

Louvain School of Management

Comparison between a Fixed Reserve Requirement and an Operating Reserve Demand Curve for Ancillary Services Procurement

A stylised, Belgium-inspired comparison of reserve-procurement designs

Auteur·e(s) : Ludovic Leroy

Promoteur·rice(s) : Prof. dr. Bert Willems

Année académique 2025-2026

Travail de fin d'études (TFE) en vue d'obtenir le titre de

Master en sciences de gestion

Horaire de jour

Abstract

Operating reserves sit at the intersection of short-run system operation and long-run resource adequacy. How reserves are procured, and in particular whether they are demanded through a hard, administratively fixed requirement or through a value-based Operating Reserve Demand Curve (ORDC), shapes dispatch, reserve volumes, price formation, reliability and system cost. This thesis builds a stylised, Belgium-inspired energy-and-reserve co-optimisation model to compare these two designs, and reads the comparison through the theoretical tension between the Hogan private-good pricing view and the Joskow–Tirole public-good view of reserves. The ORDC used in the model is derived as an *inverse demand function for reserves*, obtained from the derivative of expected welfare with respect to the reserve margin, which recovers the LOLP · VOLL scarcity adder of Hogan while leaving room for the public-good component emphasised by Joskow–Tirole. The two designs are compared on dispatch, reserve volumes, energy and reserve prices, scarcity hours and operating cost (WP1), under sensitivities on net-load uncertainty, renewable penetration and the value of lost load (WP2); an extension traces the cost–reliability frontier across a small family of ORDC shapes (WP3), in the spirit of the calibration exercise of (Cartuyvels and Papavasiliou, 2022). The thesis explicitly flags as a boundary condition the technology-specific flexibility premium identified in (He and Willems, 2025), which no single-price reserve-auction design can reimburse on its own.

Preface

This thesis was prepared in the *Louvain School of Management* (LSM) under the supervision of Prof. dr. Bert Willems, to whom I am deeply grateful for his guidance and for the freedom he gave me in shaping the topic. I also thank the researchers whose work on operating-reserve economics forms the backbone of this document, in particular Anthony Papavasiliou and the members of his research group, whose Belgium-focused publications made this thesis possible. Finally, I thank my family and friends for their patience and support throughout this project.

Use of AI tools. Artificial-intelligence tools were used in the preparation of this document *solely* for spelling and grammar correction. The economic analysis, modelling, simulations, interpretation of the results and writing of the substantive content are entirely the author's own work.

Contents

List of abbreviations and symbols	III
List of Figures and Tables	V
1 Introduction	1
1.1 Background and motivation	1
1.2 The broad policy question and the thesis scope	1
1.3 What this thesis does	2
1.4 Structure of the thesis	3
2 Literature review	4
2.1 Reliability as an economic problem: public good and knife edge	5
2.2 Scarcity pricing and the Operating Reserve Demand Curve	7
2.3 Reserve sizing under uncertainty: from rules of thumb to probabilistic methods	9
2.4 Value of lost load, co-optimisation and the design space	10
2.5 Belgian ORDC case studies and the calibration question	10
2.6 European implementation constraints and complementary evidence	11
2.7 The flexibility-premium critique: a scope limit	12
2.8 Research gap and positioning of the thesis	13
3 Methodology	14
3.1 The model	14
3.2 Stylised Belgium-inspired test system	19
3.3 Baseline: co-optimisation with a fixed reserve requirement	19
3.4 Piecewise-linear implementation as a soft reserve constraint	20
3.5 Opportunity-cost channel between energy and reserves	21
3.6 The right marginal price for consumers	21
3.7 Cost recovery: fixed fee versus energy add-on	21
3.8 Scenario matrix and extension	22
4 Results and discussion	23
4.1 WP1 — Core mechanism comparison	23
4.2 WP2 — Sensitivity to the economic environment	29
4.3 WP3 — ORDC shape extension	35
4.4 Discussion	36

5 Conclusion	38
5.1 Answers to the four sub-questions	38
5.2 Answers to the work packages	39
5.3 Limitations	40
5.4 Future work	40
BIBLIOGRAPHY	41
A Model formulation in detail	44
A.1 Sets and indices	44
A.2 Parameters	45
A.3 Decision variables	45
A.4 Common constraints (both modes)	45
A.5 Mode A — fixed reserve requirement (Section 3.3)	46
A.6 Mode B — welfare-derivative ORDC (Sections 3.1–3.4)	46
A.7 Equilibrium / interpretive identities	47
A.8 Numerical implementation	47
B Input data and scenario tables	48
B.1 Thermal fleet	48
B.2 Demand and renewable profiles	48
B.3 Forecast-error distribution and reserve sizing	49
B.4 ORDC discretisation	49
B.5 Reference values and scenario grids	50
B.6 Realised-reliability evaluation	50
C Supplementary results and derivations	51
C.1 Detailed comparison of the first-order conditions with Hogan (2013) and Joskow and Tirole (2007)	51
C.2 Detailed WP1 consumer-side and welfare-decomposition analyses	53
C.3 The flexibility-premium critique: full exposition	55
C.4 WP2 renewable-share sweep — detailed results	58
C.5 WP2 VOLL sweep — detailed results	58
C.6 WP3 — ORDC shape extension	60

List of abbreviations and symbols

Abbreviations

ACE	Area Control Error
aFRR	automatic Frequency Restoration Reserve
BRP	Balance Responsible Party
BSP	Balancing Service Provider
CCGT	Combined-Cycle Gas Turbine
CREG	Commission for Electricity and Gas Regulation (Belgium)
CRM	Capacity Remuneration Mechanism
EB GL	Electricity Balancing Guideline (EU Reg. 2017/2195)
EENS	Expected Energy Not Served
ERCOT	Electric Reliability Council of Texas
FCR	Frequency Containment Reserve
LOLE	Loss of Load Expectation
LOLP	Loss of Load Probability
LP	Linear Program
mFRR	manual Frequency Restoration Reserve
MILP	Mixed-Integer Linear Program
MISO	Midcontinent Independent System Operator
ORDC	Operating Reserve Demand Curve
PJM	Pennsylvania–New Jersey–Maryland Interconnection
RES	Renewable Energy Sources
RQ	Research Question
TSO	Transmission System Operator
UC	Unit Commitment
VOLL	Value of Lost Load

Symbols

$p_{g,t}$	Energy output of unit g at time t (MW)
$r_{g,t}$	Upward reserve held by unit g at time t (MW)
s_t	Reserve shortfall at time t (MW)
R_t^{req}	Administratively fixed reserve requirement (MW)
R_t^{ref}	Reference reserve level used in the ORDC (MW)
λ_t	Energy clearing price at time t (€/MWh)
μ_t	Reserve clearing price at time t (€/MWh)
C_g	Marginal energy cost of unit g (€/MWh)
V	Value of lost load (€/MWh)
$\pi_t(x)$	Probability of reserve shortage larger than x at time t

List of Figures and Tables

List of Figures

4.1	Hourly dispatch on the baseline instance under the fixed-requirement design. The ORDC dispatch is identical on this calibration because no ORDC shortfall slab is used.	24
4.2	Reserve procurement on the baseline instance. Both mechanisms procure exactly the reference requirement in every hour; the ORDC is willing to pay enough at the margin that no shortfall slab is activated.	25
4.3	Energy and reserve prices on the baseline instance. The two mechanisms overlap hour by hour. The right panel shows the observed (μ_t, λ_t) pairs against the opportunity-cost lines $\lambda_t = \mu_t + C_g$	25
4.4	Planned scarcity, slack and ORDC-shortfall cost on the baseline instance. The value is zero in every hour for both mechanisms.	26
4.5	Knife-edge sweep at the peak hour. The fixed-rule reserve price jumps to VOLL when the target crosses the reservable capability; the ORDC price remains bounded by the welfare-derivative curve.	26
4.6	Reserve procurement on the stress instance. Both mechanisms hit the same reservable-capability ceiling during the daytime stress block; the fixed rule records this as slack, while the ORDC records it as use of shortfall slabs.	27
4.7	Energy and reserve prices on the stress instance ($\sigma_t \times 1.25$). The fixed-rule reserve price saturates at VOLL over the stress block; the ORDC replaces that jump with a bounded probability-weighted scarcity price.	27
4.8	Planned scarcity-cost decomposition on the stress instance. The fixed rule prices the unmet reserve requirement at VOLL; the ORDC prices the same reserve shortfall with probability-weighted shortfall slabs.	28
4.9	Price distributions on the stress instance. Energy prices remain similar; reserve prices reveal the mechanism difference.	28
4.10	WP2 σ -sweep: planned operating cost. Fixed-requirement and ORDC coincide for $\alpha \leq 1$ and diverge sharply once the joint reserve+capacity constraint binds at $\alpha \geq 1.25$. The fixed-requirement cost grows quadratically with α because every unmet MW of reserve is priced at the full VOLL; the ORDC grows almost linearly because the marginal welfare cost $V\pi_t(r)$ is bounded below the cap.	30

