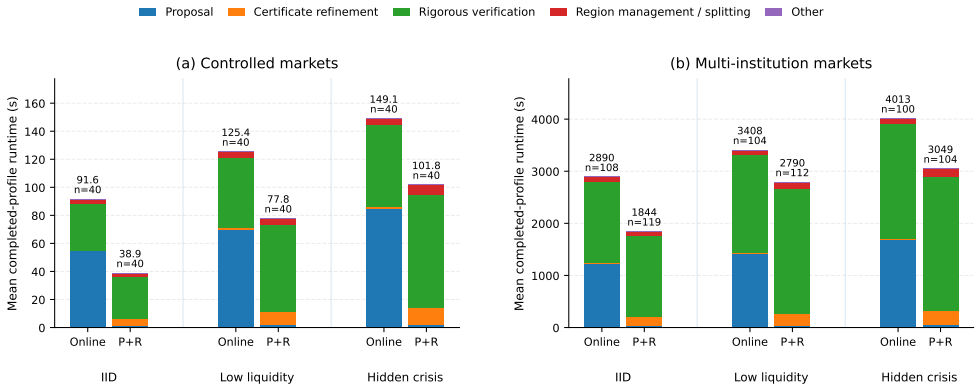


Figure 8. Runtime decomposition for Online and Predict+Repair under IID, low-liquidity, and hidden-crisis conditions. Bars show mean completed-profile time by component; each market contributes the eligible seed run with median profiling total.

On the controlled markets, Predict+Repair reduced mean completed-profile runtime from 91.6 to 38.9 s under IID sampling, from 125.4 to 77.8 s under low liquidity, and from 149.1 to 101.8 s under the hidden-crisis shift. These changes correspond to reductions of 57.6%, 37.9%, and 31.7%. Online synthesis spent 54.4, 69.8, and 84.7 s on per-box proposal optimization, whereas the amortized proposal required only 1.0, 1.8, and 2.3 s. Bounded refinement added 4.9–12.1 s, and rigorous verification became the dominant component of Predict+Repair, but the avoided online proposal cost remained larger than the additional refinement and verification effort.

The multi-institution profiles show the same mechanism at a larger scale. The method-specific completed-profile means were 2890.3 and 1843.6 s for Online and Predict+Repair under IID sampling, 3407.8 and 2790.2 s under low liquidity, and 4012.7 and 3049.3 s under the hidden-crisis shift. Because the profiled counts differ between methods, these values are interpreted as aggregate component summaries rather than paired instance-level reductions. Across the three conditions, the learned proposal component was 97.3–98.0% smaller, while bounded refinement and rigorous verification consumed more time on the harder shifts. The largest measured saving therefore comes from amortizing repeated per-box certificate optimization, which offsets the additional repair and verification work.

6.5. Financial Interpretation of the Enclosed Equilibria

From the retained enclosure, for the synthetic one-period fire-sale model, we selected 48 representative equilibrium states, evenly distributed across the three post hoc classes defined by Equation (12): 16 Mild Liquidation, 16 Partial Deleveraging, and 16 Fire-Sale states. The class labels are descriptive and do not affect the solver or verifier. Every selected state satisfied the common tolerance $\varepsilon = 10^{-3}$; the largest observed maximum regret was 5.58×10^{-4}. Table 11 summarizes their financial profiles.

Table 11. Financial profiles of retained equilibrium states in the synthetic fire-sale model. Values are means ± sample standard deviations over 16 states per class.

Equilibrium Class	States	Aggregate Liquidation (%)	Minimum Price (% of Initial)	Equity Loss (%)	Pressured Institutions (Count)
Mild Liquidation	16	27.90 ± 9.70	80.78 ± 5.57	17.28 ± 8.62	0.69 ± 0.79
Partial Deleveraging	16	48.51 ± 10.10	64.08 ± 8.18	31.52 ± 10.52	1.19 ± 0.75
Fire-Sale	16	72.16 ± 8.38	46.14 ± 8.90	48.57 ± 8.00	2.00 ± 0.73

The class means differ materially across all four reported indicators. Relative to Mild Liquidation, the Fire-Sale states increased aggregate liquidation by 44.26 percentage

points, reduced the minimum asset price by 34.63 percentage points, and increased mean system equity loss from 17.28% to 48.57%, a difference of 31.29 percentage points. Partial Deleveraging occupied an intermediate mean range for these quantities. Among the selected Fire-Sale states, the hidden-crisis condition produced the lowest mean minimum price, 38.99% of its initial level, together with a mean system equity loss of 52.79%. These profiles show that the complete enclosure preserves materially different liquidation and price-impact regimes under one equilibrium criterion, including regimes that cannot be represented simultaneously by a single selected equilibrium.

