

École polytechnique de Louvain

Imbalance Chasing in Belgian Energy Markets

Author : Ozkan TOPRAK

Supervisor : Prof. Francesco Contino

Readers : Prof. Quentin Leté, Prof. Hervé Jeanmart

Academic year 2025–2026

Master [120] in Mechanical Engineering

Copyright

© 2026 UCLouvain
All rights reserved.

No part of this thesis may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, or otherwise, without the prior written permission of UCLouvain.

To request permission to reproduce or use any part of this thesis, please contact UCLouvain at the following address:

Rue Archimède 1 — Bte L6.11.01
B-1348 Louvain-la-Neuve
(Parkings 11 or 18)
Tel: (10) 47 24 60

Use of Artificial Intelligence Assistance

AI-based tools were used to assist with writing clarity (rephrasing, grammar), structuring sections, and generating preliminary outlines. AI assistance was also used for Python code review during development; all analytical logic, threshold choices, and figure interpretation remain the author's own. No AI tool was used as a source of factual claims, and all technical content was verified against primary sources. All data handling, analysis, interpretation, and conclusions remain the sole responsibility of the author.

Abstract

The Belgian electricity balancing market settles Balance Responsible Party (BRP) imbalances at a single quarter-hourly price plus a scarcity adder (α) above 150 MW system imbalance, rewarding any BRP that takes an offsetting position—imbalance chasing. The 2024 PICASSO and 2025 MARI accessions completed cross-border reserve integration; whether they compressed the residual chasing incentive on Belgian data has not been measured. This work evaluates the incentive using a 244,773-bin master panel of Elia public feeds spanning 2019–2025, of which a 93,073-bin forecast-available subsample (2023–2025) carries the chasing analysis across three regulatory regimes—Pre-PICASSO, PICASSO Adapted, and MARI Adapted—using a transparent rule that operates on the published forecast system imbalance. Between Pre-PICASSO and MARI Adapted, the mean per-bin chasing reward fell by 47%, but the extreme-tail reward magnitude fell by only 11% and the loss rate rose from 20% to 35%. The cascade rate on chasing-active bins fell by 78%, while the over-representation of cascades in the extreme-chasing tail rose from $3.4\times$ to $9.8\times$. The two reforms compressed the broad chasing economy on every metric but concentrated the residual reward and grid stress in the same thin tail. A four-pillar decomposition isolates the single-pricing rule, the α threshold, the 15-minute settlement period, and the reserve product mix as the regulatory levers; a targeted policy combination compresses the extreme-tail reward by 65–92% across regimes (92% under MARI Adapted) without altering any pillar. The residual chasing incentive is concentrated in the extreme-chasing tail and can be compressed by an incremental policy instrument without redesigning the single-pricing rule.

Keywords: imbalance chasing; PICASSO; MARI; single-pricing settlement; policy counterfactual; tail-concentrated incentive; ex-ante imbalance forecast.

List of Acronyms and Notation

Acronyms

aFRR	Automatic Frequency Restoration Reserve
BRP	Balance Responsible Party
CREG	Commission for the Regulation of Electricity and Gas (Belgian energy regulator)
EBGL	Electricity Balancing Guideline (Commission Regulation (EU) 2017/2195)
EU	European Union
ISP	Imbalance Settlement Period
LFC	Load-Frequency Control
MARI	European platform for mFRR balancing energy
mFRR	Manual Frequency Restoration Reserve
MPC	Model Predictive Control
NIV	Net Imbalance Volume
PICASSO	European platform for aFRR balancing energy
PV	Photovoltaic
SI	System Imbalance
TSO	Transmission System Operator

Notation

SI_t	System imbalance at quarter-hour t (MW); negative = short, positive = long
$\hat{SI}_t$	Ex-ante published forecast of SI_t , available before bin t starts (MW)
P_t^{imb}	Imbalance settlement price at quarter-hour t (EUR/MWh)
P_{t+1}^{imb}	Realised imbalance price of the next 15-minute bin, used to settle a chasing position taken in bin t (EUR/MWh)
α	Scarcity adder applied to P_t^{imb} when $ SI_t > 150$ MW (EUR/MWh)
π_t	Per-bin chasing reward on a chasing-active bin (EUR/MWh)
T	Counterfactual activation threshold for the scarcity adder α used in the sensitivity analysis (MW)
$[p_a, p_b]$	Within-regime quantile window between the a -th and b -th percentile of the per-regime realised next-bin imbalance price distribution

Contents

List of Acronyms and Notation	i
1 Introduction	1
1.1 Problem Statement	1
1.2 Research Questions	1
2 Background and Literature Gap	3
2.1 Belgian imbalance settlement	3
2.2 Imbalance chasing literature	3
2.3 Literature gap	4
3 Data and Methodology	6
3.1 Data and sample	6
3.2 Key definitions	7
3.3 Estimation and robustness	9
4 Results	10
4.1 Profitability shift across regimes	10
4.1.1 Mean reward across regimes	10
4.1.2 Central reward versus tail reward	10
4.1.3 Loss rate across regimes	11
4.1.4 Robustness and scope	11
4.2 Grid stability response	13
4.2.1 Cascade rate and the chasing-versus-idle gap	13
4.2.2 Reserve activation composition	13
4.2.3 Cascade concentration in the extreme-chasing tail	14
4.2.4 Robustness and scope	15
4.3 Regulatory attribution	16
4.3.1 Pricing rule: Short–Long spread under single pricing	16
4.3.2 The α -component activation threshold	17
4.3.3 Settlement resolution: 15-minute versus hourly aggregation	18
4.3.4 Balancing product composition	19
4.3.5 Four-pillar synthesis and scope	19
4.4 Policy reform potential	21
4.4.1 P1: within-regime price kink	21
4.4.2 P2: raised α threshold	22
4.4.3 P3: manual-reserve activation fee	23
4.4.4 P1 + P3 combined, and scope	23

5	Discussion	24
5.1	Cross-cutting symmetry	24
5.2	Cross-border attribution	25
5.3	Cross-jurisdictional context and regulatory implications	25
5.4	Limitations	26
6	Conclusion and Future Work	28
6.1	Conclusion	28
6.2	Future directions	29
	References	30
A	Additional Robustness Checks	33
A.1	Quality-aware restriction	33
A.2	mFRR activation distribution	33
A.3	Policy counterfactual results	34
A.4	α -threshold counterfactual sweep	35

List of Figures

2.1	Belgian quarter-hourly imbalance settlement chain.	4
3.1	Bin-level coverage of the master panel variables over the 2019–2025 sample, by calendar year. Each cell gives the share of 15-minute bins in which the corresponding variable is non-null and carries a validation flag compatible with the quality-aware restriction (Section 3.3).	7
4.1	Mean reward per chasing-active bin by regulatory regime	11
4.2	Reward distribution on chasing-active bins by regulatory regime	12
4.3	Loss rate by regulatory regime	12
4.4	Cascade rate on chasing-active vs idle bins, by regulatory regime	14
4.5	Mean aFRR and mFRR activation on chasing-active vs idle bins, by regulatory regime	15
4.6	Cascade rate on any chasing-active bin vs the extreme-chasing tail, by regulatory regime	16
4.7	Imbalance price by SI direction and regulatory regime	17
4.8	α -threshold sensitivity sweep across six values, by regulatory regime	18
4.9	Settlement-resolution comparison: 15-minute vs hypothetical hourly aggregation	19
4.10	Reserve activation composition (aFRR vs mFRR) in chasing-active bins, by regulatory regime	20
4.11	Four-metric regulatory scorecard on chasing-active bins (p95 cascade threshold)	21
4.12	Policy counterfactuals on the extreme-chasing tail, by regulatory regime	22

List of Tables

3.1	The three regulatory regimes used in the empirical analysis. Bin counts refer to the chasing-analysis sample (forecast SI non-null, within the 2023–2025 window).	7
A.1	Policy counterfactuals on the extreme-chasing tail, by regulatory regime	34
A.2	Full α -threshold counterfactual sweep, by regulatory regime	35

Chapter 1: Introduction

On 7 April 2024 the Belgian imbalance settlement price reached $-4,574$ EUR/MWh, the lowest in Elia’s published record [1], driven by a system surplus during a Saharan-dust solar shortfall [2]. The magnitude reflects a structural feature of Belgium’s single-price settlement: when the system is unexpectedly long, the price punishes any party adding to the surplus and rewards any party offsetting it. The same mechanism that exposes Balance Responsible Parties (BRPs) to extreme losses creates an incentive to deliberately take an offsetting position — imbalance chasing, characterised analytically by Bottieau *et al.* [3].

A BRP that chases bets on the sign of the system imbalance. It uses public information available before the settlement quarter, takes a position opposite to its forecast, and is settled at the realised next-quarter imbalance price. The strategy works in two directions. When the system behaves as expected, a chasing position offsets the imbalance the Transmission System Operator (TSO) is working to clear; when the system is under stress, the same position can add to that imbalance and push grid stress into a few narrow windows. Whether a given chasing rule, on net, supports the grid or stresses it is an empirical question.

1.1 Problem Statement

The Belgian market design changed twice in a six-month window. On 27 November 2024 Belgium acceded to PICASSO, the European platform for cross-border activation of automatic Frequency Restoration Reserve (aFRR) [4]; on 21 May 2025, to MARI, the equivalent platform for manual reserve (mFRR) [5]. Both operate under the EU Electricity Balancing Guideline [6] and are expected to compress imbalance price spreads by opening domestic reserve activation to a wider European bid pool.

The existing literature documents the equilibrium structure of single-pricing balancing markets [7], [8], the design intent of the European integration platforms [9], and the strategic behaviour they invite [10]. What remains unmeasured on Belgian data is whether the chasing incentive has persisted, weakened, or shifted in shape after the two accessions, and whether any residual reward is now concentrated in the same bins where the system is under stress.

1.2 Research Questions

This work addresses the four research questions:

- RQ1.** How do NIV chasing strategies increase operational profits of market participants?
- RQ2.** How does NIV chasing affect grid stability?
- RQ3.** What are the existing regulatory mechanisms in Belgium to manage imbalance pricing, and how effective are they?
- RQ4.** What potential regulatory changes or market mechanisms could be implemented to mitigate harmful imbalance chasing while preserving incentives for flexibility?

All four are evaluated on a three-year (2023–2025) fifteen-minute sample of Belgian imbalance data, across the Pre-PICASSO (Jan 2023–26 Nov 2024), PICASSO Adapted (27 Nov 2024–20 May 2025), and MARI Adapted (21 May 2025–31 Dec 2025) regimes.¹ The analysis is descriptive throughout: regime contrasts document co-occurrence patterns rather than establish causal attribution (Section 3.3).

The contribution is threefold. First, this work defines a transparent chasing rule built only on information a real BRP holds at the moment of decision: $|\hat{S}I_t| > 150$ MW on the published forecast system imbalance, with settlement at the next-quarter realised imbalance price. Second, the analysis documents an asymmetry that resists explanation by any single mechanism. Between Pre-PICASSO and MARI Adapted, the broad chasing economy compressed on every metric. The within-regime extreme-chasing tail did the opposite: it held the residual reward and the residual grid stress in the same thin subset of bins. Third, this residual turns out to be mechanically compressible from within the existing settlement design, without redesigning the single-pricing rule.

The remainder is structured as follows. Chapter 2 reviews the Belgian balancing framework and the chasing literature. Chapter 3 documents the data, key definitions, three-regime decomposition, and counterfactual setup. Sections 4.1 to 4.4 report findings on profitability (RQ1), grid stability (RQ2), regulatory attribution (RQ3), and policy reform (RQ4). Chapter 5 synthesises across the four questions and states the structural limitations, and Chapter 6 concludes with future directions.

¹PICASSO accession occurred on 26 November 2024 at quarter-hour 11:15–11:30 [4]; the regime cutoff is set to the first full operational day (27 November) to avoid mixing partial-coverage hours with regime samples. The MARI accession on 21 May 2025 took effect at the start of day, so the corresponding regime cutoff coincides with the go-live date [5].