4.11	WP2 σ -sweep: 95th-percentile reserve clearing price. The fixed-requirement design clears at the VOLL cap as soon as the requirement breaches the reservable capability, illustrating the Joskow–Tirole knife-edge in price (Joskow and Tirole, 2007). The ORDC translates the same physical scarcity into a smooth marginal-welfare price bounded by the curve’s cap. . . .	31
4.12	WP2 σ -sweep: expected energy not served. The two curves overlap at every α . Realised reliability is set by the fleet’s reservable capability, not by the pricing rule; the wedge in Figure 4.10 is therefore a pricing wedge, not a quantity wedge.	32
4.13	WP2 public-good share sweep: 95th-percentile reserve price. Fixed-requirement and Tirole-ORDC are by construction invariant in V_{pub}/V because they price scarcity at the full V (or at $V\pi_t$ derived from it). The Hogan-ORDC flattens as V_{pub} grows because its heights scale with $V_{\text{priv}} = V(1 - V_{\text{pub}}/V)$; at $V_{\text{pub}}/V = 0.6$ the Hogan tail price is below 200 €/MWh, less than half the Tirole price.	33
4.14	WP2 public-good share sweep: planned operating cost. The Hogan-ORDC reports the lowest cost because it is the cheapest-to-clear design, but, as Figure 4.15 shows, that is a property of how the design measures cost, not of the social cost it delivers.	34
4.15	WP2 public-good share sweep: planned + realised expected cost (planned cost plus social VOLL applied to realised EENS). When the Monte-Carlo loss is priced at the <i>social</i> V , the Hogan-ORDC’s apparent saving in Figure 4.14 narrows. The remaining difference reflects the slack-price wedge (the part of scarcity the Hogan design under-prices), not a difference in physical reliability, since EENS is identical across the three designs at every V_{pub}/V	35
C.1	Consumer-facing reserve markup on the baseline instance. The unscaled add-on μ_t is identical under the two mechanisms because the underlying reserve prices are identical. . . .	53
C.2	Consumer-facing reserve markup on the stress instance. The fixed-rule add-on inherits the VOLL cliff; the ORDC add-on remains tied to $V\pi_t(r)$	54
C.3	Cost recovery on the baseline instance. Scheme A is an average uplift that exactly recovers reserve-capacity payments; scheme B is an unscaled volumetric add-on that transmits the marginal reserve scarcity signal.	55
C.4	Cost recovery on the stress instance. The fixed-rule VOLL spike dominates both the reserve-capacity payment and the raw volumetric add-on; the ORDC keeps both at the probability-weighted scarcity level.	56
C.5	Peak-hour ORDC decomposed into private-good and public-good value components. Both components use the same shortage-probability shape and add into the total inverse demand curve for reserves.	57
C.6	Peak-hour welfare-derivative ORDC, discretised into eight slabs between R_t^{ref} and $R_t^{\text{thr}} = 0.05R_t^{\text{ref}}$	58
C.7	WP2 renewable-share sweep: planned operating cost. Energy cost falls with β as renewable energy displaces thermal output, while the reserve requirement rises with the wind- and solar-side error contributions. For $\beta \leq 1$ the energy saving dominates; beyond $\beta = 1$ the joint reserve+capability constraint binds and the fixed-requirement design pays the same VOLL-cliff scarcity penalty as in the σ sweep.	59

C.8	WP2 renewable-share sweep: expected energy not served. EENS scales with β because the additional wind- and solar-side error widens the residual-load distribution.	60
C.9	WP2 VOLL sweep at $\alpha = 1.25$: planned operating cost. Fixed-requirement cost is exactly linear in V (slope = EENS $\cdot$ slack hours) because every unmet MW pays the full V . The ORDC cost rises sub-linearly because V only enters through the curve heights $V\pi_t(r)$, which are bounded by the survival probability π_t	61
C.10	WP2 VOLL sweep: 95th-percentile reserve price. The fixed-requirement clearing price is identically equal to V once the constraint binds; the ORDC clearing price rises in proportion to $V\pi_t(R^{\text{thr}})$, much more slowly.	62
C.11	WP2 VOLL sweep: expected energy not served. Invariant in V , because in the binding-capability regime V enters the operating problem only through the price of unmet reserve, not through its quantity.	63
C.12	WP3 cost–reliability frontier. Coloured dots are the 45 shape variants of the WP3 grid; colour encodes the cap price $\bar{p}$. Black-circled points are the Pareto-non-dominated subset (the WP3 frontier proper). The blue diamond is the welfare-derivative ORDC and the red cross the fixed-requirement design, both run on the same stress instance. The fixed-requirement design lies far off-frontier on the cost axis; the welfare-derivative ORDC sits inside the frontier at the EENS ≈ 402 MWh/day band.	64
C.13	WP3 shape grid grouped by reference level R^{ref} . Each R^{ref} level produces a near-flat EENS band with cost varying along the band as a function of the cap and threshold. The Pareto frontier of Figure C.12 therefore consists almost entirely of $R^{\text{ref}} = 1.15 R^{\text{req}}$ variants.	65
C.14	WP3 cost–reliability scatter coloured separately by the cap price $\bar{p}$, the threshold R^{thr} and the reference R^{ref} . Confirms that R^{ref} is the dominant lever (right panel separates the EENS bands cleanly), the cap $\bar{p}$ is the dominant cost lever within each band (left panel) and R^{thr} has near-zero effect on either dimension (middle panel).	65

List of Tables

4.1	WP1 headline metrics. Money figures are in €; prices are in €/MWh; LOLE is in hours and EENS in MWh.	29
B.1	Stylised thermal fleet parameters used in WP1–WP3.	48
B.2	Hourly load and renewable availability used in WP1 (rounded to the nearest MW).	49
B.3	Reference values of the economic parameters and the WP2 sensitivity grids.	50
B.4	WP3 ORDC shape grid (45 variants, run on the $\alpha = 1.25$ stress instance).	50

Chapter 1

Introduction

1.1 Background and motivation

Electricity systems must balance supply and demand continuously. Because demand and renewable generation are uncertain, the transmission system operator (TSO) sets aside dispatchable capacity, the *operating reserves*, to absorb forecast errors and contingencies. With variable renewable generation dominating new capacity additions in Europe, the value of being flexible rises, and how reserves are procured and how they are priced has become a first-order market-design problem (Hogan, 2013; Joskow and Tirole, 2007).

The puzzle is concrete. Two reserve-procurement designs coexist in the literature and in practice. The first, the European norm, is an administratively *fixed reserve requirement*: the TSO sets the minimum MW of each reserve product, and that level enters market clearing as a hard constraint. The second, the ERCOT/PJM benchmark and an active Belgian reform proposal, is the *Operating Reserve Demand Curve* (ORDC): reserves are valued by a downward-sloping demand curve whose height approximates the product of the loss-of-load probability (LOLP) and the value of lost load (VOLL) (CREG, 2021; Elia, 2020; Papavasiliou, 2020; Papavasiliou and Smeers, 2017). The two designs generate sharply different price signals, scarcity rents and consumer-facing marginal prices, even when they procure exactly the same physical reserve.

The stakes are not abstract. Two foundational papers disagree on what reserves *are*. Hogan (2013) reads them as a priced private good worth $LOLP \cdot VOLL$. Joskow and Tirole (2007) read them, in part, as a public good against systemic collapse that no individual consumer internalises (Joskow and Tirole, 2007, p. 81). Anyone arguing for an ORDC in Europe implicitly takes a position on this tension, often without saying so. The puzzle the thesis solves is whether the two readings can be reconciled inside a single reserve-demand rule, or whether they describe genuinely distinct economic objects.

1.2 The broad policy question and the thesis scope

The broad policy question. Behind the day-to-day debate over reserve products and clearing rules sits a larger question: *is the price of operating reserves the right place to encode the social value of reliability, and if so, how should that price be constructed?* Europe has historically said no: reliability is

a quantity target enforced by a hard reserve requirement. The United States, through ERCOT and PJM, has progressively said yes: reliability is a price, embedded in an ORDC whose height is approximately $\text{LOLP} \cdot \text{VOLL}$. Behind these two practical answers sit the two theoretical readings just introduced: Hogan’s private-good pricing view and Joskow–Tirole’s public-good view. The broad question is whether the two readings live inside one rule or describe genuinely distinct objects.

