
PENALTY-FREE SDDP: FEASIBILITY CUTS FOR ROBUST MULTI-STAGE STOCHASTIC OPTIMIZATION IN ENERGY PLANNING

A PREPRINT

Guilherme Freitas
PSR
Rio de Janeiro, Brazil
guilhermefreitas@psr-inc.com

Luiz Carlos da Costa Junior
PSR
Rio de Janeiro, Brazil
luizcarlos@psr-inc.com

Tiago Andrade
PSR
Rio de Janeiro, Brazil
tiago.andrade@psr-inc.com

Alexandre Street
Department of Electrical Engineering
PUC-Rio
Rio de Janeiro, Brazil
street@puc-rio.br

ABSTRACT

Multi-stage decision problems under uncertainty can be efficiently solved with the Stochastic Dual Dynamic Programming (SDDP) algorithm. However, traditional implementations require all stage problems to be feasible. Feasibility is usually enforced by adding slack variables and penalizing them in the objective function, a process that depends on case-specific calibration and often distorts the economic interpretation of results. This paper proposes the *Penalty-Free SDDP*, an extension that introduces a *Future Feasibility Function* alongside the traditional Future Cost Function. The new recursion handles infeasibilities automatically, distinguishing between temporary and *truly infeasible* cases, and propagates feasibility information across stages through dedicated feasibility cuts. The approach was validated in a large-scale deterministic case inspired by the Brazilian hydrothermal system, achieving equivalent feasibility to the benchmark solution while eliminating miscalibrated artificial penalties. Results confirm its robustness and practicality as a foundation for future stochastic, multi-stage applications.

Keywords Penalty-Free SDDP · feasibility cuts · multi-stage stochastic optimization · Benders decomposition · hydrothermal scheduling · energy planning.

1 Introduction

Multistage decision problems under uncertainty arise in several domains, including energy systems, finance, logistics, and water management. These problems require sequential decisions that depend on both current information and uncertain future outcomes. Solving them in their extensive form is computationally intractable because of the exponential growth of the scenario tree. To overcome this limitation, *Stochastic Dual Dynamic Programming* (SDDP) [1] is the state-of-the-art algorithm, relying on decomposition to approximate the cost-to-go functions, which represents a decision policy that guide optimal decisions over time.

A general assumption in traditional SDDP is that of *relatively complete recourse*, which guarantees that every realization of uncertainty leads to a feasible problem. In practice, this assumption is rarely satisfied and feasibility is typically enforced by introducing slack variables penalized in the objective function. Although this artifice allows the modified problem to be numerically solved, it mixes real economic costs with artificial penalties, which lack clear interpretation and must be subjectively calibrated. The calibration process itself is not trivial, especially when violations involve quantities of different natures such as hydraulic limits, energy balance, or environmental targets. As a consequence, the

modified problem can produce distorted price signals and misleading economic interpretations, ultimately leading to suboptimal planning decisions.

The situation is particularly problematic when constraint violations are associated with low-probability events. Their impact becomes diluted in the expected cost, making it difficult to guarantee constraint satisfaction unless penalties are increased to unrealistic values or risk-averse formulations are introduced, both of which further bias the results. In the former case, where penalties are artificially inflated to enforce compliance, convergence tends to deteriorate, and numerical instability may arise.

As modern planning problems grow in complexity, driven by renewable integration, regulatory requirements, and sustainability goals, there is a growing need for methodologies that can handle feasibility in a rational and economically consistent way.

This paper proposes a general extension to SDDP that eliminates the need for artificial penalties by introducing *Benders feasibility cuts*. The proposed *Penalty-Free SDDP* expands the classical algorithm with an additional recursive function, the *Future Feasibility Function (FFF)*, which quantifies the value of preserving flexibility to meet future constraints. The method maintains the decomposition structure and convergence properties of SDDP, while ensuring model robustness and interpretability. Although the proposed framework is general and applicable to any multi-stage stochastic optimization problem, its performance is illustrated here through an application to long-term energy planning under operational constraints.

2 Background and Related Work

The foundation of the proposed approach lies in the decomposition principles introduced by Benders [2], later extended to stochastic optimization by Van Slyke and Wets [3]. In this framework, a large-scale problem is decomposed into a master problem and a set of subproblems. Information from the subproblems is then summarized by linear inequalities, known as *Benders cuts*, that iteratively refine the master problem. Two types of cuts are typically generated: *optimality cuts*, when the subproblem is feasible, and *feasibility cuts*, when it becomes infeasible after the master's decision.