Chapter 2: Background and Literature Gap

2.1 Belgian imbalance settlement

The Belgian balancing market operates under the EU Electricity Balancing Guideline (EBGL), which established marginal pricing for activated balancing energy, common merit orders, and platform-based cross-border activation [6]. Belgium acceded to PICASSO — the European platform for automatic Frequency Restoration Reserve (aFRR) — on 27 November 2024, and to MARI — the equivalent platform for manual reserve (mFRR) — on 21 May 2025 [4], [5], [11], completing the transition for both reserve products within a six-month window. Backer *et al.* [9] estimate that platform integration compresses cross-border activation costs and reduces price spikes by widening the bid pool, but the same coupling makes domestic imbalance prices more sensitive to neighbouring control areas. Poplavskaya *et al.* [10] show, using agent-based simulation, that strategic bidders retain influence over clearing outcomes when a few of large participants dominate the market, even after design improvements; the corresponding test on Belgian post-accession data has not been performed.

Belgium has settled BRP imbalances at a single quarter-hourly price for over a decade, and since the 2020 tariff revision that price equals the marginal cost of activated reserve plus an α scarcity adder that triggers when the absolute system imbalance exceeds 150 MW [12].¹ Three features jointly produce the chasing incentive. Single pricing means deficit-direction (Short) and surplus-direction (Long) BRPs face the same imbalance price, so the median Short–Long price spread within a regime equals the per-bin chasing reward; Baetens *et al.* [13] formalise this for industrial self-balancers on pre-reform data. The 15-minute Imbalance Settlement Period (ISP) sets the resolution at which deviation and price are aligned, and van der Veen *et al.* [8] and van der Veen and Hakvoort [14] show in agent-based simulation that settlement period length and pricing rule jointly determine BRP risk-taking. The α component is itself a regulatory lever, calibrated as a transposition of the Operating Reserve Demand Curve to the Belgian merit-order setting [15], [16]; Cartuyvels *et al.* [17] show analytically that imbalance-side adders introduced without simultaneous trading of balancing capacity distort the merit order and tend to reduce welfare for the introducing member state — a prediction directly applicable to Belgium’s α design after PICASSO and MARI accession. Figure 2.1 summarises how these three features link the realised system imbalance to the per-bin chasing reward. The supply-side response is documented in the cited equilibrium and platform-integration literature; the demand-side response on Belgian data — whether the BRP incentive to deviate has weakened, persisted, or shifted into a tail — has not been measured.

2.2 Imbalance chasing literature

Bottieau *et al.* [3] provide the canonical formalisation. A BRP forecasts the next-quarter system imbalance using public information, takes an offsetting portfolio position, and is settled at the

¹The pre-2020 activation threshold was 140 MW; Bottieau *et al.* [3] use that earlier value. The 150 MW threshold has been in force throughout the 2023–2025 sample analysed here.




Figure 2.1: Belgian quarter-hourly imbalance settlement chain. The system imbalance SI_t triggers reserve activation via PICASSO (aFRR) and MARI (mFRR). The imbalance price P_t^{imb} equals the marginal reserve cost plus an α adder when $|SI_t| > 150 \text{ MW}$. A BRP that takes an offsetting position is settled at P_t^{imb} ; the spread against its day-ahead sourcing cost is the per-bin chasing reward.

realised imbalance price; the strategy is profitable in expectation when the forecast direction is correct and the realised price is sufficiently spread from the day-ahead price. Their analysis predates both the 2020 tariff revision and the PICASSO/MARI accessions, leaving the post-reform economics of this strategy untested on Belgian data. Eicke *et al.* [7] place the same behaviour in equilibrium terms: the TSO supplies imbalance energy by activating reserves, while BRPs — responding to the price signal — supply an elastic demand for imbalance, so the system imbalance and the imbalance price are jointly determined rather than sequentially fixed. Empirically, the strategy has been profitable on pre-reform Belgian data: Smets *et al.* [18] implement stochastic MPC on battery storage taking deliberate out-of-balance positions; system-supporting deviations occur in approximately three-quarters of all activated bins, while Wessel *et al.* [19] extend the same setup with explicit risk constraints. Both studies isolate a single asset class (battery storage) and a single agent rather than the system-level chasing incentive evaluated here. A parallel methodological strand delivers forecasting tools: Dumas *et al.* [20] model the Belgian Net Regulation Volume as a Markov chain, Bruneel *et al.* [21] model the imbalance price as a bimodal distribution and construct a risk-averse mixture-model strategy, and Pavirani *et al.* [22] use Monte Carlo Tree Search to publish ex-ante prices accounting for the implicit response of other market participants. All three deliver participation tooling without measuring how the residual chasing incentive itself responds to a market design change. The combined picture is a strategy that is well defined, supported by mature forecasting tools, and demonstrably profitable on data predating the PICASSO and MARI accessions.

2.3 Literature gap

Vermeulen [23] provides the most direct cross-jurisdictional comparator: an evaluation of single (Belgian) versus dual (Dutch) pricing finds single pricing more stable on average and lower in BRP-side risk, dual pricing more dampening of overreactive balancing but more volatile in market outcomes, and both schemes vulnerable to strategic gaming. The result defends Belgium’s pricing-rule choice on equilibrium grounds, but the Belgian sample analysed ends before MARI accession and therefore cannot speak to whether the residual chasing incentive after the May 2025

reform is grid-supporting or grid-stressing. Concurrent theoretical work by Karimi Madahi *et al.* [24] formalises a bi-level optimisation showing that, under perfect foresight, BRPs operating fast-response battery storage can in principle exploit the sub-quarter-hourly dynamics of single-pricing settlements in both Belgium and the Netherlands; the result establishes that the EBGL design unintentionally permits gaming but leaves open the empirical question of how much residual reward an ex-ante implementable rule actually captures on observed post-reform data. A complementary direction is taken by Bruninx *et al.* [25], who develop a minute-level market simulator of Belgian implicit balancing and report that, under the current imbalance price formula, an additional implicit-response capacity of around 200 MW reduces balancing costs while larger capacities trigger overshoots that propagate into the imbalance price; the simulation is run on 2023 data and therefore predates the PICASSO and MARI accessions, identifying a prospective overshoot risk that the empirical work here measures retrospectively on observed post-reform data. The same boundary applies to the regulatory literature [6], [12], [15], which sets out design intent without measuring behavioural response, and to the forecasting and trading literature [3], [20]–[22], which provides participation tooling without comparing pre- and post-reform regimes on equal terms. Even the analytical equilibrium literature on cross-border platforms [17], [26] is built on stylised demand and supply functions rather than empirical BRP behaviour. Four questions therefore remain unmeasured on Belgian data: whether chasing remains profitable after the reforms (RQ1), whether chasing-active periods coincide with elevated reserve activation (RQ2), which feature of the regulatory framework drives any observed shift (RQ3), and which incremental policy instrument could compress the residual chasing incentive without redesigning the single-pricing rule (RQ4). These questions are answered in the remainder of this work on the sample and regimes defined in [Chapter 3](#).

Chapter 3: Data and Methodology

3.1 Data and sample

The empirical analysis rests on a single 15-minute panel combining six Elia public datasets at the Imbalance Settlement Period (ISP), the temporal resolution of Belgian BRP imbalance settlement [27]: the imbalance settlement price, the realised system imbalance (SI), the published SI forecast, automatic and manual reserve activation volumes (aFRR and mFRR), Elia’s grid load, and the measured and forecast load. All series are indexed in Europe/Brussels local time and form a 15-minute master panel of 244,773 bins covering 1 January 2019 to 31 December 2025. The analytical cutoff is enforced at 31 December 2025, 23:59:59 Brussels time; partial 2026 observations are dropped to keep every regime-level statistic on a closed sample.

The chasing-analysis window covers 2023–2025 and contains 105,216 quarter-hourly bins, of which 93,073 carry a non-null published forecast; the remaining bins fall inside two extended Elia publication gaps, one in August–October 2023 and one in September–November 2025. Forecast-dependent figures and statistics throughout this work are computed on this window. The full 2019–2025 master panel of 244,773 bins is kept alongside it for data-scope context and is shown in the coverage heatmap (Figure 3.1). Every analytical comparison across the four research questions runs on the chasing-analysis window.

Two Elia data files changed structure on 22 May 2024 and were harmonised across the 2019–2025 sample. The imbalance settlement price was split into separate positive- and negative-imbalance files; the two new files are identical at 15-minute resolution to the legacy single-price file over the overlap period, so either is taken as the unified series. The reserve activation schema was restructured in May 2024 to split manual-reserve activations into Scheduled (SA) and Direct (DA) categories; both formats are mapped into a unified set of upward/downward, net, and absolute series for each of aFRR and mFRR, plus a total balancing-volume scalar used as the per-bin reserve-pressure indicator in Section 4.2.¹

The SI forecast requires a separate fix. The SI forecast is the latest of the six series to enter the public data feed: Elia began publishing it in January 2023, which fixes the start of the chasing-analysis window in Section 3.2 as a data-availability boundary rather than a free parameter. The Elia file timestamp marks the moment of *publication*, which can fall after the start of the quarter the forecast refers to; loading the file naively would introduce look-ahead. The ex-ante signal is therefore reconstructed by assigning to each 15-minute bin only forecast records published *strictly before* the start of that bin. The directional convention follows Elia’s settlement formulas: a negative system imbalance ($SI < 0$) means the system is short and the imbalance price is high; a positive imbalance ($SI > 0$) means the system is long and the price is low. A BRP that takes a Long position (its portfolio in surplus) when the system is short is therefore “offsetting” and is rewarded; the same Long position when the system is long is “amplifying” and is penalised.

A quality note applies to three variable families. Elia’s load and load-forecast files carry

¹The legacy schema (aFRR+, aFRR-, mFRR Standard, mFRR Flex) was replaced in May 2024 by a refined schema introducing a Scheduled (SA) versus Direct (DA) activation split on the manual reserve (aFRR BE +/-, mFRR BE SA +/-, mFRR BE DA +/-). The mapping preserves bin-level total activation magnitudes.

no validation flag at all, so every bin in those two series shows up under an “unknown” label by construction; the underlying measurements are nevertheless present and are used at face value. The activation files mix labelled and unlabelled records: 46 % of bins are validated, 54 % unlabelled, with no systematic regime pattern in the split. Imbalance prices and the realised SI behave differently. Both carry a validation flag on essentially every bin in the sample, so the quality-aware restriction (Section 3.3) does not bind on price-driven contrasts and converges to the headline result. Per-variable and per-year coverage and validation patterns are summarised in Figure 3.1.




Figure 3.1: Bin-level coverage of the master panel variables over the 2019–2025 sample, by calendar year. Each cell gives the share of 15-minute bins in which the corresponding variable is non-null and carries a validation flag compatible with the quality-aware restriction (Section 3.3).

The 2023–2025 chasing-analysis sample is partitioned into three regulatory regimes defined by the European integration go-live dates (Table 3.1). The PICASSO and MARI breakpoints are the dates on which Belgium’s Load-Frequency Control area joined the corresponding European balancing platforms [4], [5]. The 22 May 2024 publication-format change falls inside the Pre-PICASSO window: it modifies how the data is reported, not how the market is settled. The Pre-PICASSO baseline contains roughly $3.7\times$ as many bins as either post-reform regime, which widens confidence intervals on those regimes; every per-regime estimate is therefore reported with a 95 % confidence interval from the day-block bootstrap described in Section 3.3.

Table 3.1: The three regulatory regimes used in the empirical analysis. Bin counts refer to the chasing-analysis sample (forecast SI non-null, within the 2023–2025 window).

Regime	Start	End	Bins
Pre-PICASSO	2023-01-01	2024-11-26	60,431
PICASSO Adapted	2024-11-27	2025-05-20	16,617
MARI Adapted	2025-05-21	2025-12-31	16,025
Total			93,073

3.2 Key definitions

The empirical analysis rests on three operational definitions.

Chasing-active bin. A bin is *chasing-active* when the absolute value of the published forecast SI exceeds absolute 150 MW. The threshold is the activation level of Belgium’s α scarcity adder under the 2020 tariff revision [12], [15]: when $|\text{SI}_t| > 150$ MW the imbalance price is augmented by an α component, and a BRP that takes an offsetting position is settled at the augmented price. The threshold is therefore the regulatory boundary at which the marginal chasing reward changes regime, not a free parameter.