What this thesis does. The thesis isolates a controlled comparison: it rebuilds both designs inside one common energy-and-reserve co-optimisation on a stylised, Belgium-inspired test system, and compares them on dispatch, reserve procurement, energy and reserve price distributions, scarcity incidence, expected energy not served (EENS), and total operating cost. The ORDC is not borrowed as a calibration formula; it is derived in Chapter 3 as the *inverse demand function for reserves* obtained from the derivative of expected welfare with respect to the reserve margin. Two first-order optimality conditions, one with respect to dispatched energy Q_t and one with respect to reserve r_t , simultaneously pin down the reserve price $\mu_t = V \cdot \pi_t(r_t)$ and the energy price $\lambda_t = C_g + \mu_t$ at the dispatch margin. The closed-form height $V \cdot \pi_t(r_t) - \hat{M}C_t$ recovers the $\text{LOLP} \cdot \text{VOLL}$ adder of Hogan (2013) exactly, while the public-good content of Joskow and Tirole (2007) is absorbed into the calibration of $V = V_{\text{priv}} + V_{\text{pub}}$ and into the tail of the shortfall distribution.

What the thesis finds. The simulations support a qualified answer to the broad question: the Hogan and the Joskow–Tirole readings are *compatible at the level of equation form but distinguishable at the level of calibration*. In an abundant-reserve regime the two designs are observationally identical. In a stressed regime they procure the same physical reserve and reach the same EENS, but they price the same scarcity entirely differently: a VOLL-cliff under the fixed rule, a probability-weighted scarcity adder under the ORDC, with an order-of-magnitude wedge in planned cost. The wedge lives in the consumer-facing reserve add-on, not in physical reliability, and a Hogan-only calibration under-prices scarcity by exactly $V_{\text{pub}} \pi_t(r)$. Section 1.3 states these findings as the answers to WP1–WP3 with the specific numbers.

1.3 What this thesis does

The comparison is organised in three work packages, reported by the results chapter and answered by the conclusion in the same order. *WP1 (mechanism comparison)* compares the two designs side by side on a baseline 24-hour instance and on a $\sigma_t \times 1.25$ stress instance, reporting dispatch, reserve volumes, the joint distribution of (λ_t, μ_t) , scarcity hours, reliability and cost. On the baseline the two designs are observationally identical (same dispatch, prices and reliability; 5.114 M€ total cost); on the stress instance they procure the same physical reserve and reach the same EENS, but the fixed rule’s μ_t^{p95} saturates at VOLL (8500 €/MWh) while the ORDC tracks $V \pi_t(r)$ (519 €/MWh), producing a 32.93 M€ vs. 6.37 M€ planned-cost wedge. *WP2 (sensitivity)* sweeps the four economic ingredients of the welfare-derivative ORDC (σ , β , V , V_{pub}/V). EENS is identical between the two designs at every grid point because the binding constraint is the fleet’s reservable capability; the cost wedge scales with each of σ , β and V ; the public-good share enters only the calibration of V , so a Hogan-style ORDC built from

V_{priv} alone under-prices scarcity by exactly $V_{\text{pub}} \pi_t(r)$. *WP3 (ORDC shape extension)* traces the cost–reliability frontier across a 45-point $(R^{\text{ref}}, R^{\text{thr}}, \bar{p})$ grid on the stress instance, in the spirit of Cartuyvels and Papavasiliou (2022): any reasonable ORDC shape Pareto-dominates the fixed rule at equal reliability by roughly a factor of three on cost, and the welfare-derivative ORDC sits inside the empirical Pareto front at $\approx 7.8\%$ higher clearing cost than a flatter linear ramp of the same R^{ref} . The preference for the welfare-derivative form is therefore a preference for an analytical property (the explicit $V = V_{\text{priv}} + V_{\text{pub}}$ decomposition and the operational reconciliation between Hogan (2013) and Joskow and Tirole (2007)), not for a numerical cost optimum.

Scope limit. The comparison is between two single-price reserve-demand rules. As He and Willems (2025) show, no such rule reimburses the technology-specific flexibility premium flexible units earn from waiting to commit, which the thesis flags as a scope limit rather than studies.

1.4 Structure of the thesis

Chapter 2 reviews the literature in eight blocks, organised around the Hogan-vs-Joskow–Tirole compatibility question of Section 1.2: reliability economics as a public-good and knife-edge problem (Joskow and Tirole, 2007); scarcity pricing and the LOLP · VOLL rule (Hogan, 2013); reserve sizing under uncertainty; VOLL and co-optimisation; the Belgian case-study literature (Cartuyvels and Papavasiliou, 2022; Papavasiliou and Smeers, 2017); the European implementation debate (Papavasiliou, 2020); the flexibility-premium critique of He and Willems (2025); and the research gap and positioning of the thesis.

Chapter 3 presents the model first (Section 3.1): technology, welfare integrand, the two FOCs with respect to Q_t and r_t , the optimal reserve and energy quantities, and the explicit comparison of the resulting optimality conditions with those of Hogan (2013) and Joskow and Tirole (2007) in the thesis’s own notation. It then specialises the model to the Belgium-inspired test system (Section 3.2), states the fixed-requirement baseline LP (Section 3.3), implements the ORDC as a piecewise-linear soft constraint (Section 3.4), makes the opportunity-cost channel $\lambda_t - C_g = \mu_t$ explicit (Section 3.5), formalises the consumer-facing marginal price (Section 3.6), distinguishes the two cost-recovery schemes (Section 3.7), and lays out the scenario matrix for WP2 and WP3 (Section 3.8).

Chapter 4 reports WP1–WP3 on the stylised system. Chapter 5 answers WP1–WP3, discusses the limitations, and identifies future work.

Chapter 2

Literature review

This chapter reviews the literature that frames the comparison between a fixed reserve requirement and an Operating Reserve Demand Curve (ORDC). It is organised around the central question raised by the supervisor of this thesis and recovered analytically in Chapter 3: are the Hogan (2013) private-good pricing view and the Joskow and Tirole (2007) public-good view of operating reserves *compatible* inside a single reserve-demand rule, or are they two genuinely distinct objects? The review treats this as the organising thread, not as a peripheral observation, and reads the rest of the literature through it.

The chapter is in eight sections. Section 2.1 develops the public-good and knife-edge arguments of Joskow and Tirole (2007), which are the theoretical case for treating reserves as more than an administrative quantity and the theoretical critique of fixed reserve rules. Section 2.2 develops the scarcity-pricing and ORDC mechanism of Hogan (2013), derives its LOLP · VOLL pricing rule, and confronts it with the public-good view; this is where the *Hogan-vs-Joskow-Tirole compatibility* paragraph lives, and where the welfare-derivative construction of Chapter 3 is positioned as the constructive answer. Section 2.3 discusses reserve sizing under uncertainty, from heuristic rules of thumb to probabilistic and dynamic methods; the value of better reserve modelling rises with renewable uncertainty, which motivates the WP2 sensitivity grid. Section 2.4 covers the value of lost load and the co-optimisation of energy and reserves, including the opportunity-cost channel that equation (3.20) of Chapter 3 formalises. Section 2.5 reviews the two Belgium-focused ORDC case studies that directly inspire the thesis model, namely Papavasiliou and Smeers (2017) for the case-study construction and Cartuyvels and Papavasiliou (2022) for the shape-calibration question that drives WP3. Section 2.6 discusses the European implementation debate, where a day-ahead-only ORDC is shown to be institutionally incomplete (Papavasiliou, 2020), and reports complementary empirical and regulatory evidence. Section 2.7 elevates the working paper of He and Willems (2025), which the supervisor of this thesis presented as the natural extension of Joskow and Tirole (2007), and identifies the flexibility-premium limit that no single-price reserve-auction design, including the ORDC of this thesis, can reimburse on its own. Section 2.8 states the research gap and positions the contribution of this thesis against the three work packages WP1–WP3.

2.1 Reliability as an economic problem: public good and knife edge

The reference text for this thesis on the economics of reserves is Joskow and Tirole (2007). The paper studies optimal pricing and investment when a large share of retail consumers is price-insensitive and reliability must be managed through a combination of market design and rationing rules. Three of its results matter directly for the comparison this thesis carries out and for the welfare-based construction in Chapter 3.

Result 1: operating reserves are a public good. Joskow and Tirole (2007) distinguishes two types of shortage. Under *orderly rationing*, available generation keeps its full economic value V (the value of lost load), because load is curtailed in a controlled way and the social cost of the shortage equals V times the energy not served. Under *network collapse*, by contrast, the tripping of a single unit or transmission line cascades into the rest of the system and “available plants are almost valueless” (Joskow and Tirole, 2007, p. 78). Operating reserves are the instrument that prevents the second regime, and their benefit is system-wide rather than private to any one consumer. Network users take system reliability as exogenous to their own decisions and therefore underprovide reserves in any decentralised equilibrium: “the possibility of system collapses make operating reserves a public good” (Joskow and Tirole, 2007, p. 81). This is the first-order economic justification for an administrative reserve requirement existing at all, and the reason a comparison of a hard requirement with an ORDC must be framed as a comparison *between two ways of valuing a reliability good that no single consumer fully internalises*, not as a comparison of efficient against inefficient.