Building upon this idea, Pereira and Pinto [1] proposed the Stochastic Dual Dynamic Programming (SDDP) algorithm as a multi-stage extension of the L-shaped method. SDDP combines Benders-type decomposition with Monte Carlo sampling to approximate the stochastic process, recursively constructing a piecewise-linear representation of the expected future cost at each stage through a set of linear cuts, known as the *Future Cost Function (FCF)*. This combination of sampling and decomposition enables the solution of complex multi-stage stochastic problems while avoiding the exponential growth of the scenario tree that would result from an explicit formulation.

Several works have investigated alternatives to the penalty-based approach, including the generation of feasibility cuts within the SDDP framework [4, 5, 6, 7, 8]. These studies show that feasibility cuts can accelerate convergence and improve numerical stability in problems that are ultimately feasible. However, most formulations rely on simplified or problem-specific assumptions, do not address the continuation of the algorithm in truly infeasible cases, or mix feasibility cuts with risk-based formulations that alter the original objective function. As a result, a general penalty-free extension of SDDP that maintains its decomposition structure and robustness has not yet been fully established [9, 10, 11].

The present work addresses this gap by formulating a *Penalty-Free SDDP* algorithm that explicitly integrates feasibility cuts into the recursive process while preserving the economic meaning of the cost function. This approach ensures feasibility handling without the need for artificial penalties and provides a consistent framework applicable to a broad class of multi-stage stochastic optimization problems.

3 Penalty-Free SDDP Framework

3.1 Classical SDDP with penalized slacks

In the classical penalty-based formulation, feasibility is ensured at each stage by introducing nonnegative slack variables that are penalized in the objective function. Let x_t be the stage- t decision, $s_t \geq 0$ the slack vector associated with potentially violated constraints, and let $w_t \geq 0$ denote the penalty weights. Using a generic stagewise functional $\mathcal{R}_t[\cdot]$ (e.g. expectation or risk measure), the stage problem can be formulated as

$$\begin{aligned}
 \text{FCF}_t(x_{t-1}, \xi_{t-1}) = \\
 \min_{x_t, s_t} \quad & c_t^\top x_t + w_t^\top s_t + \mathcal{R}_t[\widetilde{\text{FCF}}_{t+1}(x_t, \xi_t)] \\
 \text{s.t.} \quad & A_t x_t + s_t \geq b_t - D_t x_{t-1}, \\
 & x_t \geq 0, \\
 & s_t \geq 0.
 \end{aligned} \tag{1}$$

Here, $\widetilde{\text{FCF}}_{t+1}$ is the piecewise-linear approximation of the Future Cost Function constructed from Benders optimality cuts. The penalty vector w_t is chosen “large enough” to discourage violations, but its calibration is problem-dependent and lacks direct economic meaning. As discussed in the introduction, this artifice allows the modified problem to be solved at every iteration while blending economic costs with artificial penalties, which complicates interpretation and may affect numerical behavior, especially when low-probability violations are present.

3.2 Penalty-Free SDDP formulation

The proposed formulation extends the traditional SDDP recursion by splitting the solution process into two coordinated stages: a *feasibility problem*, responsible for identifying and quantifying violations of the system constraints, and an *optimality problem*, which computes the economic decision based on the feasible domain defined by the previous step. This separation allows the algorithm to treat feasibility and optimality independently, removing artificial penalties from the cost function while maintaining the decomposition structure of SDDP.

Step 1: Feasibility problem (strict, all scenarios). At stage t , feasibility is enforced *for all* realizations $\xi_t \in \Xi_t$. Nonnegative slacks s_t absorb constraint violations and are weighted by $w_t > 0$ to express preferences when true infeasibilities occur. Unlike the penalized formulation of classical SDDP, these weights do not establish a financial trade-off between economic costs and constraint violations; instead, they serve only to prioritize among different types of violations (e.g., hydraulic, balance, or environmental) in situations where full feasibility cannot be achieved.