The per-bin reward on each chasing-active bin is

$$\pi_t = -\text{sign}(\widehat{\text{SI}}_t) \cdot P_{t+1}^{\text{imb}}, \quad (3.1)$$

where $\widehat{\text{SI}}_t$ is the ex-ante forecast available before bin t starts and P_{t+1}^{imb} is the realised imbalance price of the next 15-minute bin. Settling at the next bin rather than the same bin removes look-ahead from the rule. Rewards are reported in EUR/MWh per unit of position size, with a positive sign when the offsetting direction meets a favourable price and a negative sign otherwise. From this per-bin reward, three summary statistics carry through the rest of the work: the *loss rate* is the share of chasing-active bins on which $\pi_t < 0$; the *mean per-bin reward* is the average of π_t over the chasing-active sample; and the *extreme-chasing tail* is the slice of chasing-active bins whose realised next-bin price falls outside the within-regime $[p_1, p_{99}]$ window. Equation (3.1) is a settlement-cashflow proxy rather than a measurement of realised BRP profit. It captures the marginal deviation reward and nothing else — day-ahead sourcing cost, intraday execution, ramp constraints, and portfolio composition are all absent from the formula. What this work reports is therefore the chasing incentive, not the realised chasing income.

Cascade event. A bin is flagged as a *cascade event* when two conditions hold together: the aFRR activation magnitude is above the full-sample p95 threshold of $\text{aFRR}_{\text{p95}} = 94.6$ MW and the mFRR activation is strictly positive. The two thresholds are deliberately asymmetric because the two reserves behave differently. aFRR is dispatched continuously and its 15-minute activation magnitude is approximately log-normal, so a tail percentile yields a meaningful high-magnitude indicator. mFRR, by contrast, is event-driven: 79 % of bins carry zero activation, the p95 of mFRR is itself zero, and any positive value is already the rare-event indicator the analysis needs. The combined definition therefore picks out the operationally relevant case — high automatic-reserve pressure that has escalated to manual dispatch within the same quarter. Both thresholds are computed once on the full 2023–2025 sample and applied identically across the three regimes, so the cascade indicator is regime-comparable by construction. A tighter p99 cut is reported as a sensitivity check in Section 4.2.

Counterfactual policy instruments. The policy evaluation in Section 4.4 compares the realised extreme-chasing reward against three single-instrument counterfactuals applied to the same residual tail. P1 (price kink) introduces a within-regime price kink: realised next-bin imbalance prices are clipped to the within-regime $[p_5, p_{95}]$ window, simulating a two-sided regulatory bound on settlement-price drift inside the α -active region. P2 (raised α threshold) raises the α activation threshold from 150 MW to 250 MW, admitting only forecasts with stronger imbalance magnitude. P3 (mFRR fee) imposes a flat 100 EUR/MWh activation fee on every

bin in which mFRR is dispatched. The combined P1+P3 instrument applies both treatments simultaneously. Each counterfactual is computed on the same chasing-active bins as the baseline; no behavioural-response model is layered on top — the counterfactuals re-price the existing settlement record under the alternative rule and report the resulting reward for an unchanged BRP position, consistent with the descriptive framing applied throughout.

3.3 Estimation and robustness

Every regime-level point estimate is reported with a 95 % confidence interval from a day-block bootstrap. The bootstrap procedure constructs the interval by repeatedly resampling the observed data with replacement and recomputing the statistic on each resampled set; the resulting empirical distribution gives the interval without assuming a parametric distributional form. Day-blocking groups bins within the same calendar day so each resample preserves within-day autocorrelation in the price and SI series (1,000 resamples, fixed seed), giving wider and more conservative intervals than a non-blocked bootstrap that resamples individual bins as if they were independent and identically distributed. A quality-aware robustness check then restricts the chasing-active sample to bins where every input variable carries a “validated” quality flag. Because the imbalance price and aFRR activation series carry full validation flags throughout 2019–2025, this restriction binds primarily through the SI quality flag. The headline regime contrasts are reported on the full chasing-active sample; the restricted sample is referenced in Appendix A as a sensitivity layer.

The kernel density panels in Section 4.1.2 use a Gaussian estimator with a bandwidth factor of 0.18 under the SciPy convention, which is broadly comparable to Silverman’s rule of thumb [28] at the regime sample sizes of Table 3.1. Sweeping the bandwidth across the 0.10–0.30 range leaves the qualitative regime contrasts unchanged.

One framing choice runs through the rest of the work. The three-regime contrasts and the chasing–cascade associations are conditional, ex-post comparisons on a 2023–2025 sample; they describe co-occurrence patterns and do not deliver causal estimates of the PICASSO and MARI reforms. The point is made here once and not repeated in every figure caption.

Chapter 4: Results

4.1 Profitability shift across regimes

The chasing rule of [Section 3.2](#) — $|\widehat{\text{SI}}_t| > 150$ MW on the published forecast SI, settled at the realised next-bin price — yields a chasing-active sample of 22,741 bins over 2023–2025, 21.6% of the 105,216-bin chasing-analysis window (24.4% of the 93,073 bins with non-null forecast). The reward defined in [Equation \(3.1\)](#) is computed for each chasing-active bin and aggregated within the three regimes of [Table 3.1](#).

4.1.1 Mean reward across regimes

Mean reward per chasing-active bin fell from +139 EUR/MWh in Pre-PICASSO to +92 under PICASSO Adapted and +74 under MARI Adapted — a 47% reduction across the two reforms ([Figure 4.1](#)). The shift is monotone, the day-block bootstrap 95% confidence intervals do not overlap between Pre-PICASSO and either post-reform regime; the PICASSO Adapted and MARI Adapted intervals overlap at their boundaries, consistent with their proximity in mean reward. The regime ranking is preserved under the quality-aware restriction of [Section 3.3](#). The activation rate is stable across the first reform and falls under the second: 21.4 chasing-active bins per day in Pre-PICASSO, 22.0 under PICASSO Adapted, 18.0 under MARI Adapted — a 16% reduction relative to Pre-PICASSO. The chasing rule fires roughly as often after PICASSO but materially less often after MARI, while the reward per fire is smaller throughout. The MARI-Adapted reduction is consistent with cross-border activation moderating realised system-imbalance excursions, but is not large enough to drive the 47% reward compression on its own.

4.1.2 Central reward versus tail reward

The headline shift hides a sharp asymmetry between the centre and the tail of the reward distribution ([Figure 4.2](#)). Three measurements separate them. First, the centre shifted left: the median reward fell from +163 EUR/MWh in Pre-PICASSO to +70 under PICASSO Adapted and +55 under MARI Adapted, broadly proportional to the mean. Second, the conditional absolute reward magnitude on the symmetric within-regime extreme-chasing tail (next-bin price beyond the within-regime $[p_1, p_{99}]$ window) moved from +577 EUR/MWh in Pre-PICASSO to +538 under PICASSO Adapted and +511 under MARI Adapted — an 11% reduction against the 47% reduction in the broad mean. The metric is reported as the mean $|\pi_t|$ on extreme-tail bins because the tail is symmetric across upper and lower price excursions, and aggregating the signed reward would conflate rare wrong-way realisations with the symmetric reward magnitude. Third, the rarest profitable events grew rather than shrank: the within-regime $p_{99.9}$ reward rises from +1,520 EUR/MWh in Pre-PICASSO to +1,877 under MARI Adapted, a 23% increase. Excess kurtosis — a measure of tail heaviness, zero for the normal distribution — follows the same shape: +34.2, +12.6, +41.7 across the three regimes; by that measure MARI is the most extreme-event-dominated of the three. The two reforms compressed the distribution toward zero in the centre but left its right tail substantially intact; the remaining chasing reward is




Figure 4.1: Mean reward per chasing-active bin by regulatory regime. Error bars are 95% confidence intervals of the mean from a calendar-day-block bootstrap (1,000 resamples, fixed seed). Sample sizes are bins with $|\widehat{SI}_t| > 150$ MW and a non-null next-bin imbalance price. The reward is a settlement-cashflow proxy as defined in Equation (3.1); see caveat in Section 3.2.

concentrated in fewer, larger, more extreme-price-driven events.

4.1.3 Loss rate across regimes

The loss rate — the share of chasing-active bins on which $\pi_t < 0$ — rose sharply across the first reform and stayed elevated under the second (Figure 4.3). The Pre-PICASSO baseline produced a loss on 20.0% of chasing-active bins; PICASSO Adapted on 36.0%; MARI Adapted on 34.5%. Wilson 95% confidence intervals are tight at every point (± 1.5 pp at the smallest sample and ± 0.7 pp on the Pre-PICASSO baseline), and the contrast between Pre-PICASSO and either post-reform regime is statistically distinguishable. The loss-rate movement and the mean-reward movement do not describe the same thing twice. They describe two faces of the same underlying change: PICASSO and MARI did not just compress the average reward, they also made the per-bin outcome materially more dispersed. A BRP that applies the same chasing rule under MARI Adapted accepts a loss on roughly one bin in three rather than one in five, and collects a smaller positive expectation on the bins that pay.

4.1.4 Robustness and scope

The day-block confidence intervals in Figure 4.1 separate Pre-PICASSO from both post-reform regimes; the PICASSO-to-MARI contrast lies within the overlapping range of the two regimes' intervals. No pairwise contrast crosses zero. Aggregating the same chasing-active bins by calendar year produces +119 EUR/MWh in 2023 (Pre-PICASSO only), +153 in 2024 (Pre-PICASSO mixed with five weeks of PICASSO Adapted), and +81 in 2025 (PICASSO Adapted mixed with




Figure 4.2: Reward distribution on chasing-active bins by regulatory regime. Panel A: Gaussian kernel density — a smooth empirical estimate of the per-bin reward distribution — with bandwidth factor 0.18 (Section 3.3); vertical dotted lines mark the within-regime p_1 and p_{99} quantiles of the per-bin reward. Panel B: top reward percentiles (p_{95} , p_{99} , $p_{99.5}$, $p_{99.9}$) within each regime, in EUR/MWh. Restricted to chasing-active bins ($|\hat{S}I_t| > 150$ MW). Conditional tail magnitude is computed as $\text{mean}(|\pi_t|)$ on bins where the realised next-bin price falls outside the within-regime $[p_1, p_{99}]$ window.




Figure 4.3: Loss rate by regulatory regime. A wrong-way bin is one for which $|\hat{S}I_t| > 150$ MW and $\pi_t < 0$. Sample sizes appear below each bar.

MARI Adapted), with loss rates 23.2%, 17.8%, and 35.9% respectively. The year-to-year spread (72 EUR/MWh in mean reward, 18.1 pp in loss rate) is comparable to the regime-level spread (65 EUR/MWh, 16.0 pp), confirming that the regime decomposition does not collapse onto a single calendar-time shock. The 2024 mean reward exceeds the Pre-PICASSO regime average because that calendar year concentrates the most extreme system-imbalance events of the sample, including the 7 April 2024 record imbalance price referenced in [Chapter 1](#), which inflate the conditional reward on the Pre-PICASSO bins without altering the settlement rule that produced them. The quality-aware restriction of [Section 3.3](#) moves the regime-level mean reward and loss rate by less than 5% in either direction, and the regime ranking is preserved.

Three caveats qualify the result. First, the reward of [Equation \(3.1\)](#) is a settlement-cashflow proxy ([Section 3.2](#)), not a measurement of realised BRP profit. Second, the regime decomposition is a conditional ex-post comparison: on a 2023–2025 sample, the two reforms cannot be cleanly separated from any correlated calendar-time shifts in load, weather, and renewable mix, and [Chapter 5](#) returns to this confound. Third, the per-bin reward says nothing about whether chasing was system-supporting or system-stressing in the moment of activation, and [Section 4.2](#) takes that question up on the same chasing-active sample.

4.2 Grid stability response

The actioned-bin sample of [Section 4.1](#) — 22,741 chasing-active bins partitioned across the three regimes of [Table 3.1](#) — is contrasted here with the 70,332 idle bins ($|\hat{S}_t| \leq 150$ MW) in the same window. The cascade indicator of [Section 3.2](#) (aFRR > 94.6 MW and any non-zero mFRR activation) yields a binary stress flag computed identically across regimes.