Result 2: three operating-reserve regimes. Proposition 8 of Joskow and Tirole (2007, p. 82) characterises the welfare optimum across three regimes as forecast demand grows. In the *off-peak* regime, capacity is slack and the optimal reserve coefficient is fixed at the upper level r_H defined by the first-order condition $h_i v / (1 + r_H)^2 = s$ (Joskow and Tirole, 2007, eq. 10, p. 79), so consumers pay $(1 + r_H)s$ for one MWh of served load (Joskow and Tirole, 2007, p. 81). In the *reserve-shedding* regime, generation capacity is binding, all load is still served, and reserves are gradually reduced as forecast demand rises; the regulated reserve price corresponds to a quasi-rent $h_i v / (1 + r_i)^2 - s$ (Joskow and Tirole, 2007, eq. 10, p. 79). In the *load-shedding* regime, reserves are pinned at the lower floor r_L defined by $[h_i / (1 - H_i)] (1 + r_L)^{-1} = 1$, and consumers pay $v / (1 + r_L)$ conditional on being served, with generators earning a quasi-rent $(1 - H_i) v / (1 + r_L) - s$ (Joskow and Tirole, 2007, p. 81). This is not yet an ORDC formula but it is its conceptual precursor: the marginal value of reserves rises with system tightness and reaches v only at the load-shedding boundary, which is the shape that the welfare-derivative ORDC of equation (3.11) reproduces. Importantly, Joskow and Tirole (2007) do not draw this object as a downward-sloping reserve demand curve and do not refer to it as an ORDC, since neither term appears in the paper; they treat the optimum as a quantity choice plus a regulator-set reserve price. The thesis reads their object as the welfare-derivative inverse demand function for reserves, which is exactly what the formal ORDC derivation of Chapter 3 obtains.

Result 3: the knife-edge effect of rigid reserve rules. Joskow and Tirole (2007, p. 81) observe that in the reserve-shedding region, “the supply curve and the total demand curve (energy plus reserves) are vertical and identical. Hence, a small mistake in the choice of reserve ratio creates wild swings in the market price (from $(1 + r_i)s$ to $v/(1 + r_i)$ conditionally on being served).” In the conclusion they restate this property as the *knife-edge problem* (Joskow and Tirole, 2007, p. 82): “the market price and the associated scarcity rents available to support investments in generating capacity are extremely sensitive to small mistakes or discretionary actions by the system operator.” Mechanically, in the reserve-shedding region the reserve supply and the total demand are both vertical, so the equilibrium price is indeterminate within a wide band between marginal cost and a near- v scarcity price; a tiny operator action picks anywhere in that band. This is the strongest theoretical argument against a fixed reserve requirement, and the property that the soft-constraint formulation of Chapter 3 (Section 3.4) is explicitly designed to remove. It also explains why the sensitivity analysis in WP2 stresses tail-price metrics rather than averages: the knife-edge lives in the tails of the price distribution.

Why both a private-good and a public-good reading are needed. The three results together imply that reserves carry *two* kinds of value in Joskow and Tirole (2007): a private-good marginal value at V per MW of avoided curtailment under orderly rationing, and a public-good externality value against systemic collapse that no individual consumer internalises. A pricing rule that captures only the first value, such as a pure LOLP · VOLL adder calibrated on individual outage costs, will undervalue reserves in the collapse-prone tail; a pure quantity rule such as a fixed requirement captures the second value administratively but creates the knife-edge instability of Result 3. This tension is the explicit motivation for the welfare-derivative construction in Chapter 3, which absorbs the public-good component into the calibration of V and into the tail of $\pi_r(\cdot)$ rather than into the form of the equation.

The Hogan energy-only complement. Hogan (2005) reaches a related conclusion from the opposite direction. If a market is to support resource adequacy without a separate capacity mechanism, scarcity pricing in energy and operating reserves must be credible and transparent; suppressed scarcity rents, whether through price caps or through informal operator intervention, create the classical *missing-money* problem because short-run market revenues no longer reflect the underlying reliability value. The paper argues, in particular, that a structured reserve demand curve is preferable to ad hoc reductions of reserve targets as the system tightens: the same scarcity information should be encoded in a price, not in a discretionary operator action. This is the conceptual bridge between Joskow–Tirole and the ORDC formulation of Hogan (2013) reviewed in the next section.

Taken together, Joskow and Tirole (2007) and Hogan (2005) establish the role of the fixed-requirement model in this thesis: a benchmark, not a first-best. The fixed-requirement model is retained because it solves a genuine public-good coordination problem; it is critiqued because, as a single-price-or-zero administrative rule, it generates the knife-edge instability and undervalues the smooth marginal reliability value that an ORDC encodes.

2.2 Scarcity pricing and the Operating Reserve Demand Curve

The ORDC mechanism is most cleanly formulated in Hogan (2013), which is the second foundational reference for this thesis and the constructive counterpart of Joskow and Tirole (2007). The argument is built in three steps that map directly onto the welfare-derivative construction of Chapter 3.

Step 1: the verbal LOLP · VOLL argument. Hogan (2013, p. 6) states the LOLP · VOLL result as economic intuition rather than as the conclusion of a formal optimisation: “Whenever there is involuntary load curtailment and the system has just the minimum of contingency operating reserves, then any increment of reserves would correspondingly reduce the load curtailment. Hence the price of operating reserves should be set at the value of lost load. At any other level of operating reserves... the value of an increment of operating reserves would be the same VOLL multiplied by the probability that net load would increase enough in the coming interval to reduce reserves to the minimum level where load would be curtailed... Hence the incremental value of operating reserves would be the analogous to the product of the loss of load probability (LOLP) and VOLL, or LOLP·VOLL.” The argument relies on a marginal-reasoning chain about contingency states rather than on a derivative of welfare; the explicit closed-form $V^R(r) = (\text{VOLL} - \widehat{\text{MC}}) \cdot \text{LOLP}(r)$ is given as Eq. 1 in Cartuyvels and Papavasiliou (2022), and the corresponding energy-price formula $\lambda = \text{MC}_{gm} + (\text{VOLL} - \widehat{\text{MC}}) \cdot \text{LOLP}(R_\Delta)$ is derived from KKT conditions in Papavasiliou and Smeers (2017, eq. 10). Equation (3.11) of Chapter 3 recovers the same closed form as the derivative of expected welfare with respect to the reserve margin, which makes the verbal argument of Hogan (2013) operational and the explicit formula of Cartuyvels and Papavasiliou (2022) explicit at the level of derivation rather than at the level of postulate.

Step 2: the curve fits naturally inside co-optimised dispatch. Hogan (2013, p. 7) argues that scarcity pricing belongs inside dispatch rather than as an ex-post overlay: “with simultaneous optimization in the economic dispatch, the scarcity prices attributed to operating reserves apply as well to energy whenever there is a tradeoff between energy dispatch and operating reserve capacity.” Hogan (2013) states this property verbally; the operational equation that makes it precise is derived from KKT stationarity of the co-optimised LP in Papavasiliou and Smeers (2017, eq. 7), which gives the arbitrage relation $\lambda_t - C_g = \mu_t$ at the dispatch margin and is reused in Section 3.5 of Chapter 3. Papavasiliou and Avila (2024) document the welfare gain from this co-optimisation against sequential clearing in a European setting.

Step 3: the curve helps resolve the missing-money problem without administrative cliffs. Hogan (2013, p. 3) restates the missing-money concern of Hogan (2005): “there is money missing in the whole-sale electricity market in the sense that average prices are not high enough to sustain new investment.” He attributes the gap to suppressed scarcity prices and absent demand bidding, and argues that the ORDC contributes “scarcity prices [that] would contribute to resolving the missing money problem for all generators actually providing energy or reserves” (Hogan, 2013, p. 7), with “higher scarcity prices [that] would lower the missing money challenge” (Hogan, 2013, p. 8). Because the ORDC delivers a non-zero scarcity adder in many hours rather than a very large adder in a handful of hours, it generates short-run rents over a broader set of operating states than a fixed reserve constraint, without relying on price caps

or on discretionary out-of-market operator interventions; the latter would re-introduce the knife-edge instability that Joskow and Tirole (2007) identifies. The thesis does not endogenise investment, so it cannot quantify the missing-money channel directly; what it can do is report short-run rents under both designs, which is one of the WP1 outputs.

Multi-product cascades. A practitioner-oriented extension by Hogan and Pope (2019) translates the same economic logic into a multi-product market design with day-ahead and real-time settlement layers and locational reserve needs. The contribution is mostly architectural for this thesis: it confirms that a stepped piecewise-linear ORDC is a standard way to embed scarcity pricing inside existing clearing engines, and that reserve-pricing choices in one settlement layer have consequences for others, an argument that becomes important in the European context (see Section 2.6).