The feasibility recursion does not rely on probabilistic aggregation (such as expectation or CVaR). Instead, it follows a worst-case principle, equivalent to a conservative risk operator that enforces feasibility across all realizations. Algorithmically, this principle is implemented by retaining the *union of feasibility cuts* encountered across sampled scenarios, guaranteeing that the solution remains feasible under every possible realization. A compact robust formulation is:

$$\begin{aligned}
 \text{FFF}_t(x_{t-1}, \xi_{t-1}) = \\
 \min_{x_t, s_t, \beta_t} \quad & w_t^\top s_t + \beta_t \\
 \text{s.t.} \quad & A_t x_t + s_t \geq b_t - D_t x_{t-1}, \\
 & \beta_t \geq \widetilde{\text{FFF}}_{t+1}(x_t, \xi_t), \quad \forall \xi_t \in \Xi_t, \\
 & x_t \geq 0, \\
 & s_t \geq 0.
 \end{aligned} \tag{2}$$

This problem computes the minimum infeasibility required to satisfy all constraints for a given state x_{t-1} , taking into account both immediate and future constraint violations. Its objective combines two terms: (i) the current infeasibility, measured by the weighted slacks $w_t^\top s_t$, and (ii) the recursive infeasibility β_t , which represents the value of the Future Feasibility Function (FFF) at the next stage. The solution provides two key quantities: the minimal slack vector s_t^* and the cumulative infeasibility β_t^* . These values quantify the degree of constraint violation at the current and future stages, respectively, and are used to restrict the feasible region in the subsequent optimality problem. In particular, s_t^* defines admissible bounds for local constraint relaxations, while β_t^* limits the propagation of infeasibility in the recursive structure.

Each time a violation is observed in any scenario, a new feasibility cut is generated and retained, progressively constructing a piecewise-linear outer approximation of the feasible domain. This “union-of-cuts” approach ensures that both s_t^* and β_t^* capture all relevant infeasibilities, maintaining strict feasibility across forward and backward passes while preserving the stagewise decomposition structure of the algorithm.

Step 2: Optimality problem with feasibility cuts. The second step solves the cost-minimization subproblem, where the previously computed infeasibility bounds s_t^* and β_t^* are enforced as upper limits for the slack and cumulative infeasibility

variables, respectively. This formulation incorporates all feasibility information into the economic optimization while keeping the objective function free of artificial penalties.

Formally, the stage problem at time t and realization ξ_t is expressed as:

$$\begin{aligned} \text{FCF}_t(x_{t-1}, \xi_{t-1}) = & \\ \min_{x_t, s_t, \beta_t} & c_t^\top x_t + \mathcal{R}_t[\widetilde{\text{FCF}}_{t+1}(x_t, \xi_t)] \\ \text{s.t.} & A_t x_t + s_t \geq b_t - D_t x_{t-1}, \\ & x_t \geq 0, \\ & 0 \leq s_t \leq s_t^*, \\ & \beta_t \leq \beta_t^*, \\ & \beta_t \geq \widetilde{\text{FFF}}_{t+1}(x_t, \xi_t), \quad \forall \xi_t \in \Xi_t. \end{aligned} \tag{3}$$

When $\beta_t^* = 0$, the current state is in the approximated FFF domain. Then, the cumulative feasibility slack β_t becomes inactive, and the problem coincides with the standard SDDP formulation with feasibility cuts. Otherwise, when $\beta_t^* > 0$, it propagates the residual infeasibility backward, allowing the model to adjust the current state variables to anticipate constraint violations. In some visited states, the problem may still be infeasible — a situation that arises when maintaining feasibility requires multi-stage anticipation or when a true structural infeasibility is present in the data but has not yet been fully detected.

The constraint $\beta_t \geq \widetilde{\text{FFF}}_{t+1}(x_t, \xi_t)$ enforces consistency with the recursive feasibility representation, ensuring that any remaining violation at future stages is recognized at the current one. Through this coupling, the algorithm becomes *feasibility-aware*: state decisions such as storage levels, thermal fuel contracts, emission budgets or other intertemporal controls are guided not only by expected economic trade-offs (captured by the FCF) but also by the accumulated infeasibility information (captured by the FFF).

The dual multipliers obtained from this problem generate *optimality cuts* that refine the Future Cost Function, while the feasibility cuts remain active to define the feasible region. As the backward iterations progress, both sets of cuts are updated and expanded, progressively improving the joint representation of cost and feasibility across stages.