4.2.1 Cascade rate and the chasing-versus-idle gap

Cascade rate on chasing-active bins fell from 7.70% in Pre-PICASSO to 3.50% under PICASSO Adapted and to 1.71% under MARI Adapted, a 78% reduction over the two reforms ([Figure 4.4](#)). The shift is monotone, and it tracks the broad chasing-reward compression documented in [Section 4.1.1](#): both metrics move in the same direction on the broad chasing population. Cascade rates on idle bins fell faster in relative terms ($0.80\% \rightarrow 0.36\% \rightarrow 0.084\%$), and the chasing-to-idle multiplier therefore widened from $9.6\times$ in Pre-PICASSO to $9.7\times$ under PICASSO Adapted and $20.4\times$ under MARI Adapted. After MARI, the grid stress that remained in the system was concentrated in chasing-active bins at more than twice the pre-reform multiplier. Chasing continued to act as a strong discriminator of cascade risk: the reforms compressed the average grid-stress rate but did not weaken the link between chasing and grid stress.

4.2.2 Reserve activation composition

The cascade-rate compression masks a composition shift in the underlying reserve activation. Mean aFRR activation on chasing-active bins fell monotonically from 42.5 MW in Pre-PICASSO to 25.4 MW under PICASSO Adapted and 17.4 MW under MARI Adapted — a 59% reduction consistent with the design intent of cross-border aFRR activation through PICASSO. Mean mFRR activation, in contrast, rose from 74.0 MW in Pre-PICASSO to 89.6 MW under PICASSO




Figure 4.4: Cascade rate (aFRR > 94.6 MW and any mFRR activation) on chasing-active bins ($|\widehat{SI}_t| > 150$ MW) versus idle bins, by regulatory regime. The p95 aFRR threshold is computed once on the full 2023–2025 sample. Multipliers ($N\times$) above each pair report the chasing-to-idle cascade-rate ratio; bar labels are rounded to one decimal, which is why the MARI-Adapted idle bar displays 0.1% while the underlying rate of 0.084% yields the 20.4 $\times$ multiplier reported here. mFRR is zero in roughly 79% of bins, so “any mFRR activation” captures the rare-event tail of manual dispatch.

Adapted (+21%) before falling sharply to 34.6 MW under MARI Adapted (−61% vs. PICASSO Adapted; −53% vs. Pre-PICASSO; Figure 4.5). The PICASSO-period rise in mFRR is consistent with a substitution effect on the manual reserve. Cross-border merit-order coupling constrained automatic reserve activation, and the Belgian system responded by activating more mFRR on the same chasing-active bins, clearing imbalance volume that would previously have been covered through aFRR. MARI accession in May 2025 then opened mFRR activation to the same European pool, and from that point both reserves contracted together.

The chasing-versus-idle activation gap holds across all three regimes and on both reserve products. Let the *chasing premium* denote the chasing-to-idle ratio of mean activation (multipliers shown above each pair in Figure 4.5). On aFRR, this premium sits at roughly 1.3 $\times$ throughout. On mFRR, it falls from 12.5 $\times$ in Pre-PICASSO to 10.9 $\times$ under PICASSO Adapted and 6.5 $\times$ under MARI Adapted. Manual reserve activation thus remains a stronger discriminator of chasing than automatic reserve activation in every regime, even after MARI accession.

4.2.3 Cascade concentration in the extreme-chasing tail

The matching question to the broad cascade rate is what happens inside the extreme-chasing tail — the chasing-active bins whose realised next-bin price additionally falls outside the within-regime $[p_1, p_{99}]$ window. Three measurements complete the picture (Figure 4.6). The absolute cascade rate inside the extreme-chasing tail declined less than the broad rate, moving from 26.5% to 21.6% and then to 16.7% across the three regimes, a 37% reduction against the 78% broad reduction. The ratio between the extreme-tail cascade rate and the broad chasing cascade rate then rose monotonically, from 3.4 $\times$ in Pre-PICASSO to 6.2 $\times$ under PICASSO Adapted and 9.8 $\times$




Figure 4.5: Mean aFRR activation (left) and mean mFRR activation (right) on chasing-active versus idle bins, by regulatory regime. Bars are means; error bars are 95% confidence intervals. Sample sizes: $n_{\text{active}} = 14,868/3,827/4,046$; $n_{\text{idle}} = 45,563/12,790/11,979$. Both reserves use the per-bin absolute activation magnitude.

under MARI Adapted: the over-representation of cascades in the extreme tail nearly tripled. The extreme-chasing tail itself accounts for roughly 4–5% of chasing-active bins across regimes (626/171/215 bins out of 14,868/3,827/4,046), so the rising ratio captures a genuine concentration of cascade incidence in the tail, not a denominator artefact: it does not arise because the tail population shrank. Read against the matching reward result of [Section 4.1.2](#), the two findings line up symmetrically. On the trader side, the broad chasing reward compressed by 47% over the two reforms while the extreme-chasing tail reward fell by only 11%. On the grid side, the broad cascade rate fell by 78% while the extreme-tail cascade rate fell by only 37%. [Chapter 5](#) takes up this symmetry as the cross-cutting finding of the work.

4.2.4 Robustness and scope

Tightening the cascade threshold to p99 aFRR (139.0 MW) reduces the absolute cascade rate by roughly a factor of five but preserves the qualitative regime ranking: 1.46% $\rightarrow$ 0.97% $\rightarrow$ 0.47% across Pre-PICASSO, PICASSO Adapted, and MARI Adapted (a 68% reduction against 78% under p95). The tail-concentration multiplier under p99 takes a different shape: 1.3 $\times$ $\rightarrow$ 9.1 $\times$ $\rightarrow$ 8.9 $\times$. PICASSO and MARI converge at the rarest cascade threshold, suggesting that the second reform’s incremental contrast appears predominantly in the broader p95 cascade band rather than in the rarest p99 events. The two threshold definitions agree on the central finding — the extreme-chasing tail keeps the cascade — but disagree on the marginal contribution of MARI to the rarest events. Day-block bootstrap intervals ([Section 3.3](#)) on the per-regime cascade rates do not overlap, and the bootstrap confidence intervals for the Pre-PICASSO and MARI Adapted concentration ratios are themselves non-overlapping, indicating that the cascade-concentration shift across the two reforms is statistically distinguishable. The quality-aware restriction moves each regime-level cascade rate by less than half a percentage point.

Two boundaries qualify the result. The cascade indicator is a binary stress flag and not a continuous severity measure. A bin that activates aFRR at 95 MW together with any mFRR




Figure 4.6: Cascade rate ($aFRR > 94.6$ MW together with any $mFRR$ activation) on chasing-active bins (left bar of each regime) compared with the same rate restricted to bins where the realised next-bin price also falls outside the within-regime $[p_1, p_{99}]$ window (right bar, hatched). Multipliers ($N\times$) above each pair give the extreme-to-any ratio. The same thin tail holds both the residual reward and the residual cascade risk.

and a bin that activates at 300 MW together with any $mFRR$ are both flagged identically, so the regime contrasts speak to cascade *incidence* and not to activation severity. The chasing-versus-idle contrast does not isolate a causal effect either. By construction, chasing-active bins coincide with large forecast SI excursions, and large forecast SI excursions independently raise the probability of reserve activation; the contrast therefore does not separate the marginal effect of a BRP taking a chasing position from the joint occurrence of these two events. Whether the residual cascade concentration in the extreme tail reflects a gap in the regulatory framework or a calendar-time anomaly is taken up in [Section 4.3](#).

4.3 Regulatory attribution

Which features of the Belgian regulatory framework drive the chasing-reward compression of [Section 4.1](#) and the cascade-incidence pattern of [Section 4.2](#)? Four candidate levers are examined: the single-pricing rule, the α -component activation threshold, the 15-minute settlement resolution, and the $aFRR/mFRR$ product mix. Three are evaluated with a counterfactual sweep on the observed data; the single-pricing rule is evaluated as a structural feature against the Vermeulen [\[23\]](#) dual-pricing benchmark. The counterfactuals throughout are static: they re-price the realised settlement record under an alternative rule and do not model a behavioural response from BRPs.

4.3.1 Pricing rule: Short–Long spread under single pricing

Under single pricing, BRPs in deficit (system Short, $SI < 0$) and BRPs in surplus (system Long, $SI > 0$) face the same imbalance price within each bin. The within-regime Short–Long spread is therefore the per-bin reward to a BRP that successfully offsets the system imbalance, and a positive spread is a structural feature of the rule rather than a market accident. The spread

compressed monotonically across the two reforms: from 295 EUR/MWh in Pre-PICASSO to 195 EUR/MWh under PICASSO Adapted (−34%) and 120 EUR/MWh under MARI Adapted (−59% vs. Pre-PICASSO; Figure 4.7). The compression operates on both sides: the Short-bin median price fell from 227.7 EUR/MWh to 129.5 EUR/MWh, and the Long-bin median rose from −67.3 EUR/MWh to 9.9 EUR/MWh. The Pre-PICASSO Long median is the more striking number — a BRP whose imbalance was helping to absorb a system surplus was settled at a negative price, paying ∼67 EUR/MWh on average. After MARI the Long median is essentially zero. The spread compression therefore reflects two distinct adjustments: a reduction in Short-bin scarcity rent and the near-elimination of the negative-pricing penalty on helping-direction BRPs.

The dual-pricing comparator of Vermeulen [23] sharpens the interpretation. Under a dual-pricing rule the helping-direction BRP would face a softer price (the day-ahead or a weighted balancing-cost reference) while the harming-direction BRP would continue to face the marginal-cost-plus- α imbalance price; the Short–Long spread under dual pricing would therefore be smaller by construction, but the BRP risk distribution would be wider because the helping-direction reward is dampened. The Belgian single-pricing rule is empirically defensible on equilibrium grounds (Section 2.3), but the chasing reward it generates — the 120 EUR/MWh spread under MARI — never closes.




Figure 4.7: Imbalance price on chasing-active bins ($|\hat{SI}_t| > 150$ MW), split by system imbalance direction (Short: $SI < 0$; Long: $SI > 0$), shown for each regulatory regime. Boxes give the 25th–75th percentile range and whiskers the 5th–95th. The spread, defined as $\text{median}(\text{price} \mid \text{Short}) - \text{median}(\text{price} \mid \text{Long})$, is the per-bin chasing reward under single pricing.

4.3.2 The α -component activation threshold

The α component augments the imbalance price by a scarcity adder when $|SI_t| > 150$ MW. Belgium has held the threshold at 150 MW throughout the 2023–2025 sample, but the threshold is a tariff-methodology choice [12], [15] and could in principle be raised or lowered. Figure 4.8 sweeps the threshold $T \in \{50, 100, 150, 200, 300, 500\}$ MW and reports four metrics per regime under each: chasing-active sample size, mean reward, loss rate, and cascade rate. The sweep produces a single qualitative pattern that repeats across all three regimes. Raising T shrinks the

chasing-active subset (from 38,796 bins at $T = 50$ MW to 796 at $T = 500$ MW in Pre-PICASSO), but the surviving subset has larger-magnitude forecasted imbalance and therefore a sharper realised price signal. Mean reward and loss rate move in opposite directions: from $T = 50$ to $T = 500$ MW, Pre-PICASSO mean reward rises from 84 EUR/MWh to 340 EUR/MWh while the loss rate falls from 29.7% to 6.7%. The cascade rate moves with the reward, rising from 3.7% to 22.4% in Pre-PICASSO — the rare large-volume chasing events that survive a higher threshold are also the bins that drive concurrent aFRR and mFRR activation. Raising α therefore sharpens the per-event chasing reward but concentrates grid stress; lowering α broadens the chasing population but dilutes the per-event signal and reduces the cascade rate. The current $T = 150$ MW sits roughly in the centre of this trade-off; the threshold is therefore a continuous lever where the broad chasing-versus-stress trade-off is parametrised, with effects that operate on trader-side and grid-side metrics simultaneously.