Taken together, the additional analyses connect the algorithmic results with the financial interpretation of EquiPartition. The independent cross-check matched 156 of 157 recorded components, verifier numerical failures remained between 0.52% and 1.65%, the repair and internal-cell sweeps showed diminishing returns, and the parent-warm-start control quantified the value of learned initialization beyond certificate reuse. The complexity sweep quantifies the cost of richer certificates, and the runtime decomposition shows that the largest measured saving relative to neutral-start Online synthesis comes from amortizing per-box proposal optimization. The resulting enclosure also retains financially distinct liquidation states. The main computational benefit appears when proposal and repair cost less than the search-tree expansion and repeated online optimization avoided by certifying larger regions.

6.6. Limitations and Future Work

The empirical evaluation uses synthetic normalized one-period fire-sale games with box-constrained liquidation rates, common-asset price impact, and the parameter shifts specified in Section 4.1.

The main computational limitation of EquiPartition is the growth of the tensor-product Bernstein representation with the joint strategy dimension. This representation enables whole-box verification and guarantees valid responsibility weights throughout a continuous region, but increasing the atom count or polynomial degree also increases the number of certificate coefficients and the cost of regional bounding. A practical extension is to use sparse or factorized responsibility weights, activate deviation atoms only where they contribute certified margin, and adapt the polynomial degree to the local switching geometry. Piecewise-affine and exact adjustable-policy constructions provide useful principles for developing such compact state-dependent representations [51,52].

The exclusion logic in Equation (5) requires a rigorous regional lower bound on G_B rather than polynomial payoffs specifically. The current verifier obtains this bound by Bernstein conversion on supported piecewise-polynomial cells. For a non-polynomial payoff, the same outer-search theorem applies when a validated enclosure routine supplies outward-rounded bounds for each deviation gain and their weighted sum on every cell. If such an enclosure is unavailable, the verifier returns `NotCertified` or `NumericalFailure`, and the region remains in conservative search.

For a shared feasibility constraint $F_i(z_i, x_{-i}; c) \leq 0$, atom k is regionally legal on cell C only when the verifier establishes

$$\overline{F}_{k,C} \geq \sup_{x \in C} F_{i_k}\big(d_k(x_{-i_k}), x_{-i_k}; c\big), \qquad \overline{F}_{k,C} \leq 0.$$

A shared-constraint extension therefore authorizes exclusion only when every used atom passes this feasibility test and the aggregate deviation certificate also satisfies $\underline{G}_B(\Pi) > \varepsilon$. The present experiments use individually box-constrained actions, for which Equation (8) supplies legality directly. Existing formulations for rational, mixed-integer,

and quasi-linear generalized Nash problems provide compatible tools for bounding these coupled constraints [17–19].

7. Conclusions

This study introduced EquiPartition, a complete approximate-equilibrium enclosure method for strategic financial markets in which deviation responsibility may change across institutions inside a search region. The central contribution is a player-adaptive deviation certificate that combines a finite set of legal unilateral adjustments through state-dependent Bernstein weights. This representation captures institution switches without forcing the outer solver to reproduce the switching geometry through repeated axis-aligned subdivision. Certificate proposals may be obtained online or amortized across related market instances, but every regional exclusion remains subject to the same rigorous outward-rounded verifier. Unsuccessful proposals are repaired under a fixed budget or returned to conservative retain-or-split search. Consequently, prediction quality affects efficiency, whereas the complete outer-enclosure guarantee follows from verifier-authorized deletion and retained terminal boxes.