4 Algorithmic implementation

The Penalty-Free SDDP algorithm extends the standard forward–backward procedure by explicitly solving a feasibility problem before the optimality problem at each stage, obtaining the slack bounds s_t^* and β_t^* and enforcing them in the cost minimization. Convergence reflects the stabilization of both feasibility and cost approximations.

4.1 Forward pass

In the forward pass, a set of scenario paths Ω^{fwd} is sampled. For each path $\omega \in \Omega^{\text{fwd}}$ and for $t = 1, \dots, T$, we visit the corresponding conditional state and solve, in this order:

1. the *feasibility problem* (2) at $(x_{t-1}(\omega), \xi_{t-1}(\omega))$, obtaining the minimum slack vector $s_t^*(\omega)$ and the cumulative infeasibility $\beta_t^*(\omega)$;
2. the *optimality problem with feasibility cuts* (3) at the same state, now enforcing $s_t \leq s_t^*(\omega)$ and $\beta_t \leq \beta_t^*(\omega)$.

The forward pass therefore *directly obtains* the violations (via s_t^* and β_t^*) and *visits new states* $x_t(\omega)$ in the conditional scenario tree, building the set of trial states $\mathcal{S}_t := \mathcal{S}_t \cup \{x_t(\omega)\}$ for each stage t .

4.2 Backward pass

In the backward pass, the algorithm proceeds from $t = T$ down to $t = 1$. For each stage t and for every trial state $x_t \in \mathcal{S}_t$ generated in the forward pass, we solve:

- the feasibility problem (2) (at x_t and across the sampled realizations at $t+1$) obtaining the minimum slack vector s_t^* and the cumulative infeasibility β_t^* ; and, from its dual multipliers, *generate feasibility cuts* that update $\widetilde{\text{FFF}}_t$;

- the optimality problem (3) (with $s_t \leq s_t^*$ and $\beta \leq \beta_t^*$ inherited from the feasibility problem above) and, from its dual multipliers, *generate optimality cuts* that update $\widetilde{\text{FCF}}_t$.

Both sets of cuts are stored, refining the piecewise-linear approximations of the FCF and FFF at each stage. When no new feasibility cuts are produced in the backward pass, the feasible region is considered stable at the trial states visited.

4.3 Convergence criteria

Convergence is declared when the following conditions hold simultaneously:

1. **Feasibility stability:** no new feasibility cuts are generated during the backward pass;
2. **Optimality:** the standard optimality gap criterion is satisfied, i.e., $|Z_{\text{up}} - Z_{\text{low}}| \leq \varepsilon$.

This guarantees that no additional feasibility information can alter the policy at the visited states and that the cost approximation has converged within the prescribed tolerance.

4.4 Algorithm summary

The overall procedure of the Penalty-Free SDDP algorithm is summarized in Algorithm 1. The pseudocode highlights the alternation between feasibility and optimality steps, as well as the convergence check based on the stability of feasibility cuts and the optimality gap criterion.

Algorithm 1 Penalty-Free SDDP

```

1: Initialize: empty sets of optimality and feasibility cuts
2: repeat
3:    $\mathcal{S}_t \leftarrow \emptyset$ , for all  $t = 1:T$ 
4:   Forward pass:
5:   for each sampled path  $\omega \in \Omega^{\text{fwd}}$  do
6:     for  $t = 1$  to  $T$  do
7:       Solve (2) at  $(x_{t-1}(\omega), \xi_{t-1}(\omega))$ ; obtain  $s_t^*(\omega)$  and  $\beta_t^*$ 
8:       Solve (3) with  $0 \leq s_t \leq s_t^*(\omega)$  and  $\beta_t \leq \beta_t^*$ 
9:       Record trial state:  $\mathcal{S}_t \leftarrow \mathcal{S}_t \cup \{x_t(\omega)\}$ 
10:    end for
11:  end for
12:  Backward pass:
13:  for  $t = T$  down to 1 do
14:    for each  $x_t \in \mathcal{S}_t$  do
15:      for each  $\xi_t \in \Xi_t$  do
16:        Solve (2) at  $(x_t, \xi_t)$  to obtain  $s_t^*$  and  $\beta_t^*$  and add feasibility cuts to update  $\widetilde{\text{FFF}}_t$ 
17:        Solve (3) at  $(x_t, \xi_t)$  with  $0 \leq s_t \leq s_t^*$  and  $\beta_t \leq \beta_t^*$  and add optimality cuts to update  $\widetilde{\text{FCF}}_t$ 
18:      end for
19:    end for
20:  end for
21:  Update  $Z_{\text{up}}$  and  $Z_{\text{low}}$ 
22: until no new feasibility cuts added in backward pass and  $|Z_{\text{up}} - Z_{\text{low}}| \leq \varepsilon$ 
    
```

This implementation ensures that the recursive process stop only when no additional feasibility information can modify the state decisions and the cost approximation has converged within tolerance.