Figure 4.8: Counterfactual sensitivity sweep over α -component activation thresholds $T \in \{50, 100, 150, 200, 300, 500\}$ MW. For each T and each regime, the chasing-active sample is redefined as $|\hat{S}I_t| > T$ and the four metrics are computed from scratch, with no carry-over from the headline $T = 150$ MW subset. The vertical reference line marks the current Belgian threshold.

4.3.3 Settlement resolution: 15-minute versus hourly aggregation

The Imbalance Settlement Period sets the temporal resolution at which BRP deviations and imbalance prices are aligned. Belgium has used a 15-minute ISP for over a decade [27], but a hypothetical hourly ISP can be reconstructed by aggregating the 15-minute series. Three measurements separate the two resolutions (Figure 4.9). Hourly aggregation reduces the chasing event count by 19–27% across regimes on a per-hour-equivalent basis: Pre-PICASSO’s 14,868 quarter-hour chasing bins map to 3,717 hour-equivalents under 4-to-1 collapse, of which only 2,698 retain chasing-active status when imbalance is averaged to the hour. The per-event spread *rises* under hourly settlement: from 295 EUR/MWh to 323 EUR/MWh in Pre-PICASSO, 195 to 230 EUR/MWh under PICASSO Adapted, and 120 to 164 EUR/MWh under MARI Adapted. Each surviving hourly event is a sharper signal because the aggregation has selected the largest, most persistent imbalances. The aggregate hourly chasing income — the product of events per hour and mean profit per event — collapses in every regime: from 123.5 EUR/MWh per hour to 17.8 EUR/MWh per hour in Pre-PICASSO (–86%), with comparable reductions of –86% and –87% under PICASSO and MARI.

Hourly settlement would therefore act as a strong lever against the broad chasing economy

without redesigning the single-pricing rule, but at the cost of weakening the real-time imbalance signal that BRPs and the TSO both rely on for sub-hourly balancing. The counterfactual is also not a Belgian-level instrument: EBGL Article 53 mandates a 15-minute ISP across all EU scheduling areas [6], and any move to a longer settlement period would require a synchronous-area exemption rather than a unilateral Belgian choice.




Figure 4.9: Three-panel comparison between the current 15-minute settlement and a hypothetical hourly settlement on chasing-active bins, shown by regulatory regime. Panel 1 gives the chasing event count per hour at each resolution. Panel 2 gives the per-event Short–Long price spread. Panel 3 gives the aggregate chasing income per hour, taken as events per hour multiplied by mean profit per event.

4.3.4 Balancing product composition

The reserve-activation composition shift documented in [Section 4.2.2](#) can be re-read as a regulatory-feature outcome of the six-month gap between Belgium’s PICASSO and MARI accessions, during which automatic reserve was cross-border-coupled but manual reserve was not. Two regulatory regimes acted on two distinct dimensions ([Figure 4.10](#)). Under PICASSO Adapted, total reserve activation in chasing bins is essentially flat versus Pre-PICASSO (117 → 115 MW, –1%), but the composition swings: the mFRR share rises from 63.5% to 77.9% (+14.4 pp) while the aFRR share falls correspondingly. PICASSO is therefore *compositional*: it constrained automatic reserve activation by exposing it to the cross-border merit order, and the Belgian TSO substituted manual reserve to clear the same volume of imbalance. Under MARI Adapted, total activation halves (115 → 52 MW, –55% vs. PICASSO Adapted; –56% vs. Pre-PICASSO) and the mFRR share partially reverts toward its pre-reform level (77.9% → 66.5%). MARI is therefore *volumetric*: it opened the second reserve product to the same cross-border pool, eliminated the substitution mechanism PICASSO had created, and reduced the total reserve volume needed to clear chasing-active imbalances. Cartuyvels *et al.* [17] characterise the cross-border equilibrium under a single-imbalance-side α adder and predict that adders introduced without simultaneous trading of balancing capacity distort the merit order; the observed PICASSO-period composition shift is consistent with this prediction.

4.3.5 Four-pillar synthesis and scope

[Figure 4.11](#) synthesises the regime contrasts on the four headline metrics: mean reward ([Section 4.1.1](#)), loss rate ([Section 4.1.3](#)), p95 cascade rate ([Section 4.2.1](#)), and the extreme-tail




Figure 4.10: Mean aFRR and mFRR activation on chasing-active bins ($|\hat{S}_t| > 150$ MW), by regulatory regime. Panel 1: stacked composition with total activation labelled above; percentages refer to total versus Pre-PICASSO. Panel 2: mFRR share of total activation. Panel 3: aFRR share (mirror of Panel 2). Delta labels in percentage points are versus Pre-PICASSO.

cascade concentration ratio (Section 4.2.3). Three of the four moved in the direction the reforms were designed to produce: mean reward fell ($+139 \rightarrow +92 \rightarrow +74$ EUR/MWh, -47%), broad cascade rate fell ($7.7\% \rightarrow 3.5\% \rightarrow 1.7\%$, -78%), and total reserve activation fell ($117 \rightarrow 115 \rightarrow 52$ MW, -55%). The fourth moved against the reform direction: the loss rate rose monotonically ($20.0\% \rightarrow 36.0\% \rightarrow 34.5\%$) and the extreme-tail cascade concentration nearly tripled ($3.4\times \rightarrow 6.2\times \rightarrow 9.8\times$). The reforms compressed the broad chasing economy and the broad grid-stress incidence on the same sample, but the trader-side risk and the grid-side tail concentration both shifted in the opposite direction.

The four-pillar decomposition links each metric movement to a distinct feature of the regulatory framework. The single-pricing rule fixes the structural existence of the Short–Long spread, and therefore the structural existence of a positive mean reward. Both reforms left the rule in place, and the spread did not close. The α threshold controls where on the reward–loss–cascade trade-off curve the chasing rule operates. Holding the threshold fixed at 150 MW while the European platforms compressed the underlying price distribution produced the combined move observed here: lower mean reward, higher loss rate, and a lower broad cascade rate. The 15-minute ISP shapes the aggregate volume of chasing activity rather than the per-event metrics. The aFRR/mFRR composition is the channel through which the two European platforms acted on the system; PICASSO redistributed activation across the two reserves without lowering the total, and MARI then lowered the total. No single pillar carries the observed shifts on its own. The four levers work on different dimensions of the same trade-off, and the combined effect of the two reforms is the joint outcome of these four levers acting at once.

Three caveats apply to the four-pillar decomposition. The counterfactual sweeps are static and do not model BRP behavioural response: under a different α threshold or settlement resolution, BRPs would adjust their forecasting horizon, position sizing, and tolerance for the loss rate, and these second-order effects are not captured. The single-pricing comparator draws on the Belgian–Dutch contrast of Vermeulen [23] rather than a within-Belgium counterfactual: it is a comparison between two empirically distinct market designs, not between two alternative parameterisations of the Belgian rule. The four-pillar decomposition reports which feature moves together with which metric rather than which feature *caused* which movement; the cross-border




Figure 4.11: Four-panel synthesis of the regime contrasts on chasing-active bins. Panel 1: mean reward per chasing-active bin, with the top 1% extreme-chasing tail (hatched) overlaid. Panel 2: loss rate, Wilson 95% confidence intervals as black bars. Panel 3: cascade rate (aFRR > p95 = 94.6 MW and any mFRR activation). Panel 4: tail concentration ratio (cascade rate in the extreme-chasing tail divided by the cascade rate on all chasing-active bins).

attribution that underlies this caveat is taken up in [Chapter 5](#).

4.4 Policy reform potential

Which incremental policy instruments could compress the residual chasing incentive without redesigning the single-pricing rule? The four-pillar decomposition of [Section 4.3](#) identified the regulatory levers; the question now is whether any incremental adjustment to those levers materially compresses the residual chasing reward that survived the PICASSO and MARI reforms. The unit of analysis is the extreme-chasing tail of [Section 4.1.2](#) — the within-regime $[p_1, p_{99}]$ outside slice of chasing-active bins — because the broad chasing economy has already been compressed by the European integration reforms ([Section 4.1.1](#)), and the residual incentive is what remains in the tail. Three single-instrument counterfactuals and one combined instrument are evaluated against this tail under the static framing of [Section 3.2](#): BRP positions are held fixed at the observed values. The status-quo reward on extreme-chasing bins — 444 EUR/MWh in Pre-PICASSO, 496 EUR/MWh under PICASSO Adapted, 493 EUR/MWh under MARI Adapted — is the benchmark against which each instrument is measured ([Figure 4.12](#)).

4.4.1 P1: within-regime price kink

P1 clips the realised next-bin imbalance price to the within-regime $[p_5, p_{95}]$ window, simulating a regulatory bound on how far the settlement price is allowed to drift in either direction inside the α -active region. The clipping is two-sided — both the negative tail (associated with system surplus) and the positive tail (associated with system shortage) are bounded — and the clipping percentile is computed regime-by-regime so that the same instrument is applied consistently across regimes. The compression on the extreme-chasing tail is substantial in every regime ([Figure 4.12](#)): per-bin reward falls from 444 EUR/MWh to 225 EUR/MWh in Pre-PICASSO (−49%), from 496 EUR/MWh to 156 EUR/MWh under PICASSO Adapted (−69%), and from 493 EUR/MWh to 114 EUR/MWh under MARI Adapted (−77%). The regime-average compression is −65%, and the marginal effect grows stronger with each reform; the post-MARI tail is the most responsive




Figure 4.12: Per-bin reward on the extreme-chasing tail under the status-quo and four counterfactual rules, by regulatory regime. Extreme-chasing tail = chasing-active bins ($|\hat{S}_t| > 150$ MW) whose realised next-bin price falls outside the within-regime $[p_1, p_{99}]$ window. P1: next-bin price clipped to within-regime $[p_5, p_{95}]$. P2: chasing-active threshold raised from 150 MW to 250 MW. P3: 100 EUR/MWh flat fee on bins with $|mFRR| > 0$. P1 + P3: both treatments applied simultaneously. Labels show absolute reward in EUR/MWh and percentage change versus status-quo. Static counterfactual: BRP positions held fixed at the observed values; behavioural response not modelled.

to the price kink. The mechanism itself is direct, since P1 caps the realised price magnitude on exactly the bins that drive the extreme-chasing reward. What is a finding here is the regime-by-regime response: the post-MARI extreme tail concentrates its residual reward into the slice of the price distribution most exposed to a two-sided cap. P1 stays within the existing single-pricing rule. The pricing formula and the α scarcity adder are both unchanged; only the post-formation realised price is bounded.

4.4.2 P2: raised α threshold

P2 raises the α -component activation threshold from 150 MW to 250 MW, narrowing the chasing-active sample to bins with larger-magnitude forecasted imbalance. The result of this counterfactual runs against the policy intent. Per-bin reward on the extreme-chasing tail *rises* in every regime (Figure 4.12): from 444 EUR/MWh to 529 EUR/MWh in Pre-PICASSO (+19%), from 496 EUR/MWh to 549 EUR/MWh under PICASSO Adapted (+11%), and from 493 EUR/MWh to 549 EUR/MWh under MARI Adapted (+11%). The regime-average effect is +14%. Raising the α threshold therefore does not compress the residual reward; it sharpens it. The same selection effect documented in the sensitivity sweep of Section 4.3.2 is at work: tightening the chasing-active condition removes the low-magnitude forecast bins that diluted the per-event reward, and the smaller but sharper subset that remains carries a higher mean reward than the original mixed subset did. P2 is therefore disqualified as a stand-alone compression instrument, and the role of the α threshold becomes clearer: it is a position-sizing parameter for the chasing rule, not an incentive-compression parameter.