Across the synthetic one-period fire-sale protocols, EquiPartition-Predict+Repair reduced mean proof-tree size by 83.3% and 83.9%, unresolved volume by 75.8% and 69.0%, and mean capped runtime by 40.5% and 20.2% relative to Nash-B&B. It completed 113 of 120 multi-institution instances, compared with 103 for Nash-B&B, while recorded-reference coverage remained 100% with zero intersecting exclusions. Certificate-level comparisons and ablations associated the largest structural gain with player-adaptive responsibility. Selection controls located the largest regional benefit in responsibility-switching geometry, while Adaptive-Action and K-Adaptability led the single-dominant controls. The independent cross-check matched 156 of 157 recorded components, verifier numerical failures remained below 1.65%, and the repair and internal-cell sweeps showed diminishing returns beyond the default budgets. Parent-WarmStart+Repair captured part of the amortization benefit, whereas Predict+Repair achieved lower proof-tree size and runtime on both market groups and higher completion on the multi-institution group. The runtime decomposition identified avoided per-box certificate optimization as the largest measured source of savings.

Author Contributions: Conceptualization, J.C. and F.X.; methodology, J.C., K.Y., D.L. and F.X.; software, K.Y. and J.C.; validation, J.C., K.Y. and Y.W.; formal analysis, J.C. and K.Y.; investigation, J.C., K.Y. and D.L.; data curation, K.Y. and J.C.; writing—original draft preparation, J.C. and K.Y.; writing—review and editing, all authors; visualization, K.Y. and J.C.; supervision, Y.W. and F.X.; project administration, F.X. All authors have read and agreed to the published version of the manuscript.

Funding: This research was supported by the Natural Science Foundation of Xinjiang Uygur Autonomous Region (Grant No. 2026D01C579), the University-level Scientific Research Fund Project of Xinjiang College of Science & Technology (Grant No. 2026-KYPT37), and JST SPRING, Japan Grant Number JPMJSP2108.

Data Availability Statement: The raw data supporting the conclusions of this article will be made available by the authors on request.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Eisenberg, L.; Noe, T.H. Systemic Risk in Financial Systems. *Manag. Sci.* **2001**, *47*, 236–249. [CrossRef]
2. Bichuch, M.; Feinstein, Z. Endogenous Inverse Demand Functions. *Oper. Res.* **2022**, *70*, 2702–2714. [CrossRef]