5 Case Study: Application to Energy Planning

The proposed methodology was applied to a large-scale planning model inspired by the operational challenges of the Brazilian hydrothermal system. This system comprises dozens of interconnected reservoirs with complex physical interdependencies, where operators must reconcile multiple and often conflicting objectives — such as maximizing energy production, ensuring environmental and navigational flows, and meeting strict regulatory and ecological requirements. These *hydraulic operating conditions* impose binding physical limits on the system and represent some

of the most demanding aspects of large-scale hydrothermal scheduling, particularly when inflows are uncertain and their satisfaction competes with economic efficiency.

In current operational models, these conditions are typically represented either as hard constraints or through large artificial penalty terms. While the latter approach guarantees numerical feasibility, it introduces severe calibration challenges and distorts the interpretation of marginal costs. The Penalty-Free SDDP formulation proposed in this work was designed precisely to overcome these limitations, ensuring compliance with physical and environmental requirements without contaminating the economic objective with arbitrary penalties.

To provide a controlled validation of the proposed methodology, we focus here on a deterministic formulation of the problem. This setup can be solved as a single optimization problem that serves as a consistent reference solution, allowing the specific effects of feasibility cuts to be isolated from the stochastic variability that would otherwise arise in a fully stochastic implementation. Future developments will naturally extend this analysis to stochastic formulations, where inflow and renewable-generation uncertainties can further amplify the relevance of proper feasibility handling.

The test system represents a long-term hydrothermal scheduling model comprising 75 interconnected reservoirs arranged in cascades, 109 thermal plants, and 32 aggregated renewable plants. The planning horizon spans 4.5 years with monthly resolution, plus an additional buffer year to mitigate the so-called “End of the World” effect.

Within the Brazilian hydrothermal context, these constraints define the physical, regulatory, and environmental operating conditions for hydropower plants. They include minimum total outflows, ecological flow requirements, navigability support and structural safety limits that ensure compliance with federal and state regulations. Altogether, they account for 662 *hard constraints* that embody nonnegotiable regulatory and ecological priorities within system operation.

To evaluate the performance of the proposed approach, three formulations were compared:

- (i) the traditional SDDP model using penalized slack variables to enforce HOCs;
- (ii) a single large-scale deterministic problem serving as the benchmark solution; and
- (iii) the proposed Penalty-Free SDDP model with feasibility cuts.

All formulations were solved under identical convergence criteria of 0.5%. The main computational results and performance comparisons are summarized in Table 1.

Table 1: Performance comparison of the three formulations.

Solution Method	Operation Cost [M\$]	Violation Cost [M\$]
Single problem	31,872.67	1.38
SDDP	31,877.09	2.59
Penalty-free SDDP	31,877.14	1.37

The Penalty-Free SDDP produced violations that were almost identical to those obtained in the single large-scale problem, affecting the same set of reservoirs, as shown in Figure 1. It correctly identified situations in which the strict enforcement of HOCs became infeasible under specific inflow conditions, reporting them explicitly through the feasibility function rather than masking them with large penalty terms. In contrast, the traditional SDDP showed violations at 14 reservoirs, compared to 11 in the single-problem benchmark, suggesting difficulty to calibrate the constraint and premature convergence. These results demonstrate that the use of feasibility cuts improves the interpretability and reliability of the solution, enabling the algorithm to distinguish between temporary and truly infeasible constraints without requiring iterative penalty tuning or repeated recalibration of model parameters.