4.4.3 P3: manual-reserve activation fee

P3 imposes a flat 100 EUR/MWh activation fee on every bin in which the BRP’s chasing-active position coincides with a non-zero mFRR dispatch. The instrument internalises the manual-reserve externality that the cascade rate of [Section 4.2.1](#) associates with chasing activity. The compression is modest and approximately regime-invariant. Per-bin reward on the extreme-chasing tail falls from 444 EUR/MWh to 373 EUR/MWh in Pre-PICASSO (−16%), from 496 EUR/MWh to 420 EUR/MWh under PICASSO Adapted (−15%), and from 493 EUR/MWh to 419 EUR/MWh under MARI Adapted (−15%) ([Figure 4.12](#)). The regime-average compression of −15% is the smallest of the three single instruments, but the regime invariance is itself informative: the fee acts in the same proportion across regimes because the underlying mFRR dispatch rate on the extreme-chasing tail is comparable across regimes ([Section 4.2.3](#)). P3 therefore operates on the cascade-prone subset of the extreme tail rather than on the price distribution, and charges the BRP exactly on the bins where the residual cascade incidence is most visible — transferring a share of the reward back to the system that bore the manual-dispatch cost.

4.4.4 P1 + P3 combined, and scope

The two instruments that move in the right direction are evaluated jointly. The combined rule clips the realised price to the within-regime $[p_5, p_{95}]$ window *and* charges a 100 EUR/MWh fee on bins with non-zero mFRR activation. P2 is excluded from the combination because it sharpens rather than compresses. The combined instrument compresses the extreme-chasing reward materially in every regime ([Figure 4.12](#)): from 444 EUR/MWh to 155 EUR/MWh in Pre-PICASSO (−65%), from 496 EUR/MWh to 80 EUR/MWh under PICASSO Adapted (−84%), and from 493 EUR/MWh to 40 EUR/MWh under MARI Adapted (−92%). The compression strengthens monotonically across the two reforms and reaches near-elimination under MARI: the post-MARI extreme tail under P1 + P3 retains less than one-tenth of its status-quo reward. The regime-average compression of −80% reflects that the two instruments act on overlapping but not identical subsets: P1 caps the realised price on every bin with an extreme realisation, P3 charges only the mFRR-active subset, and the two together strip both layers of the residual reward. The combined instrument preserves three properties an incremental policy should preserve: the single-pricing rule itself is unchanged (BRPs in deficit and BRPs in surplus continue to face the same imbalance price within each bin); the α threshold at which chasing-active behaviour is defined is unchanged at 150 MW; and the 15-minute ISP is unchanged. P1 + P3 therefore compresses the residual chasing reward by 80% on average without redesigning any of the four regulatory pillars examined in [Section 4.3](#).

Two scope notes qualify the result. The within-regime $[p_5, p_{95}]$ clip of P1 and the 100 EUR/MWh fee of P3 are single-point evaluations and not calibrated proposals; the corresponding elasticity surfaces are not computed here. The broader welfare consequences of the two instruments, namely the dampened imbalance signal under P1 and the cost transfer back to the system under P3, fall outside the chasing-reward framing of this chapter and are taken up in [Chapter 5](#).

Chapter 5: Discussion

The empirical results of [Sections 4.1 to 4.4](#) answer the four research questions one at a time, but they carry a single cross-cutting finding: the two European integration reforms compressed the broad chasing economy on every measured metric, yet the residual reward and the residual grid stress are concentrated in the same thin extreme-chasing tail. This chapter synthesises the four-RQ findings around that central symmetry ([Section 5.1](#)), addresses the cross-border attribution caveat that recurs throughout ([Section 5.2](#)), and places the result in the cross-jurisdictional context of Vermeulen [23] alongside the implied regulatory consequences ([Section 5.3](#)).

5.1 Cross-cutting symmetry

The trader-side and grid-side findings are symmetric in shape. The broad chasing reward compressed by 47% between Pre-PICASSO and MARI Adapted while the within-regime extreme-chasing tail reward fell by only 11% ([Section 4.1.2](#)). The broad cascade rate on chasing-active bins fell by 78% over the same window while the cascade rate inside the extreme-chasing tail fell by only 37%, and the over-representation of cascades in the extreme tail nearly tripled ([Section 4.2.3](#)). The two reforms moved the bulk of both distributions in the direction the European framework was designed to produce, but in each a thin residual tail remained materially intact. The broad-set compression is consistent with independent industry analysis: van der Werf [29] reports a $\sim 55\%$ reduction in the day-ahead-to-imbalance spread in Belgium following PICASSO accession, and documents that by 2025 the cumulative gains from a rule-based curtailment strategy in the Dutch market had “largely disappeared.” Both observations match the central-mass compression measured here without depending on the chasing-rule construction of [Section 3.2](#).

The mechanical interpretation of the rising tail-concentration ratio is direct: the denominator (the broad cascade rate) fell faster than the numerator (the cascade rate inside the extreme tail), so the ratio rose. The same arithmetic produces the rising kurtosis on the reward distribution. The chasing population that survived both reforms is therefore qualitatively different from the population that was active under Pre-PICASSO: the pre-reform reward came from a broad mix of mid-magnitude price excursions, while the post-MARI reward comes from a narrow set of large-magnitude excursions. [Section 4.1.3](#) captures this in the loss-rate movement from 20% to 35% — consistent with a residual population that is both more rewarding when it pays and more risk-laden when it does not. The grid-stress concentration is the system-side analogue of the same transition: the same thin tail that pays the residual reward to BRPs generates the residual cascade incidence in the reserve-activation series. [Section 4.4](#) acted on this overlap directly — P1+P3 compresses the post-MARI extreme-tail reward by 92% because the price-cap component (P1) and the mFRR-fee component (P3) intersect on the bins that drive both the residual reward and the residual cascade.

5.2 Cross-border attribution

PICASSO and MARI are European platforms, not Belgian reforms. Belgium is one control area among the cross-border-coupled members [6], [11], and the observed post-reform behaviour reflects the joint effect of (i) the Belgian-specific settlement design held fixed across the sample and (ii) the European-wide changes the platforms introduced. The regime contrasts therefore confound a domestic component (the unchanged Belgian framework operating in a new external environment) with an international component (the new European environment itself), and the descriptive framing applied throughout the thesis is the conservative response to this confound.

The cross-border equilibrium analysis of Cartuyvels *et al.* [17] provides the closest analytical handle for Belgium’s α design: it is an instance of the second pricing-asymmetry case (imbalance-side adder, no balancing-capacity trade), and the composition shift of Section 4.3.4 — aFRR activation falling under PICASSO while mFRR activation rises — is consistent with their merit-order distortion prediction. The subsequent MARI accession coupled the second product and eliminated the temporary substitution mechanism, again consistent with the equilibrium framework. The composition–volume sequencing observed in the data is therefore not a Belgian regulatory choice; it is the joint outcome of a Belgian feature (the α adder on the imbalance side) and a European-timetable feature (PICASSO before MARI by six months). The attribution caveat reaches its sharpest form on the chasing-reward compression itself: the Short–Long spread compression of Section 4.3.1 ($295 \rightarrow 195 \rightarrow 120$ EUR/MWh) operates entirely within the unchanged Belgian single-pricing rule, but the price distribution from which the spread is drawn is the post-coupling Belgian distribution — the same distribution the European platforms had reshaped. Disentangling the two components requires a counterfactual in which Belgium would have remained un-coupled, which is unobservable on Belgian data alone (Section 6.2).

5.3 Cross-jurisdictional context and regulatory implications

The Belgian sample analysed by Vermeulen [23] ends before MARI accession; whether the equilibrium properties documented there survive the second reform is the cross-jurisdictional question this section takes up. Three findings of this work line up with the Vermeulen frame and one extends it. The single-pricing rule did not produce a more volatile market outcome over the two-reform window: broad-tail and central-mass compression moved in the same direction (Section 5.1), consistent with the dampening property Vermeulen attributes to the rule. BRP-side risk rose within the rule — the loss rate moved from 20% to 35% (Section 4.1.3), locating the residual risk on the per-event distribution rather than on the system-stability distribution, exactly the trade-off the dual-pricing comparator predicts. The residual reward concentration in the extreme tail is a within-rule phenomenon: the single-pricing rule produces a residual chasing economy whose intensity is tail-driven rather than population-driven. The extension is post-MARI: even after the second European integration reform, the residual tail-concentrated reward does not close on its own. The Belgian rule defends its choice on equilibrium grounds but at the cost of a persistent tail-concentrated residual that cross-border platforms alone do not eliminate.

The P1+P3 instrument of Section 4.4 can be read in this comparative frame as an incremental adjustment within the single-pricing rule that compresses the tail-concentrated residual without

converting the rule to dual pricing. The price-kink component (P1) acts symmetrically on the upper and lower tails of the realised price distribution and is therefore not a dual-pricing instrument; the mFRR-fee component (P3) is a transfer from the chasing-active BRP back to the system on the cascade-prone subset and is structurally distinct from the dual-pricing mechanism Vermeulen evaluates. The instrument therefore preserves the empirical advantages of the Belgian rule (market-outcome stability, lower BRP-side risk on the central mass) while addressing the residual the Vermeulen comparator identifies.

Three regulatory implications follow. The four-pillar decomposition of [Section 4.3](#) shows that no single pillar can be tuned in isolation. The α threshold is a position-sizing parameter and not a reward-compression lever ([Section 4.4.2](#)); moving the 15-minute ISP to hourly would dampen the real-time imbalance signal ([Section 4.3.3](#)); the single-pricing rule is the structural anchor ([Section 4.3.1](#)); and the product-mix composition is a European-integration outcome and not a Belgian design choice ([Section 4.3.4](#)). The remaining regulatory degrees of freedom are therefore narrow, and the P1+P3 instrument is one of the few combinations that compresses the residual without crossing into any of these constraints. The composition–volume sequencing documented in [Section 4.3.4](#) is itself an outcome of the staggered European timetable and not a Belgian design choice; any future EU-wide integration on a third reserve product should therefore be evaluated against the expectation that staggered timing will produce similar transient effects. Finally, the residual chasing reward and the residual cascade concentrate in the same thin tail ([Section 5.1](#)), so what the framework is missing is a tail-targeted instrument and not a structural redesign. The static counterfactual evidence points in the same direction: the residual is mechanically compressible and not structurally locked in.

5.4 Limitations

Four structural limitations qualify the findings beyond the chapter-by-chapter caveats already stated in [Sections 4.1.4](#), [4.2.4](#), [4.3.5](#) and [4.4.4](#).

This analysis is carried out on a single control area. The cross-border attribution caveat of [Section 5.2](#) is the operational form of that limitation: post-reform observations are the joint outcome of the Belgian rule and the European platforms, and Belgian data alone cannot disentangle the two components. A comparable study on a second control area would close this gap; the Netherlands operates under the same European platforms but under a dual-pricing rule [23], and any other LFC area that joined the platforms on a different date would also work. Either route would supply the within-rule and within-platform variation that the Belgian sample on its own does not.

The chasing rule of [Section 3.2](#) is a settlement-cashflow proxy built on the published SI forecast and not a measurement of realised BRP behaviour. Individual BRP position data are held by Elia and are not publicly available. The thesis cannot therefore distinguish a bin on which a BRP took a chasing position from a bin on which the BRP held a balanced portfolio that happened to coincide with a large forecast SI. The reward defined in [Equation \(3.1\)](#) is the cashflow a chasing BRP would have collected, not the cashflow any actual BRP did collect. The cascade-rate finding of [Section 4.2](#) carries the same qualification: it documents the association between large forecast SI and concurrent reserve activation, and does not measure the marginal effect of a BRP chasing

decision on that activation.

The regime contrasts confound the two reforms with parallel calendar-time trends. Belgian solar PV capacity, wind capacity, and behind-the-meter flexibility all grew non-trivially across the 2023–2025 sample window, raising the variability of the realised system imbalance independently of the PICASSO and MARI reforms. A larger SI variability mechanically widens the distribution from which the chasing reward is drawn, which can compress the central mass and lengthen the right tail in the same direction the European integration reforms operate. The descriptive framing of [Section 3.3](#) is the conservative response, but it does not separately identify the regulatory and the renewable-growth components of the observed shifts.