3. Braouezec, Y.; Kiani, K. Preventing Price-Mediated Contagion Due to Fire Sales Externalities: Strategic Foundations of Macroprudential Regulation. *Oper. Res.* **2025**, *73*, 40–60. [CrossRef]
4. Bertschinger, N.; Hoefer, M.; Krogmann, S.; Lenzner, P.; Schuldenzucker, S.; Wilhelmi, L. Equilibria and Convergence in Fire Sale Games. In *Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems*; International Foundation for Autonomous Agents and Multiagent Systems: Richland, SC, USA, 2023; pp. 215–223. [CrossRef]
5. Hoefer, M.; Wilhelmi, L. Seniorities and Minimal Clearing in Financial Network Games. In *Proceedings of the Algorithmic Game Theory: 15th International Symposium, SAGT 2022*; Lecture Notes in Computer Science; Springer: Cham, Switzerland, 2022; Volume 13584, pp. 187–204. [CrossRef]
6. Ioannidis, S.D.; de Keijzer, B.; Ventre, C. Strong Approximations and Irrationality in Financial Networks with Derivatives. In *Proceedings of the 49th International Colloquium on Automata, Languages, and Programming*; Leibniz International Proceedings in Informatics; Schloss Dagstuhl–Leibniz-Zentrum für Informatik: Dagstuhl, Germany, 2022; Volume 229, pp. 76:1–76:18. [CrossRef]
7. Besting, L.; Hoefer, M.; Huth, L. Computing Tarski Fixed Points in Financial Networks. In *Proceedings of the 43rd International Symposium on Theoretical Aspects of Computer Science*; Leibniz International Proceedings in Informatics; Schloss Dagstuhl–Leibniz-Zentrum für Informatik: Dagstuhl, Germany, 2026; Volume 364, pp. 14:1–14:18. [CrossRef]
8. Kanellopoulos, P.; Kyropoulou, M.; Zhou, H. Forgiving Debt in Financial Network Games. In *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence*; International Joint Conferences on Artificial Intelligence Organization: Freiburg, Germany, 2022; pp. 335–341. [CrossRef]
9. Kanellopoulos, P.; Kyropoulou, M.; Zhou, H. Debt Transfers in Financial Networks: Complexity and Equilibria. In *Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems*; International Foundation for Autonomous Agents and Multiagent Systems: Richland, SC, USA, 2023; pp. 260–268. [CrossRef]
10. Hoefer, M.; Ventre, C.; Wilhelmi, L. Algorithms for Claims Trading. In *Proceedings of the 41st International Symposium on Theoretical Aspects of Computer Science*; Leibniz International Proceedings in Informatics; Schloss Dagstuhl–Leibniz-Zentrum für Informatik: Dagstuhl, Germany, 2024; Volume 289, pp. 42:1–42:17. [CrossRef]
11. Wang, D.; Gao, Z.; Sheng, L. Distributed Finite-Time and Fixed-Time Nash Equilibrium Seeking for Non-Cooperative Game with Input Saturation. *Mathematics* **2023**, *11*, 2295. [CrossRef]
12. He, L.; Cheng, H.; Zhang, Y. Centralized and Decentralized Event-Triggered Nash Equilibrium-Seeking Strategies for Heterogeneous Multi-Agent Systems. *Mathematics* **2025**, *13*, 419. [CrossRef]
13. Dragotto, G.; Scatamacchia, R. The Zero Regrets Algorithm: Optimizing over Pure Nash Equilibria via Integer Programming. *Inf. J. Comput.* **2023**, *35*, 1143–1160. [CrossRef]
14. Schwarze, S.; Stein, O. A Branch-and-Prune Algorithm for Discrete Nash Equilibrium Problems. *Comput. Optim. Appl.* **2023**, *86*, 491–519. [CrossRef]
15. Kirst, P.; Schwarze, S.; Stein, O. A Branch-and-Bound Algorithm for Nonconvex Nash Equilibrium Problems. *SIAM J. Optim.* **2024**, *34*, 3371–3398. [CrossRef]
16. Černý, J.; Das Gupta, S.; Kroer, C. Spatial Branch-and-Bound for Computing Multiplayer Nash Equilibrium. *Proc. AAAI Conf. Artif. Intell.* **2026**, *40*, 16752–16760. [CrossRef]
17. Harks, T.; Schwarz, J. Generalized Nash Equilibrium Problems with Mixed-Integer Variables. *Math. Program.* **2025**, *209*, 231–277. [CrossRef]
18. Choi, J.; Nie, J.; Tang, X.; Zhong, S. Generalized Nash Equilibrium Problems with Quasi-Linear Constraints. *SIAM J. Optim.* **2026**, *36*, 1589–1618. [CrossRef]
19. Nie, J.; Tang, X.; Zhong, S. Rational Generalized Nash Equilibrium Problems. *SIAM J. Optim.* **2023**, *33*, 1587–1620. [CrossRef]
20. Nie, J.; Ranestad, K.; Tang, X. Algebraic Degrees of Generalized Nash Equilibrium Problems. *Sci. China Math.* **2025**, *68*, 1747–1776. [CrossRef]
21. Nikaidô, H.; Isoda, K. Note on Noncooperative Convex Games. *Pac. J. Math.* **1955**, *5*, 807–815. [CrossRef]
22. Ulbrich, M.; Fritz, J. On Generalized Nash Equilibrium Problems in Infinite-Dimensional Spaces Using Nikaido–Isoda Type Functionals. *Optim. Methods Softw.* **2024**, *39*, 779–809. [CrossRef]
23. Zeng, R. Proximal Analytic Center Cutting Plane Algorithms for Variational Inequalities and Nash Economic Equilibrium. *Mathematics* **2024**, *12*, 177. [CrossRef]
24. Nie, J.; Sun, D.; Tang, X.; Zhang, M. Solving Polynomial Variational Inequality Problems via Lagrange Multiplier Expressions and Moment-SOS Relaxations. *Comput. Optim. Appl.* **2025**, *90*, 361–394. [CrossRef]
25. Ben-Tal, A.; Goryashko, A.; Guslitzer, E.; Nemirovski, A. Adjustable Robust Solutions of Uncertain Linear Programs. *Math. Program.* **2004**, *99*, 351–376. [CrossRef]
26. Hanasusanto, G.A.; Kuhn, D.; Wiesemann, W. *K*-Adaptability in Two-Stage Robust Binary Programming. *Oper. Res.* **2015**, *63*, 877–891. [CrossRef]
27. Malaguti, E.; Monaci, M.; Pruente, J. *K*-Adaptability in Stochastic Optimization. *Math. Program.* **2022**, *196*, 567–595. [CrossRef]