Endogenous calibration from uncertainty distributions. Zhou and Botterud (2014) develop a *dynamic* reserve demand curve in which $\text{LOLP}(r)$ is computed from a combination of conventional-outage, load-forecast-error and wind-forecast-error distributions. Their co-optimised results show that a probabilistic reserve demand curve reduces extreme price spikes and reserve-price volatility compared with vertical reserve requirements, while still valuing reliability. This is one of the cleanest examples of an endogenous calibration of the curve and is the methodological precursor of the calibration logic used in WP3 (Cartuyvels and Papavasiliou, 2022). It also matters conceptually because it shows that the LOLP entering the ORDC need not be a fixed engineering parameter; it can be made an endogenous function of the operating state, which is what Chapter 3 does through the distribution f_t that governs $\pi_t(\cdot)$.

Reconciling Hogan and Joskow–Tirole inside a single reserve-demand rule. The two papers do not engage with each other directly: Joskow and Tirole (2007) do not refer to ORDCs or to Hogan, and Hogan (2013) does not use the “public good” or “Joskow–Tirole” labels. They nonetheless describe the same economic object from two different angles. The ORDC framework treats reserves, in form, as a priced good whose marginal value is $\text{LOLP} \cdot \text{VOLL}$; Joskow and Tirole (2007, p. 81) treats reserves, in part, as a public good whose full social value is not internalised by any individual willingness to pay. The two views are not straightforwardly nested. A Hogan-style ORDC calibrated purely from individual-outage LOLP and a private VOLL understates the protection reserves provide against *systemic* collapse, in which “available plants are almost valueless” (Joskow and Tirole, 2007, p. 78) and the relevant marginal cost of unserved energy is much higher than the orderly-rationing V . A purely public-good treatment, on the other hand, loses the price signal that makes reserves tradeable and re-introduces the knife-edge problem of fixed quantity rules. The thesis takes the position that the two views can be reconciled at the level of equation form, but only after the distinction is made operational at the level of calibration: the welfare-derivative ORDC of Chapter 3 is the constructive object on which this reconciliation rests, and the rest of this chapter is read through it.

Chapter 3 gives a constructive answer. The ORDC used in the thesis is derived as the inverse demand function obtained from the derivative of expected welfare with respect to the reserve margin (equation (3.11)). The functional form $V \cdot \pi_t(r_t) - \hat{\text{MC}}_t$ recovers the Hogan pricing rule *exactly*; the public-good content of Joskow and Tirole (2007) is absorbed into the *calibration* of V (interpreted as the social

VOLL, including the systemic-collapse externality) and into the tail of the shortfall distribution π_t (which steepens precisely at the reserve levels where systemic risk dominates). The Hogan and Joskow–Tirole readings are therefore *compatible at the level of equation form, but distinguishable at the level of calibration*; this is the precise sense in which the thesis treats them as two readings of one object rather than two competing objects. The flexibility-premium critique reviewed in Section 2.7 will identify the one ingredient that even this reconciled curve cannot capture without a multi-stage extension.

2.3 Reserve sizing under uncertainty: from rules of thumb to probabilistic methods

Before comparing a hard requirement with an ORDC, it is important to recognise that the *fixed-requirement family* is itself a design space. A review of probabilistic methods for defining reserve requirements (Dowell et al., 2016) maps the existing literature along two axes: the uncertainty representation (static historical heuristics versus probabilistic forecasts conditioned on system state) and the risk metric (LOLP, LOLE, EENS, or reserve-shortfall based). Their review highlights that rule-of-thumb reserve rules are only one point in a much wider space, and that economic reserve sizing ultimately requires combining a reliability metric with a monetary cost of unreliability, typically VOLL. This framing informs two choices in the thesis: the reliability proxy is made explicit and kept constant across scenarios, and VOLL is treated as a first-class sensitivity parameter.

The same question is approached from the stochastic-programming angle in (Papavasiliou et al., 2011). The two-stage framework commits reserves against representative wind-uncertainty scenarios and is shown to outperform deterministic reserve rules, including the classical 3+5 rule, particularly when wind penetration is high. The authors’ broader message is that the value of better reserve modelling rises with renewable uncertainty. That observation is directly actionable in this thesis: the sensitivity analysis explicitly varies both the forecast-error level and the share of renewable generation, and the model is expected to exhibit a similar property, namely that the gap between the fixed-requirement and ORDC designs should widen as uncertainty increases.

A useful intermediate design, which stays inside the administrative reserve-requirement family while borrowing from the probabilistic literature, is the dynamic dimensioning approach of (De Vos et al., 2019). The proof of concept on the Belgian system replaces the static annual sizing rule by a daily, risk-based requirement computed from an estimated next-day imbalance distribution. Dynamic sizing reduces the average reserve procurement without deteriorating the target reliability level and, importantly, adapts reserve levels to system conditions: more in high-risk hours, less in low-risk hours. For the thesis this has a modest but important implication: the fixed-requirement benchmark can be made more credible by allowing it to be time-varying rather than flat, and the qualitative contrast with the ORDC must be interpreted as a contrast between *administrative* reserve valuation (even if refined) and *market-based* reserve valuation.

2.4 Value of lost load, co-optimisation and the design space

Both the ORDC and any economic sizing rule depend critically on the value of lost load. Evidence in (Ovaere et al., 2019) shows that VOLL is substantially non-uniform across consumer types, times of day and interruption characteristics. Using detailed Norwegian, British and American data, they find that refined VOLL information can materially change reliability decisions, with up to a 43% reduction in operational cost in a five-node test system and 2–18% in a 118-node system, mainly through lower preventive redispatch when expected interruption costs are recognised as lower than the uniform benchmark would suggest. Two implications follow for this thesis. First, using a single VOLL is a simplifying assumption, not a neutral fact, and must be accompanied by a sensitivity analysis. Second, cost and reliability must be interpreted jointly: a change in VOLL often lowers cost while shifting who bears interruption risk, and an ORDC calibration that lowers cost at the price of higher shortfall is not necessarily an improvement.

Co-optimisation of energy and reserves is the other structural design choice on which the thesis takes a position. Results in (Papavasiliou and Avila, 2024) show, in a European-oriented setting, that jointly clearing energy and reserves avoids the opportunity-cost errors inherent in sequential clearing and generates sizeable welfare gains relative to the status quo. While their analysis is broader than the fixed-requirement versus ORDC question, it justifies one of the four methodological choices stated in Chapter 1: co-optimising energy and reserves in a single problem isolates the reserve-procurement question from the confounding effect of clearing architecture. Co-optimisation also makes explicit the *opportunity-cost channel* that links reserve and energy markets. A MW held as reserve is a MW not sold as energy, so in a joint clearing the reserve price μ_t and the energy price λ_t satisfy the complementarity that each flexible unit is paid, at the margin, the larger of its energy margin and its reserve margin. A fixed reserve requirement hides this link behind a discontinuous shadow price; an ORDC exposes it, which is why the same change in net load affects *both* markets when reserves are priced by a demand curve. This opportunity-cost channel is central to the interpretation of the price results in Chapter 4 and to the policy discussion of cost recovery in Chapter 5. The wider consequences of reserve-market design on energy prices are documented in (Frew et al., 2021), which shows in an ERCOT-like production-cost study that reserve eligibility rules, scarcity-pricing assumptions and reserve-quantity rules jointly shape both reserve and energy prices, and that the most visible effects typically live in the *tails* of the price distribution rather than in averages. That motivates the reporting strategy used in the results chapter, namely price distributions, scarcity hours and tail metrics alongside aggregate cost.

2.5 Belgian ORDC case studies and the calibration question

The Belgium-focused literature is the closest match to the thesis setting and supplies the case-study inspiration. A key Belgian case study (Papavasiliou and Smeers, 2017) builds a bottom-up model of the Belgian electricity market, validates it against 21 months of historical day-ahead prices and quantities, and uses it to estimate ORDC scarcity adders of the LOLP times VOLL form, with extensions for reserve products with different response times. Three elements of their analysis are directly reused here. First, reserve availability is computed from committed headroom, ramping capability and response time, rather

than from installed capacity, otherwise the comparison with an ORDC is not credible. Second, the ORDC adder is explicitly linked to reserve scarcity and VOLL, which grounds the piecewise-linear formulation used in Chapter 3. Third, the authors document that under realistic Belgian conditions, CCGTs were not able to recover their investment costs, while an ORDC-based scarcity adder improved their economic viability. Even though investment is not endogenised in this thesis, that result gives the interpretation of higher scarcity rents under ORDC a flexibility-remuneration flavour in the discussion chapter.