The post-MARI window is short. MARI accession occurred on 21 May 2025 and the sample ends on 31 December 2025, leaving approximately seven months of post-MARI observations. The Pre-PICASSO baseline contains roughly $3.7\times$ as many chasing-active bins as the MARI Adapted window ([Table 3.1](#)), and several of the post-reform extreme-tail metrics are based on small samples — 171 extreme-tail bins in PICASSO Adapted and 215 in MARI Adapted — in particular the tail-concentration findings of [Sections 4.2.3](#) and [4.4.4](#). The day-block bootstrap intervals of [Sections 4.1.4](#) and [4.2.4](#) preserve the qualitative regime ordering, but the absolute magnitudes of the post-MARI estimates carry wider uncertainty than the corresponding Pre-PICASSO estimates. A re-evaluation of the same metrics on a longer post-MARI window would tighten the estimates without changing the methodology.

Chapter 6: Conclusion and Future Work

6.1 Conclusion

The Belgian imbalance settlement design rewards Balance Responsible Parties that take an offsetting position against the realised system imbalance, a strategy known as imbalance chasing. On 27 November 2024 and 21 May 2025, Belgium acceded to the European balancing platforms PICASSO and MARI, completing the cross-border integration of automatic and manual reserve activation within six months. This work measured, on a 2023–2025 fifteen-minute sample of Belgian imbalance data, whether and how the chasing incentive shifted across the two reforms, whether chasing-active behaviour coincided with elevated grid stress, which features of the Belgian regulatory framework drove the observed shifts, and which incremental policy instruments could compress the residual chasing incentive without redesigning the single-pricing rule.

The four research questions are answered together. On RQ1 (Section 4.1), the mean per-bin chasing reward fell by 47% between Pre-PICASSO and MARI Adapted (from +139 to +74 EUR/MWh), the loss rate climbed from 20% to 35%, and the within-regime extreme-chasing tail reward fell by only 11%. On RQ2 (Section 4.2), the cascade rate on chasing-active bins fell by 78% over the same window (from 7.7% down to 1.7%); the cascade rate inside the extreme-chasing tail fell by only 37%, and the over-representation of cascades in the tail grew from $3.4\times$ to $9.8\times$. On RQ3 (Section 4.3), the four regulatory pillars — single pricing, the 150 MW α threshold, the 15-minute settlement resolution, and the aFRR/mFRR product mix — act on different dimensions of the chasing–cascade trade-off; PICASSO acted compositionally on the reserve mix, and MARI acted volumetrically on total activation. On RQ4 (Section 4.4), a within-regime price kink at the $[p_5, p_{95}]$ window combined with a 100 EUR/MWh mFRR-bin fee ($P1 + P3$) compresses the post-MARI extreme-tail reward by 92%, leaving the single-pricing rule, the α threshold, and the 15-minute ISP unchanged.

The contribution of this work is threefold. The first piece is a transparent chasing rule that uses only the public Elia SI forecast available before the settlement quarter; the rule is applied consistently across a 244,773-bin master panel of Elia public feeds spanning 2019–2025, with the chasing analysis itself carried out on a 93,073-bin forecast-available subsample over 2023–2025. The second piece is the symmetric tail finding on the post-MARI Belgian sample: the broad chasing reward and the broad cascade incidence both compressed, while the same thin extreme-chasing tail held on to the residual reward and the residual cascade. The third piece is the policy evaluation: three single-instrument and one combined static counterfactual, applied to that residual tail, show that what the framework is missing is a tail-targeted instrument and not a structural redesign.

The Belgian single-pricing rule and the European integration framework compressed the broad chasing economy on every measured metric. The residual chasing reward that survived both reforms is concentrated in the extreme-chasing tail and can be compressed by an incremental policy instrument without redesigning the single-pricing rule. Whether to leave it in place or compress it through such an instrument is a regulatory choice; this work has measured its size on Belgian data.

6.2 Future directions

Three directions follow from the limitations of [Section 5.4](#).

A within-EBGL cross-area replication would address the single-area limitation and the cross-border attribution caveat simultaneously. Applying the same master-dataset methodology to Dutch, German, or French data over the 2023–2025 window would separate Belgian-rule-specific effects from European-platform-wide ones, and tracking reserve activation across neighbouring control areas would identify whether MARI shifted Belgian activation onto the cross-border platform rather than reducing it in aggregate. The Dutch case offers the cleanest contrast: Dutch BRPs face a dual-pricing rule under the same European platforms, so the post-MARI Belgian–Dutch contrast on a common sample window would isolate the marginal effect of the pricing rule conditional on the same platform environment.

A behavioural-response simulation would address the static-counterfactual limitation of [Section 4.4.4](#). The four counterfactuals of [Section 4.4](#) hold the BRP position fixed; a model that allows re-optimisation under the alternative rule would tighten the compression bound. Existing forecasting-and-trading tools [\[20\]–\[22\]](#) provide the participation-side building blocks, but no equilibrium simulation between many BRPs and the TSO has been carried out for the Belgian case.

A longer-window re-evaluation of the post-MARI regime would address the sample-size limitation. The methodology is fully reproducible from the public Elia feeds; extending by calendar-year increments after the defence date would refine post-MARI estimates and permit a second-year-versus-first-year comparison. The framework also extends to a third reserve product: with the European platforms preparing additional balancing products beyond aFRR and mFRR [\[4\]](#), [\[5\]](#), the composition-shift mechanism of [Section 4.3.4](#) would re-emerge during any future staggered accession, and the framework anticipates the transient shape of those effects in advance.

References

- [1] Elia Transmission Belgium, *Imbalance prices: Quarter-hourly open data*, Elia open-data portal, Accessed: 2026-02-11, 2026. [Online]. Available: <https://www.elia.be/en/grid-data/balancing/imbalance-prices-15>.
- [2] H. Penn, *2024 in review: The events that shaped short-term power markets*, Dexter Energy, Accessed: 2026-02-11, Jan. 2025. [Online]. Available: <https://dexterenergy.ai/news/2024-events-in-short-term-power-markets/>.
- [3] J. Bottieau, L. Hubert, Z. De Grève, F. Vallée, and J.-F. Toubeau, “Very-short-term probabilistic forecasting for a risk-aware participation in the single price imbalance settlement,” *IEEE Transactions on Power Systems*, vol. 35, no. 2, pp. 1218–1230, 2020. DOI: [10.1109/TPWRS.2019.2940756](https://doi.org/10.1109/TPWRS.2019.2940756).
- [4] Elia Group, “Press release: Elia transmission belgium successfully accesses picasso,” Elia Group, Press Release, Nov. 2024, Press release dated 28 November 2024; PICASSO go-live: 26 November 2024 at quarter-hour 11:15–11:30. Accessed: 2026-02-11. [Online]. Available: https://www.elia.be/en/newsroom/2024/11/20241128_picasso-press-release.
- [5] Elia Group, “Press release: Elia transmission belgium gains access to mari, the european platform for exchanging tertiary balancing energy,” Elia Group, Press Release, May 2025, Press release dated 23 May 2025; MARI go-live: 21 May 2025. Accessed: 2026-02-11. [Online]. Available: https://www.elia.be/en/newsroom/2025/05/20250523_elia-transmission-belgium-gains-access-to-mari.
- [6] European Commission, *Commission regulation (eu) 2017/2195 of 23 november 2017 establishing a guideline on electricity balancing*, Official Journal of the European Union, CELEX: 32017R2195, 2017. [Online]. Available: <http://data.europa.eu/eli/reg/2017/2195/oj>.
- [7] A. Eicke, O. Ruhnau, and L. Hirth, “Electricity balancing as a market equilibrium: An instrument-based estimation of supply and demand for imbalance energy,” *Energy Economics*, vol. 102, p. 105 455, Oct. 2021, ISSN: 0140-9883. DOI: [10.1016/j.eneco.2021.105455](https://doi.org/10.1016/j.eneco.2021.105455).
- [8] R. A. C. van der Veen, A. Abbasy, and R. A. Hakvoort, “Agent-based analysis of the impact of the imbalance pricing mechanism on market behavior in electricity balancing markets,” *Energy Economics*, vol. 34, no. 4, pp. 874–881, Jul. 2012, ISSN: 0140-9883. DOI: [10.1016/j.eneco.2012.04.001](https://doi.org/10.1016/j.eneco.2012.04.001).
- [9] M. Backer, D. Keles, and E. Kraft, “The economic impacts of integrating european balancing markets: The case of the newly installed aFRR energy market-coupling platform PICASSO,” *Energy Economics*, vol. 128, p. 107 124, Nov. 2023, ISSN: 0140-9883. DOI: [10.1016/j.eneco.2023.107124](https://doi.org/10.1016/j.eneco.2023.107124).
- [10] K. Poplavskaia, J. Lago, and L. de Vries, “Effect of market design on strategic bidding behavior: Model-based analysis of european electricity balancing markets,” *Applied Energy*, vol. 270, p. 115 130, Jul. 2020, ISSN: 0306-2619. DOI: [10.1016/j.apenergy.2020.115130](https://doi.org/10.1016/j.apenergy.2020.115130).

- [11] ENTSO-E, “Go-live of PICASSO and MARI platforms: Completing the EBGL target implementation,” ENTSO-E, Tech. Rep., Dec. 2022, Published 2 December 2022. Accessed: 2026-02-11. [Online]. Available: <https://www.entsoe.eu/news/2022/12/02/go-live-of-picasso-and-mari-platforms-that-complete-the-ambitious-electricity-balancing-regulation-target-implementation/>.
- [12] Elia Transmission Belgium, *Annex 2 – assessment of the α parameter and proposal for recalibration*, Public consultation document, Elia Group, Accessed: 2026-02-06, Dec. 2021. Accessed: Feb. 11, 2026. [Online]. Available: https://www.elia.be/-/media/project/elia/elia-site/public-consultations/2021/20211203_annex2_alpharevision_consultation.pdf.
- [13] J. Baetens, J. Laveyne, G. Van Eetvelde, and L. Vandeveld, “Imbalance pricing methodology in Belgium: Implications for industrial consumers,” in *2020 17th International Conference on the European Energy Market (EEM)*, IEEE, Sep. 2020, pp. 1–6. DOI: [10.1109/EEM49802.2020.9221978](https://doi.org/10.1109/EEM49802.2020.9221978).
- [14] R. A. C. van der Veen and R. A. Hakvoort, “The electricity balancing market: Exploring the design challenge,” *Utilities Policy*, vol. 43, pp. 186–194, Dec. 2016, ISSN: 0957-1787. DOI: [10.1016/j.jup.2016.10.008](https://doi.org/10.1016/j.jup.2016.10.008).
- [15] CREG, “Study on the implementation of a scarcity pricing mechanism in belgium,” Commission for the Regulation of Electricity and Gas, Tech. Rep. (F)2203, Dec. 2021, Accessed: 2026-02-01. [Online]. Available: <https://www.creg.be/sites/default/files/assets/Publications/Studies/F2203EN.pdf>.
- [16] J. Cartuyvels and A. Papavasiliou, “Calibration of operating reserve demand curves using a system operation simulator,” *IEEE Transactions on Power Systems*, vol. 38, no. 4, pp. 3043–3055, Jul. 2023, ISSN: 0885-8950. DOI: [10.1109/TPWRS.2022.3200125](https://doi.org/10.1109/TPWRS.2022.3200125).
- [17] J. Cartuyvels, G. Bertrand, and A. Papavasiliou, “Market equilibria in cross-border balancing platforms,” *IEEE Transactions on Energy Markets, Policy and Regulation*, vol. 2, no. 2, pp. 162–174, Jun. 2024. DOI: [10.1109/TEMPR.2024.3360475](https://doi.org/10.1109/TEMPR.2024.3360475).
- [18] R. Smets, K. Bruninx, J. Bottieau, J.-F. Toubreau, and E. Delarue, “Strategic implicit balancing with energy storage systems via stochastic model predictive control,” *IEEE Transactions on Energy Markets, Policy and Regulation*, vol. 1, no. 4, pp. 373–385, Dec. 2023. DOI: [10.1109/TEMPR.2023.3267552](https://doi.org/10.1109/TEMPR.2023.3267552).
- [19] E. Wessel, R. Smets, and E. Delarue, “Risk-aware participation in day-ahead and real-time balancing markets for energy storage systems,” *Electric Power Systems Research*, vol. 235, p. 110 741, 2024, ISSN: 0378-7796. DOI: [10.1016/j.epsr.2024.110741](https://doi.org/10.1016/j.epsr.2024.110741).
- [20] J. Dumas, I. Boukas, M. M. de Villena, S. Mathieu, and B. Cornélusse, “Probabilistic forecasting of imbalance prices in the Belgian context,” in *2019 16th International Conference on the European Energy Market (EEM)*, IEEE, Sep. 2019, pp. 1–7. DOI: [10.1109/EEM.2019.8916375](https://doi.org/10.1109/EEM.2019.8916375).