28. Lefebvre, H.; Malaguti, E.; Monaci, M. Adjustable Robust Optimization with Discrete Uncertainty. *Inf. J. Comput.* **2024**, *36*, 78–96. [CrossRef]
29. Kurtz, J. Approximation Guarantees for Min-Max-Min Robust Optimization and *K*-Adaptability under Objective Uncertainty. *SIAM J. Optim.* **2024**, *34*, 2121–2149. [CrossRef]
30. Divya, U.H.; Kumar, J.P. Enhanced Transfer Learning-Based CNN for Abnormal Human Activity Detection in Video Surveillance Using Spatial-Temporal Features. *Cybern. Syst.* **2026**, *57*, 1996–2025. [CrossRef]
31. Chen, J.; Ren, C.; Hu, Y.; Zhang, Y.; Lu, Y.; Li, Q.; You, T.; Rodrigues, J.J.P.C. Dual-Centralized Q-Network-Based Reinforcement Learning for Cooperative Path Planning of Multiple UAVs. *IEEE Trans. Intell. Transp. Syst.* **2025**, *26*, 13232–13246. [CrossRef]
32. Chen, J.; Shu, Q.; Lu, Y.; Zhang, Y.; Wang, Y. QCTF: A Quantized Communication and Transferable Fusion Framework for Multi-Agent Collaborative Perception. *IEEE Trans. Intell. Transp. Syst.* **2025**, *26*, 15013–15027. [CrossRef]
33. Yao, Z.; Zhu, Q.; Qin, P. Data-Driven Secure Consensus Control of Multi-Agent Systems: A Value Iteration Approach. *Circuits Syst. Signal Process.* **2026**, *45*, 8723–8739. [CrossRef]
34. Paulus, M.B.; Zarpellon, G.; Krause, A.; Charlin, L.; Maddison, C. Learning to Cut by Looking Ahead: Cutting Plane Selection via Imitation Learning. In Proceedings of the 39th International Conference on Machine Learning, Proceedings of Machine Learning Research, Baltimore, MD, USA, 17–23 July 2022; Volume 162, pp. 17584–17600.
35. Scavuzzo, L.; Chen, F.; Chételat, D.; Gasse, M.; Lodi, A.; Yorke-Smith, N.; Aardal, K. Learning to Branch with Tree MDPs. In *Proceedings of the Advances in Neural Information Processing Systems*; Curran Associates, Inc.: Red Hook, NY, USA, 2022; Volume 35, pp. 18514–18526. [CrossRef]
36. Thuerck, D.; Sofranac, B.; Pfetsch, M.E.; Pokutta, S. Learning Cuts via Enumeration Oracles. In *Proceedings of the Advances in Neural Information Processing Systems*; Curran Associates, Inc.: Red Hook, NY, USA, 2023; Volume 36, pp. 79108–79123. [CrossRef]
37. Labassi, A.G.; Chételat, D.; Lodi, A. Learning to Compare Nodes in Branch and Bound with Graph Neural Networks. In *Proceedings of the Advances in Neural Information Processing Systems*; Neural Information Processing Systems Foundation, Inc.: New Orleans, LA, USA, 2022; Volume 35, pp. 32000–32010. [CrossRef]
38. Balcan, M.F.F.; Prasad, S.; Sandholm, T.; Vitercik, E. Structural Analysis of Branch-and-Cut and the Learnability of Gomory Mixed Integer Cuts. In *Proceedings of the Advances in Neural Information Processing Systems*; Curran Associates, Inc.: Red Hook, NY, USA, 2022; Volume 35, pp. 33890–33903. [CrossRef]