A follow-on study (Cartuyvels and Papavasiliou, 2022) extends the Belgian case study in the direction that is most relevant for the extension of this thesis. They take the $V^R(r) = (\text{VOLL} - \widehat{\text{MC}}) \cdot \text{LOLP}(r)$ ORDC formula (Cartuyvels and Papavasiliou, 2022, Eq. 1) and study how the *shape* and *calibration* of the curve, namely VOLL (8300 vs 13500 €/MWh), pre- vs post-activation framing, and independent vs correlated imbalance increments, affect the cost–reliability trade-off. With these three binary choices they compare $2 \times 2 \times 2 = 8$ ORDC variants on a closed-loop system-operation simulator calibrated for Belgium. Two findings are particularly informative. First, total flexible operating cost varies only modestly across variants: Cartuyvels and Papavasiliou (2022, p. 9) report a daily range of 1,683 to 1,697 thousand € across the eight variants, “up to 14 thousand € per day... which can be considered quite stable”. Scarcity adders, by contrast, move much more: in the reference scenario the fast reserve adder ranges from 2.74 to 6.50 €/MWh across the eight variants (Cartuyvels and Papavasiliou, 2022, Tables III–IV), and the slow adder from 0.14 to 0.56 €/MWh. Second, the ranking of ORDC variants is system-dependent: there is no universally optimal shape, only trade-offs that a particular system makes transparent. This observation sets the boundary condition of the thesis extension: a small family of ORDC shapes is run on a single stylised system, and the conclusion is framed as a cost–reliability frontier for that system, not as a universal recommendation. The numerical pattern *flat in cost, much wider in adders* is also the structural prediction that the WP3 results in Chapter 4 are designed to test on the thesis simulator.

2.6 European implementation constraints and complementary evidence

A central message of the Belgian literature is that the ORDC is economically appealing in a European context but institutionally incomplete. One paper argues that for ORDC logic to translate into correct forward-market incentives, the value of scarce reserves must *back-propagate* from real time to day-ahead reserve and energy markets (Papavasiliou, 2020). This back-propagation requires an explicit real-time market for reserve capacity, which does not currently exist in Europe. A day-ahead-only ORDC, or a scarcity-pricing overlay applied only in imbalance settlement, affects equilibrium prices but dilutes or cancels the pay-for-performance logic that justifies the mechanism in the first place. This reference is used in the thesis to delimit scope: the model captures the operational logic of reserve valuation, not the full European back-propagation problem.

Two Belgian institutional documents refine this point. An Elia report (Elia, 2020) distinguishes the CORE-style ideal design from what is implementable in the short to medium term, arguing that full alignment with European platforms and existing legal constraints would likely require harmonised European-wide implementation, and develops an alternative, more imbalance-oriented scarcity proposal as a second-best. A CREG study (CREG, 2021) synthesises the Belgian evidence, finds that well-

calibrated scarcity adders are near zero in periods of abundant capacity, a useful validation property for the thesis model, and ranks design variants by their back-propagation properties, concluding that introducing a real-time market for reserve capacity has a larger effect than virtual trading or day-ahead co-optimisation alone. A related regulator note (CREG, 2020) provides the compact regulator-facing framing that scarcity pricing is a reserve-scarcity mechanism, not a generic imbalance penalty, and that it can coexist with a capacity mechanism. The European legal frame itself (European Commission, 2017) already contains the terminology of co-optimised and market-based allocation for balancing capacity, which places the thesis model inside the vocabulary of EU Regulation 2017/2195 without claiming to reproduce its institutional detail.

Complementary evidence outside Belgium sharpens the interpretation of the results. Evidence from ERCOT (Bajo-Buenestado, 2021) shows that realised ORDC prices are significantly reduced by high wind generation and can approach zero when wind output is high enough, an empirical warning that ORDC revenues may weaken precisely when renewable-driven flexibility needs are highest, and that the short-run ORDC is not, by itself, a complete long-run adequacy instrument. Results for PJM (Lavin et al., 2020) demonstrate that conditioning forced-outage probabilities on ambient temperature produces a dynamic ORDC that dominates the unconditional version in welfare terms, with most of the gain concentrated in stressed hours. A recent MISO white paper (MISO, 2024) provides a contemporary operator-side calibration benchmark: ORDCs in practice are multi-parameter objects with distinct Pricing VOLL, System VOLL, floors, caps and emergency circuit breakers. Finally, a recent ERCOT market-monitor report (Potomac Economics, 2025) reminds that average annual ORDC effects can be small while their operational importance is concentrated in a handful of critical hours, supporting the thesis decision to report tail metrics rather than only averages.

2.7 The flexibility-premium critique: a scope limit

The recent working paper of He and Willems (2025), which the supervisor of this thesis introduced as the natural extension of Joskow and Tirole (2007), sets an *upper bound* on what any single-price reserve-auction design can achieve. In a three-stage peak-load-pricing model with asymmetric adjustment costs they identify a technology-specific flexibility premium $v_{\theta,\omega}$, the option value flexible units earn from *waiting* to commit, and prove that no single-price reserve auction (forward-only, uniform-price, energy-reserve co-optimised or non-linear scoring) reimburses it. The only market-based design that does is a *menu of call-option contracts* indexed by strike prices that match the adjustment costs of the available technologies (their Proposition 4). For the thesis this means that the comparison between a fixed reserve requirement and an ORDC is a comparison *within* the class of single-price reserve-demand rules that He and Willems (2025) flag as structurally incomplete for long-run investment in flexible technology: the operational reconciliation between Hogan and Joskow–Tirole the thesis achieves is therefore *bounded above* by their result, and any conclusion about investment incentives is outside scope. The full three-stage construction and the proof sketch are reproduced in Appendix C.3 so the rest of the literature review can stay focused on the operational pricing question.

2.8 Research gap and positioning of the thesis

Taken together, the literature supports five conclusions and leaves one clear opening.

First, reliability is an economic problem and operating reserves are a public good whose marginal value rises smoothly with system tightness; rigid reserve targets generate a knife-edge instability that no operationally simple administrative rule can fully avoid (Hogan, 2005; Joskow and Tirole, 2007).

Second, the ORDC translates the marginal-reliability-value logic into a smooth LOLP · VOLL pricing rule that fits naturally inside co-optimised energy-and-reserve dispatch (Hogan, 2013; Hogan and Pope, 2019; Papavasiliou and Avila, 2024; Zhou and Botterud, 2014), and it does so in a form that is, on inspection, the derivative of expected welfare with respect to the reserve margin (equation (3.11)). The two foundational papers (Hogan, 2013; Joskow and Tirole, 2007) are therefore not first-order contradictory: they describe the same inverse demand function for reserves under different calibrations of V and π_r .

Third, Belgium-focused studies show that the ORDC is feasible economically and instructive operationally (Cartuyvels and Papavasiliou, 2022; Papavasiliou and Smeers, 2017), but institutionally incomplete in Europe because the back-propagation of scarcity value from real time to forward markets requires a real-time reserve-capacity market that does not currently exist (CREG, 2020, 2021; Elia, 2020; Papavasiliou, 2020).

Fourth, the calibration of the ORDC, not only the choice of ORDC versus fixed requirement, materially affects prices, scarcity hours and reliability, in ways that are system-dependent and not captured by mean cost alone (Bajo-Buenestado, 2021; Cartuyvels and Papavasiliou, 2022; De Vos et al., 2019; Dowell et al., 2016; Frew et al., 2021; Lavin et al., 2020; MISO, 2024; Ovaere et al., 2019; Papavasiliou et al., 2011; Potomac Economics, 2025).

Fifth, even a well-calibrated ORDC is structurally incomplete for the long-run remuneration of flexible technology: the technology-specific flexibility premium identified by He and Willems (2025) cannot be reimbursed by any single-price reserve-auction design, including the ORDC of this thesis. This is the explicit upper bound on what an operational reserve-pricing rule can achieve.

What this thesis adds. What the literature does not provide is a transparent, pedagogical, side-by-side comparison of the fixed reserve requirement and the ORDC on a *single* common stylised system, with the welfare-derivative construction of Chapter 3 as the analytical bridge between Hogan (2013) and Joskow and Tirole (2007). The contribution is organised in three work packages aligned with Chapter 1: WP1 isolates the reserve-demand mechanism inside one optimisation problem; WP2 stress-tests the comparison along the three structural ingredients of (3.11) (uncertainty, renewable share, VOLL); WP3 traces a cost–reliability frontier across a small family of ORDC shapes in the spirit of Cartuyvels and Papavasiliou (2022). Beyond the operational comparison, the thesis reads the two designs through the theoretical tension between Hogan (2013) and Joskow and Tirole (2007), answers the supervisor’s organising question (*are the two views compatible?*) by showing that they correspond to different calibrations of the same inverse demand function, and flags the flexibility-premium limit of He and Willems (2025) as the boundary beyond which the short-run comparison cannot speak.