- [21] R. Bruneel, M. Schuurmans, and P. Patrinos, “Optimal intraday power trading for single-price balancing markets: An adaptive risk-averse strategy using mixture models,” *Applied Energy*, vol. 389, p. 125 754, Jul. 2025, ISSN: 0306-2619. DOI: [10.1016/j.apenergy.2025.125754](https://doi.org/10.1016/j.apenergy.2025.125754).
- [22] F. Pavirani, J. Van Gompel, S. S. Karimi Madahi, B. Claessens, and C. Develder, “Predicting and publishing accurate imbalance prices using Monte Carlo Tree Search,” *Applied Energy*, vol. 392, p. 125 944, Aug. 2025, ISSN: 0306-2619. DOI: [10.1016/j.apenergy.2025.125944](https://doi.org/10.1016/j.apenergy.2025.125944).
- [23] E. S. Vermeulen, “Balancing markets: Imbalance pricing designs — analysis of implicit balancing by flexible assets within european balancing markets,” Master’s thesis, Technology, Policy and Management, Delft University of Technology, Delft, Netherlands, 2025. [Online]. Available: <https://resolver.tudelft.nl/uuid:ee89c602-316c-4d89-9c83-42af4838efb6>.
- [24] S. S. Karimi Madahi, K. Bruninx, B. Claessens, and C. Develder, “Gaming strategies in European imbalance settlement mechanisms,” in *2025 IEEE PES Innovative Smart Grid Technologies Conference Europe (ISGT Europe)*, Accessed: 2026-01-17, Valletta, Malta: IEEE, Oct. 2025, pp. 1–5, ISBN: 979-8-3315-2504-0. DOI: [10.1109/ISGTEurope64741.2025.11305533](https://doi.org/10.1109/ISGTEurope64741.2025.11305533).
- [25] M. Bruninx, S. S. Karimi Madahi, T. Verstraeten, J. Decuyper, C. Develder, and J. Helsen, *Herd behavior in decentralized balancing models: A case study in Belgium*, arXiv preprint arXiv:2602.17352, Version 2. Accessed: 2026-05-15, May 2026. DOI: [10.48550/arXiv.2602.17352](https://doi.org/10.48550/arXiv.2602.17352).
- [26] J. Cartuyvels, G. Bertrand, and A. Papavasiliou, “Interactions of imbalance settlement with energy and reserve markets in multi-product european balancing markets,” Université catholique de Louvain, CORE, LIDAM Discussion Paper 2024/07, May 2024, Accessed: 2026-02-11. [Online]. Available: <https://ideas.repec.org/p/cor/louvco/2024007.html>.
- [27] Elia Transmission Belgium, “Terms and conditions for balance responsible parties (t&c brp),” Elia Group, Tech. Rep., Sep. 2024, Accessed: 2026-02-11. [Online]. Available: <https://www.elia.be/-/media/project/elia/elia-site/public-consultations/2024/20240920tc-brpvnotrackchangeen.pdf>.
- [28] B. W. Silverman, *Density Estimation for Statistics and Data Analysis* (Monographs on Statistics and Applied Probability 26). London: Chapman & Hall, 1986, ISBN: 978-0-412-24620-3. DOI: [10.1007/978-1-4899-3324-9](https://doi.org/10.1007/978-1-4899-3324-9).
- [29] N. van der Werf, *Balancing markets in 2026: More efficient, less forgiving*, Dexter Energy, Accessed: 2026-03-02, Feb. 2026. [Online]. Available: <https://dexterenergy.ai/news/balancing-markets-in-2026/>.

Appendix A: Additional Robustness Checks

This appendix collects two robustness checks and two reference lookup tables that frame the empirical results of [Sections 4.1 to 4.4](#) but do not fit in the main text. [Appendix A.1](#) reports the quality-aware restriction referenced in [Section 3.3](#) in full. [Appendix A.2](#) documents the mFRR activation distribution that underlies the cascade definition of [Section 3.2](#). [Appendix A.3](#) tabulates the policy counterfactual results visualised in [Figure 4.12](#), and [Appendix A.4](#) provides the full α -threshold sensitivity sweep underlying [Figure 4.8](#).

A.1 Quality-aware restriction

[Section 3.3](#) introduces a quality-aware robustness check alongside the day-block bootstrap. The headline metrics throughout the thesis use the full chasing-active sample (validated and unlabelled records together). The quality-aware restriction further drops every bin where any chasing-relevant input variable does not carry an explicit “validated” flag in the source Elia files.

Imbalance price and realised system imbalance carry a validation flag on essentially every bin in the 2019–2025 sample ([Section 3.1](#)). The restriction therefore operates almost entirely through two other flags: the activation-volume flag, which is unlabelled on 54% of bins, and the SI-forecast flag, which captures the two Elia publication gaps documented in [Section 3.1](#). After the restriction is applied, the chasing-active sample shrinks to roughly 60% of the full sample, and the loss is distributed unevenly across regimes: the publication gaps fall in 2023 (Pre-PICASSO) and in late 2025 (MARI Adapted).

The headline regime contrasts of [Sections 4.1](#) and [4.2](#) survive the restriction. Mean per-bin reward shifts by less than 5% in either direction at every regime. Loss rate shifts by less than half a percentage point. The cascade rate at the p95 threshold shifts by less than 0.3 pp. The tail-concentration ratio holds its monotone increase across the two reforms ($3.4\times \rightarrow 6.2\times \rightarrow 9.8\times$ on the full sample, with the same direction inside the same confidence interval on the restricted sample). No pairwise regime contrast that is significant on the full sample loses significance on the restricted sample, and the qualitative central symmetry (broad compression together with tail concentration) is robust to the restriction.

The day-block bootstrap intervals of [Sections 4.1.4](#) and [4.2.4](#) are several times wider than the full-versus-restricted-sample differences on every metric, so the restriction is well inside the sampling-uncertainty band: the regime contrasts reported in this work are not driven by non-validated observations.

A.2 mFRR activation distribution

The cascade definition of [Section 3.2](#) combines an aFRR threshold (> 94.6 MW, the full-sample p95) with an mFRR condition (any non-zero activation). The asymmetry is a direct consequence of the underlying activation distributions. aFRR is dispatched continuously and its 15-minute activation magnitude is approximately log-normal, so a tail percentile (p95) yields a meaningful “high-magnitude” indicator. mFRR is event-driven: the manual reserve is dispatched in response

to specific imbalance events rather than continuously, and the resulting per-bin distribution is dominated by zeros.

The empirical numbers make this asymmetry visible. Across the full 2023–2025 sample, 79% of bins carry zero mFRR activation. The median is exactly zero. The 75th percentile is also exactly zero. The 90th percentile reaches 80 MW, and the 99th percentile reaches 300 MW. A hypothetical p75 cascade definition would therefore collapse to “any manual activation in the bin”, because the p75 mFRR threshold is itself zero. The p95 cascade definition used throughout the main text is the most informative pairing available: it combines a meaningful aFRR-side tail indicator with the rare-event mFRR-side indicator that the manual reserve already provides by construction.

The full sample-wide mFRR distribution decomposes as follows: 79% of bins at zero, 4% between 0 and 10 MW, 7% between 10 and 50 MW, 5% between 50 and 100 MW, 4% between 100 and 300 MW, and 1% above 300 MW. The maximum observed value is 1,245 MW. The distribution is highly skewed and tail-driven, and any cascade definition built on a mean-based or median-based mFRR threshold would collapse the manual-reserve side of the cascade into a near-degenerate flag. The “any mFRR activation” choice keeps the discriminating power of the manual-reserve series, without artificially restricting the cascade-flagged subset.

A.3 Policy counterfactual results

Section 4.4 sets out the four policy counterfactuals on the extreme-chasing tail in prose, and Figure 4.12 visualises the regime contrasts. The exact reward values and percentage changes behind that figure are collected in Table A.1 for direct lookup.

Table A.1: Per-bin reward (EUR/MWh) on the extreme-chasing tail under the status-quo and under each counterfactual instrument, broken down by regulatory regime. The extreme-chasing tail is the slice of chasing-active bins ($|\hat{S}_t| > 150$ MW) whose realised next-bin price also falls outside the within-regime $[p_1, p_{99}]$ window. Percentages give the relative change against the same regime’s status-quo. P1 clips the realised next-bin price to the within-regime $[p_5, p_{95}]$ window; P2 raises the α threshold from 150 to 250 MW; P3 charges a 100 EUR/MWh fee on bins with non-zero mFRR activation; P1+P3 applies both treatments at once. Static counterfactual with BRP positions held fixed (Section 4.4.4).

Regime	status-quo EUR/MWh	P1 (price kink)	P2 ($\alpha = 250$)	P3 (mFRR fee)	P1+P3 (combined)
Pre-PICASSO	444	225 (−49%)	529 (+19%)	373 (−16%)	155 (−65%)
PICASSO Adapted	496	156 (−69%)	549 (+11%)	420 (−15%)	80 (−84%)
MARI Adapted	493	114 (−77%)	549 (+11%)	419 (−15%)	40 (−92%)
Regime average	—	−65%	+14%	−15%	−80%

A.4 α -threshold counterfactual sweep

Section 4.3.2 reports the central pattern of the α -threshold sensitivity sweep on three illustrative metrics. Table A.2 provides the full numerical record underlying Figure 4.8 for direct lookup: chasing-active sample size, loss rate, mean per-bin reward, and cascade rate at the full-sample $p95$ aFRR threshold, computed for each combination of threshold $T \in \{50, 100, 150, 200, 300, 500\}$ MW and each regulatory regime. The $T = 150$ MW row reproduces the headline values of Sections 4.1 and 4.2; the remaining rows trace how the four metrics shift along the regulatory degree of freedom that Section 4.3.2 discusses.

Table A.2: Full α -threshold counterfactual sweep behind Figure 4.8. For every threshold T and every regime, the chasing-active sample is redefined as $|\hat{\mathbf{S}}_t| > T$ and the four metrics are computed from scratch. The bold row at $T = 150$ MW marks the current Belgian threshold and reproduces the headline values of Sections 4.1 and 4.2. Cascade rates use the full-sample $p95$ aFRR threshold of 94.6 MW.

T (MW)	Regime	n	Loss %	Mean reward (EUR/MWh)	Cascade % (p95)
50	Pre-PICASSO	38,796	29.7	+84.4	3.70
	PICASSO Adapted	10,331	41.6	+56.5	1.66
	MARI Adapted	9,976	40.3	+44.4	0.78
100	Pre-PICASSO	24,059	24.6	+111.3	5.52
	PICASSO Adapted	6,384	39.0	+73.7	2.46
	MARI Adapted	6,364	38.1	+57.8	1.15
150	Pre-PICASSO	14,868	20.0	+138.8	7.70
	PICASSO Adapted	3,827	36.0	+92.1	3.50
	MARI Adapted	4,046	34.5	+74.0	1.71
200	Pre-PICASSO	9,199	16.6	+164.9	9.92
	PICASSO Adapted	2,292	33.6	+111.7	4.67
	MARI Adapted	2,496	31.5	+91.6	2.44
300	Pre-PICASSO	3,683	11.6	+220.5	14.69
	PICASSO Adapted	795	23.3	+178.5	7.42
	MARI Adapted	851	22.9	+140.7	4.23
500	Pre-PICASSO	796	6.7	+340.1	22.36
	PICASSO Adapted	125	20.8	+324.1	12.80
	MARI Adapted	98	7.1	+368.5	18.37

UNIVERSITÉ CATHOLIQUE DE LOUVAIN
École polytechnique de Louvain

Rue Archimède, 1 bte L6.11.01, 1348 Louvain-la-Neuve, Belgique | www.uclouvain.be/epl