39. Julien, E.; Postek, K.; Birbil, S.I. Machine Learning for *K*-Adaptability in Two-Stage Robust Optimization. *Inf. J. Comput.* **2025**, *37*, 644–665. [CrossRef]
40. Dumouchelle, J.; Julien, E.; Kurtz, J.; Khalil, E.B. Neur2BiLO: Neural Bilevel Optimization. In *Proceedings of the Advances in Neural Information Processing Systems*; Neural Information Processing Systems Foundation, Inc.: Vancouver, BC, Canada, 2024; Volume 37, pp. 86688–86719. [CrossRef]
41. Scavuzzo, L.; Aardal, K.; Lodi, A.; Yorke-Smith, N. Machine Learning Augmented Branch and Bound for Mixed Integer Linear Programming. *Math. Program.* **2026**, *217*, 123–166. [CrossRef]
42. Xiao, L.; Gudepu, V.; Fumagalli, A. Joint Alarm Correlation and Root Alarm Detection in Optical Transport Networks with Graph-Attention. In *Proceedings of the 2025 IEEE 30th International Workshop on Computer Aided Modeling and Design of Communication Links and Networks*; IEEE: New York, NY, USA, 2025; pp. 1–6. [CrossRef]
43. Xiao, L.; Latha, A.G.; Gudepu, V.; Fumagalli, A. A YANG-Grounded LLM Agent Supporting Multi-Vendor OpenROADM Optical Transport Networks. In *Proceedings of the ACM SIGCOMM 2026 Conference*; Association for Computing Machinery: New York, NY, USA, 2026; pp. 2109–2115. [CrossRef]
44. Eliáš, M.; Kaplan, H.; Mansour, Y.; Moran, S. Learning-Augmented Algorithms with Explicit Predictors. In *Proceedings of the Advances in Neural Information Processing Systems*; Neural Information Processing Systems Foundation, Inc.: Vancouver, BC, Canada, 2024; Volume 37, pp. 97972–98008. [CrossRef]
45. Wu, M.; Kobayashi, N. Learning Simple Interpolants for Linear Integer Arithmetic. In *Proceedings of the Advances in Neural Information Processing Systems*; Neural Information Processing Systems Foundation, Inc.: Red Hook, NY, USA, 2025; Volume 38, pp. 68233–68255. [CrossRef]
46. Nash, J.F. Non-Cooperative Games. *Ann. Math.* **1951**, *54*, 286–295. [CrossRef]
47. Zhang, X.; Yang, C.; Han, R.; Zhang, S. Nash's Existence Theorem for Non-Compact Strategy Sets. *Mathematics* **2024**, *12*, 2017. [CrossRef]
48. Liu, Q.; Liao, Q. Computing Nash Equilibria for Multiplayer Symmetric Games Based on Tensor Form. *Mathematics* **2023**, *11*, 2268. [CrossRef]
49. Han, E.; Bandi, C.; Nohadani, O. On Finite Adaptability in Two-Stage Distributionally Robust Optimization. *Oper. Res.* **2023**, *71*, 2307–2327. [CrossRef]
50. Wei, N.; Zhang, P. Adjustability in Robust Linear Optimization. *Math. Program.* **2024**, *208*, 581–628. [CrossRef]
51. Lefebvre, H.; Malaguti, E.; Monaci, M. Exact Approaches for Convex Adjustable Robust Optimization. *Math. Program.* **2026**, *218*, 439–476. [CrossRef]