Finally, regarding execution times, our proposed framework took 67 iterations to converge, requiring approximately $1.2\times$ more iterations than the traditional SDDP to achieve convergence (56 iterations), translating to proportionally longer execution times. However, this comparison must consider that the SDDP produced some distortions in violation magnitudes and locations. Moreover, our framework directly produces implementable solutions with certified feasibility. For applications where satisfying hard system constraints is critical and understanding which constraints cannot be met provides essential information, the additional computational cost could be justified and the trade-off would be acceptable.

6 Conclusions and Future Work

This paper presented the *Penalty-Free SDDP*, a general extension of the Stochastic Dual Dynamic Programming algorithm that eliminates the need for artificial penalty terms by incorporating feasibility cuts into its recursive structure.




Figure 1: Total minimum outflow violations.

The proposed formulation separates economic optimization from constraint enforcement through an auxiliary feasibility recursion, which allows the algorithm to handle infeasibilities automatically and to explicitly identify *truly infeasible* situations that cannot be corrected through adjustments in the state decisions.

The methodology was validated in a large-scale deterministic application inspired by the Brazilian hydrothermal system, where complex hydraulic, environmental, and regulatory conditions were modeled as hard constraints. The results demonstrated that the proposed approach preserves the economic consistency of the solution while correctly diagnosing infeasible operating conditions, without requiring iterative tuning or subjective calibration of penalty parameters — a process that is inherently case-dependent and often compromises model transparency.

A key advantage of this penalty-free scheme is its robustness and generality: by decoupling feasibility from cost minimization, it provides a principled way to ensure constraint satisfaction across different systems and formulations, while maintaining the interpretability of economic signals.

Future research will focus on extending the method to fully stochastic, multi-stage settings — the natural domain of SDDP — where uncertainty in inflows and renewable generation amplifies the relevance of proper feasibility handling. Such an extension will inevitably increase the number of cuts, as feasibility cuts are generated for every stage, scenario and realization in the backward pass. Although this poses a computational challenge, it also removes the need for problem-specific penalty calibration and enables more accurate and transparent optimization processes in which feasibility and economic optimality are treated as distinct, complementary components of the decision framework.

References

- [1] M. V. F. Pereira and L. M. V. G. Pinto, “Multi-stage stochastic optimization applied to energy planning,” *Mathematical Programming*, vol. 52, pp. 359–375, May 1991.
- [2] J. F. Benders, “Partitioning procedures for solving mixed-variables programming problems,” *Numerische Mathematik*, vol. 4, pp. 238–252, Dec 1962.
- [3] R. M. Van Slyke and R. Wets, “L-shaped linear programs with applications to optimal control and stochastic programming,” *SIAM Journal on Applied Mathematics*, vol. 17, no. 4, pp. 638–663, 1969.
- [4] T. N. Santos and A. L. Diniz, “Feasibility and optimality cuts for the multistage benders decomposition approach: Application to the network constrained hydrothermal scheduling,” in *2009 IEEE Power & Energy Society General Meeting*, pp. 1–8, 2009.

- [5] C. Yuen, A. Poncet, and X. Chao, “Handling infeasibilities when applying benders decomposition to scheduling optimization,” 01 2006.
- [6] V. Guigues, “Sddp for some interstage dependent risk-averse problems and application to hydro-thermal planning,” *Computational Optimization and Applications*, vol. 57, pp. 167–203, Jan 2014.
- [7] A. Helseth and B. Mo, “Hydropower aggregation by spatial decomposition—an sddp approach,” *IEEE Transactions on Sustainable Energy*, vol. 14, no. 1, pp. 381–392, 2023.
- [8] A. Helseth, “Approximating hydropower systems by feasibility spaces in stochastic dual dynamic programming,” *Electric Power Systems Research*, vol. 234, p. 110615, 2024.
- [9] V. Guigues, A. Shapiro, and Y. Cheng, “Duality and sensitivity analysis of multistage linear stochastic programs,” *European Journal of Operational Research*, vol. 308, no. 2, pp. 752–767, 2023.
- [10] A. L. Diniz, M. E. P. Maceira, C. L. V. Vasconcellos, and D. D. J. Penna, “A combined sddp/benders decomposition approach with a risk-averse surface concept for reservoir operation in long term power generation planning,” *Annals of Operations Research*, vol. 292, pp. 649–681, Sep 2020.
- [11] C. Füllner and S. Rebennack, “Stochastic dual dynamic programming and its variants: A review,” *SIAM Review*, vol. 67, no. 3, pp. 415–539, 2025.