Chapter 3

Methodology

This chapter presents the model first, then specialises it to a test system and to two competing linear-programme implementations. Section 3.1 states the model in full: technology, welfare integrand, the two first-order optimality conditions with respect to dispatched energy Q_t and reserve r_t , the resulting optimal reserve and energy quantities, and the comparison of these optimality conditions with those of Hogan (2013) and Joskow and Tirole (2007) in the thesis's own notation. Sections 3.2–3.3 then specialise the model to the Belgium-inspired test system and to the baseline LP under a hard reserve requirement. Section 3.4 implements the welfare-derivative ORDC as a piecewise-linear soft reserve constraint. Sections 3.5–3.7 spell out the three operational consequences (the opportunity-cost channel $\lambda_t - C_g = \mu_t$, the consumer-facing marginal price, and the two cost-recovery schemes). Section 3.8 closes with the scenario matrix for WP2 and WP3.

3.1 The model

The model is a single-period welfare maximisation over dispatched energy Q_t and upward reserve r_t , with stochastic net-load shocks absorbed through expected curtailment. Its construction is structurally the welfare- maximisation formulation of Joskow and Tirole (2007, Section 4(a), eqs. 10–12), specialised to a co-optimised dispatch, and its closed form coincides with the operational ORDC formula $V^R(r) = (\text{VOLL} - \widehat{\text{MC}}) \cdot \text{LOLP}(r)$ of Cartuyvels and Papavasiliou (2022, Eq. 1) and the scarcity-adder formula of Papavasiliou and Smeers (2017, eq. 10).

3.1.1 Technology, demand and decisions

Supply. The supply side consists of a small set of thermal unit groups $g \in \mathcal{G}$ and a renewable aggregate. Each thermal group g has a constant marginal energy cost C_g (€/MWh), an installed capacity $\bar{P}_g$ (MW), a technical minimum $\underline{P}_g$ (MW), a ramp rate RR_g (MW/h), and an on/off status $u_{g,t} \in \{0, 1\}$ which is held at $u_{g,t} = 1$ in the LP relaxation used throughout the thesis. The renewable aggregate has time-varying available capacity W_t (wind) and S_t (solar), dispatched as a price-taker up to availability.

Demand. Net demand D_t is treated as inelastic and is hit by an additive forecast error ε_t with cumulative distribution $F_t(\cdot)$ and density $f_t(\cdot)$. The *complementary distribution*

$$\pi_t(x) \equiv \Pr(\varepsilon_t > x) = 1 - F_t(x), \quad (3.1)$$

i.e. the probability of a net-load shock larger than x , plays the role of the *loss-of-load probability at reserve level x* throughout the chapter. A single upward reserve product suffices, in the spirit of the simplification used in Papavasiliou and Smeers (2017); multi-product cascades (Hogan and Pope, 2019) and downward reserves are deferred to the conclusion.

Decisions. For each period t the social planner chooses the energy dispatch $p_{g,t} \geq 0$ and the upward reserve $r_{g,t} \geq 0$ of every unit, with aggregates $Q_t \equiv \sum_g p_{g,t}$ and $r_t \equiv \sum_g r_{g,t}$, and the involuntary load shed $LS_t \geq 0$. A unit's reserve cannot exceed either its committed headroom or its ramping capability over the response time Δ of the product:

$$r_{g,t} \leq \min(\bar{P}_g - p_{g,t}, \Delta RR_g) u_{g,t}. \quad (3.2)$$

Following Papavasiliou and Smeers (2017), (3.2) is what makes the comparison between a fixed requirement and an ORDC credible: a unit whose dispatch leaves no headroom, or whose ramp is too slow to respond within Δ , cannot meaningfully provide reserve, regardless of installed MW.

3.1.2 Welfare integrand

Period- t social welfare W_t has four components, in the same decomposition as Joskow and Tirole (2007, Section 4(a)):

$$W_t(Q_t, r_t) = \underbrace{B_t(Q_t)}_{\text{consumer benefit}} - \underbrace{C(Q_t)}_{\text{production cost}} - \underbrace{K(r_t, Q_t)}_{\text{cost of holding reserve}} - \underbrace{L_t(r_t)}_{\text{expected curtailment cost}}. \quad (3.3)$$

The notation is read as follows.

Consumer benefit $B_t(Q_t)$. The gross benefit from served energy. Because demand is inelastic, $B'_t(Q_t) = V$ on the relevant range; at the optimum $B'_t(Q_t) = \lambda_t$, the energy clearing price.

Production cost $C(Q_t)$. The aggregate $C(Q_t) = \sum_g C_g p_{g,t}$, with $C'(Q_t) = C_{gm}$ the marginal cost of the marginal unit.

Reserve commitment cost $K(r_t, Q_t)$. The additional cost of holding capacity in reserve rather than delivering energy. The partial $\partial K / \partial r_t$ is the energy margin a marginal flexible unit gives up by standing by, denoted $\hat{MC}_t$. The cross-partial $\partial K / \partial Q_t$ is the opportunity cost of producing one more MWh when capacity is shared: in the LP implementation it equals the shadow price of the aggregate capacity constraint $Q_t + r_t \leq \sum_g \bar{P}_g$, i.e. the reserve scarcity rent μ_t entering the energy-price equation below.

Expected curtailment cost $L_t(r_t)$. Conditional on a net-load shock ε_t , the system is short by $\max(\varepsilon_t - r_t, 0)$ MW. Under orderly rationing the unit social cost is V (Joskow and Tirole, 2007); under network collapse it is higher, since “available plants are almost valueless” (Joskow and Tirole, 2007, p. 78). Writing V for the *social* VOLL (encompassing both private WTP and the systemic-collapse externality) gives

$$L_t(r_t) = V \int_{r_t}^{\infty} (x - r_t) f_t(x) dx, \quad (3.4)$$

the expected loss-of-load energy multiplied by the social VOLL; its derivative is the channel through which the public-good content of Joskow and Tirole (2007) enters the optimality conditions.

3.1.3 First-order optimality conditions

The social planner chooses (Q_t, r_t, LS_t) subject to two binding constraints: the *energy balance* (production plus any involuntary load shed must cover demand)

$$Q_t + LS_t \geq D_t, \quad (3.5)$$

which is the operational version of the supervisor’s framing that *consumption cannot exceed production*; and the *aggregate capacity*

$$Q_t + r_t \leq \bar{K}_t \equiv \sum_g \bar{P}_g u_{g,t}. \quad (3.6)$$

The two multipliers are identified with the two clearing prices. $\lambda_t \geq 0$ on (3.5) is the *energy clearing price*, since the marginal value of one more MWh of available energy is exactly the shadow price of that constraint. $\mu_t \geq 0$ on (3.6) is the *reserve clearing price*, since reserve and energy compete for the same capacity at the margin and the shadow price of the capacity constraint is what a flexible unit must be paid to hold back one more MW from energy and offer it as reserve; equivalently, $\partial K / \partial Q_t = \mu_t$ at the binding capacity constraint.

The cost-minimisation dual of $\max W_t(Q_t, r_t)$ (after absorbing the consumer benefit $B_t(Q_t) = V Q_t$ as a constant on the served-energy range) is the programme the LP actually solves. Its Lagrangian is

$$\mathcal{L} = C(Q_t) + K(r_t, Q_t) + L_t(r_t) - \lambda_t (Q_t + LS_t - D_t) + \mu_t (Q_t + r_t - \bar{K}_t), \quad (3.7)$$

and its two KKT stationarity conditions are the first-order optimality conditions stated below.

FOC with respect to Q_t . Differentiating $\mathcal{L}$ with respect to Q_t gives the energy-price equation

$$\lambda_t = C'(Q_t) + \frac{\partial K}{\partial Q_t} = C_{gm} + \mu_t, \quad (3.8)$$

where $C_{gm} \equiv C'(Q_t)$ is the marginal cost of the marginal generator. The supplementary KKT condition $V - \lambda_t \geq 0$, with equality whenever $LS_t > 0$, caps the energy price at the social VOLL at scarcity. Read as a wedge, $\lambda_t - C_{gm} = \mu_t$ is the strictly positive gap between the energy clearing price and the marginal production cost whenever the reserve constraint binds. This is the opportunity-cost arbitrage relation derived as eq. 7 in Papavasiliou and Smeers (2017), the *simultaneous optimisation* verbal channel of Hogan (2013, p. 3), and the formal version of the supervisor’s observation that energy and reserves are *complementary products with an opportunity cost between them*. The wedge is reused as equation (3.20) in Section 3.5.