52. Thomä, S.; Walther, G.; Schiffer, M. Designing Tractable Piecewise Affine Policies for Multi-Stage Adjustable Robust Optimization. *Math. Program.* **2024**, *208*, 661–716. [CrossRef]
53. Li, S.; Ouyang, W.; Paulus, M.; Wu, C. Learning to Configure Separators in Branch-and-Cut. In *Proceedings of the Advances in Neural Information Processing Systems*; Neural Information Processing Systems Foundation, Inc.: New Orleans, LA, USA, 2023; Volume 36. [CrossRef]
54. Paulus, M.B.; Krause, A. Learning to Dive in Branch and Bound. In *Proceedings of the Advances in Neural Information Processing Systems*; Neural Information Processing Systems Foundation, Inc.: New Orleans, LA, USA, 2023; Volume 36. [CrossRef]
55. Huang, T.; Ferber, A.M.; Tian, Y.; Dilkina, B.; Steiner, B. Searching Large Neighborhoods for Integer Linear Programs with Contrastive Learning. In *Proceedings of the 40th International Conference on Machine Learning*; Proceedings of Machine Learning Research, PMLR: Honolulu, HI, USA, 2023; Volume 202, pp. 13869–13890.
56. Ye, H.; Xu, H.; Wang, H.; Wang, C.; Jiang, Y. GNN&GBDT-Guided Fast Optimizing Framework for Large-Scale Integer Programming. In *Proceedings of the 40th International Conference on Machine Learning*; Proceedings of Machine Learning Research, PMLR: Honolulu, HI, USA, 2023; Volume 202, pp. 39864–39878.
57. Shajeena, J.; Govindasamy, B.; Gnanasundaram, M.; Joel, M.R. Mobile-Le Harmonic Fusion Network for Object Recognition and SiamMoT-Based Multi-Object Tracking Using Video Surveillance. *Cybern. Syst.* **2026**, *57*, 866–896. [CrossRef]
58. Wu, X.; Dong, J.; Bao, W.; Zou, B.; Wang, L.; Wang, H. Augmented Intelligence of Things for Emergency Vehicle Secure Trajectory Prediction and Task Offloading. *IEEE Internet Things J.* **2024**, *11*, 36030–36043. [CrossRef]
59. Karaca, B.; Ünlü, K.D.; Türkan, S. Wavelet-Enhanced Sequence-to-Sequence Modeling with Attention Mechanism for Short-Term Wind Power Forecasting. *Cybern. Syst.* **2026**, *57*, 485–533. [CrossRef]
60. Yang, B.; Li, Y.; Lin, H.; Chen, X.; Ni, K.; Li, W.; Shen, J.; Liang, T. Multi-Objective Design Method of Concrete Mix Ratio Based on Machine Learning: Forward Prediction and Reverse Design. *Arch. Civ. Mech. Eng.* **2026**, *26*, 57. [CrossRef]
61. Yang, G.; Jiang, Y.; Wang, S.; Chen, K. VinsFusion-Line: Binocular Vision Inertial Navigation Real-Time SLAM System Based on Line Features. *Cybern. Syst.* **2026**, *57*, 350–375. [CrossRef]
62. Gladson, V.B.; Deepa, G.; Jothi, S.F.E.; Mahalakshmi, L. Image Steganography with Security Using Massive Threefold Attentional Residual GAN Optimized by Chaotic PSO Algorithm. *Cybern. Syst.* **2026**, *57*, 1277–1308. [CrossRef]
63. Giacobbe, M.; Kroening, D.; Pal, A.; Tautschnig, M. Neural Model Checking. In *Proceedings of the Advances in Neural Information Processing Systems*; Neural Information Processing Systems Foundation, Inc.: Vancouver, BC, Canada, 2024; Volume 37. [CrossRef]
64. Zhang, H.; Qin, Z.; Gao, S.; Clark, A. SEEV: Synthesis with Efficient Exact Verification for ReLU Neural Barrier Functions. In *Proceedings of the Advances in Neural Information Processing Systems*; Neural Information Processing Systems Foundation, Inc.: Vancouver, BC, Canada, 2024; Volume 37. [CrossRef]
65. Indu, V.T.; Priyadharsini, S.S. Deep Residual Squeeze Excited Region Convolutional Network: A Robust Deep Learning Framework for Precision Agriculture and Plant Disease Diagnosis. *Cybern. Syst.* **2026**, *57*, 1356–1385. [CrossRef]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.