FOC with respect to r_t . Differentiating $\mathcal{L}$ with respect to r_t gives $\partial K/\partial r_t + L'_t(r_t) + \mu_t = 0$. Writing $\hat{MC}_t \equiv \partial K/\partial r_t$ (the net commitment cost per MW held in reserve, i.e. the energy margin a marginal flexible unit gives up by standing by) and applying Leibniz's rule to (3.4) gives

$$-\frac{\partial L_t}{\partial r_t} = V \int_{r_t}^{\infty} f_t(x) dx = V \pi_t(r_t), \quad (3.9)$$

the same elementary derivative as the one behind Joskow and Tirole (2007, eq. 10). The second optimality condition therefore reads

$$V \pi_t(r_t) = \hat{MC}_t + \mu_t, \quad (3.10)$$

or, rearranging,

$$\text{ORDC}_t(r_t) \equiv \left. \frac{\partial W_t}{\partial r_t} \right|_{\mu_t=0} = V \cdot \pi_t(r_t) - \hat{MC}_t, \quad (3.11)$$

where the boxed equation is the central equation of the chapter. Read as a function of r_t , it traces an *inverse demand function for reserves*: the height of the curve at any r_t is the social willingness to pay for one more MW of reserve at that level, net of the commitment cost. The optimal reserve at any prevailing market reserve price μ_t is the level r_t^* at which the height of the curve clears against μ_t , i.e. $V \pi_t(r_t^*) - \hat{MC}_t = \mu_t$.

3.1.4 Optimal Q_t^* , r_t^* and the closed-form energy-price expression

Combining the two FOCs (3.8)–(3.10) delivers two closed-form expressions that organise the rest of the chapter.

Optimal reserve r_t^* . From (3.10),

$$V \pi_t(r_t^*) = \hat{MC}_t + \mu_t, \quad (3.12)$$

i.e. r_t^* equates the *marginal value of reserve* $V \pi_t(r_t^*)$ to its *marginal cost* (commitment cost plus capacity rent). With the simulator's calibration $\hat{MC}_t = 0$, the two limits are clean: at abundant reserve ($\pi_t \rightarrow 0$) the reserve price $\mu_t \rightarrow 0$ and the capacity constraint slackens; at the load-shedding boundary ($\pi_t \rightarrow 1$) $\mu_t \rightarrow V$, exactly Hogan's "price-of-reserves-equals-VOLL" rule. With $\hat{MC}_t > 0$ the same limits hold up to the constant correction.

Optimal energy quantity Q_t^* and energy price. Substituting $\mu_t = V \pi_t(r_t^*) - \hat{MC}_t$ from (3.10) into (3.8) yields

$$\lambda_t = C_{gm} + V \pi_t(r_t^*) - \hat{MC}_t, \quad (3.13)$$

i.e. the energy price is the marginal cost of the marginal generator plus the welfare-derivative reserve adder $V \pi_t(r_t^*) - \hat{MC}_t$, which vanishes at abundant reserve and reaches $V - \hat{MC}_t$ at the load-shedding boundary. With $\hat{MC}_t = 0$ in the simulator (3.13) coincides with the operational $\lambda = MC_{gm} + (V - \hat{MC})\text{LOLP}(R_\Delta)$ of Papavasiliou and Smeers (2017, eq. 10) and Cartuyvels and Papavasiliou (2022, Eq. 1).

Compatibility with Hogan (2013) and Joskow and Tirole (2007). The two FOCs (3.8)–(3.10) can be mapped explicitly onto the optimality conditions of Hogan (2013) and Joskow and Tirole (2007): Hogan’s verbal LOLP · VOLL rule is the welfare-derivative ORDC (3.11) at the load-shedding boundary; Joskow and Tirole’s eq. 10 marginal condition $h_{iv}/(1 + r_H)^2 = s$ is, after relabelling, exactly $V\pi_t(r_H^*) = \hat{M}C_t + \mu_t$ of (3.10), and their three reserve regimes map onto the abundant, intermediate and load-shedding regions of the welfare-derivative curve. The detailed point-by-point comparison is given in Appendix C.1. The substantive difference is therefore not in the optimality condition itself, which is the same FOC under different notation, but in what each author does with it: Hogan states the price rule without derivation; J-T derive the FOC but leave the intermediate regime as a knife-edge in price; the thesis takes the same FOC and prices the intermediate regime continuously along the welfare-derivative curve.

3.1.5 Orderly rationing versus systemic blackout: where the public-good content lives

Joskow and Tirole (2007) distinguish two qualitatively different ways in which a reserve shortfall can translate into a social cost. In *orderly rationing*, available generation keeps its full economic value and load is curtailed in a controlled way; the social cost of one MWh of unserved energy is V . In *network collapse*, the tripping of a single unit or transmission line cascades into the rest of the system and “available plants are almost valueless” (Joskow and Tirole, 2007, p. 78). Operating reserves are the instrument that prevents the second regime, and their benefit is system-wide rather than private to any one consumer (Joskow and Tirole, 2007, p. 81).

The thesis’s welfare-derivative ORDC absorbs both cases into the *calibration* of V . Writing the social VOLL as the additive decomposition

$$V = V_{\text{priv}} + V_{\text{pub}}, \quad (3.14)$$

where V_{priv} is the orderly-rationing private value of served energy that any individual curtailed consumer is willing to pay (Hogan, 2013) and V_{pub} is the systemic-collapse externality stressed by Joskow and Tirole (2007) that no individual consumer internalises, equation (3.11) splits into a Hogan component and a Joskow–Tirole component:

$$\text{ORDC}_t(r_t) = \underbrace{V_{\text{priv}} \cdot \pi_t(r_t)}_{\text{Hogan, private-good}} + \underbrace{V_{\text{pub}} \cdot \pi_t(r_t)}_{\text{Joskow–Tirole, public-good}} - \hat{M}C_t. \quad (3.15)$$

Both components share the same loss-probability shape $\pi_t(r_t)$; the public-good piece does not enter the *form* of the curve, only its *height*. The two readings are therefore *compatible at the level of equation form but distinguishable at the level of calibration*: a regulator who calibrates V to the private WTP V_{priv} implements a Hogan-style ORDC, while one who calibrates V to the social VOLL $V_{\text{priv}} + V_{\text{pub}}$ implements a Joskow–Tirole-aware ORDC. The dispatch implication, confirmed in Section 4.2.4 of Chapter 4, is that in capability-constrained states both calibrations procure the same physical reserve and reach the same EENS, but the Hogan calibration under-prices reserve scarcity by exactly $V_{\text{pub}}\pi_t(r_t)$, a wedge that lives entirely in the consumer-facing reserve add-on of Section 3.7. The simulator uses $V_{\text{pub}}/V = 0.20$ as the reference calibration and sweeps $V_{\text{pub}}/V \in \{0, 0.1, 0.2, 0.4, 0.6\}$ as a WP2 sensitivity, with the corner

$V_{\text{pub}}/V = 0$ recovering the pure Hogan reading and $V_{\text{pub}}/V = 0.6$ approximating a strong systemic-collapse calibration. A state-conditional $\pi_t(\cdot)$ in the spirit of Lavin et al. (2020) that distinguishes rolling from systemic blackouts through the *shape* of the shortfall distribution is flagged as future work.

Boundary of what the welfare-derivative ORDC captures. He and Willems (2025) show that, in the absence of an efficient real-time market, no single-price reserve-auction design can reimburse the technology-specific flexibility premium that flexible units earn from waiting to commit. That premium requires a multi-stage call-option menu and cannot live inside the height of $\text{ORDC}_t(\cdot)$ alone. The welfare-derivative ORDC of this thesis is therefore correctly read as a value-based reserve *procurement* rule for the operational layer; it does not, on its own, deliver the long-run investment optimum.

3.2 Stylised Belgium-inspired test system

The rest of the chapter specialises the model to a stylised, Belgium-inspired test system, following the case-study logic of Papavasiliou and Smeers (2017) and Cartuyvels and Papavasiliou (2022). Each simulation covers $T = 24$ hourly periods. The thermal fleet has three unit groups (baseload, CCGT mid-merit, OCGT peaker) with marginal costs, capacities, technical minima, ramp rates and reserve-capability products ΔRR_g calibrated to the Belgian system documented in Papavasiliou and Smeers (2017); the numbers are in Appendix B.1. The renewable aggregate is 4000 MW of wind and 6000 MW of solar, dispatched as price-takers (Appendix B.2). Load is a stylised winter-weekday Belgian profile normalised to a peak of 12000 MW.

The additive net-load forecast error $\varepsilon_t \sim \mathcal{N}(0, \sigma_t^2)$ has a time-varying standard deviation built as the quadrature sum of load, wind and solar contributions (De Vos et al., 2019; Dowell et al., 2016; Papavasiliou et al., 2011); the formula is (B.1) in Appendix B.3. The fixed reserve requirement used both as the WP1 baseline and as the ORDC reference is the probabilistic dynamic-dimensioning rule of De Vos et al. (2019), $R_t^{\text{req}} = z_{0.97} \sigma_t$ with $z_{0.97} \simeq 1.881$ ((B.2)). The reserve response time is $\Delta = 0.25$ h (mFRR-style), and the reference social VOLL is $V = 8500$ €/MWh with $V_{\text{pub}}/V = 0.20$. Full scenario grids are in Appendix B.

3.3 Baseline: co-optimisation with a fixed reserve requirement

The baseline model jointly clears energy and reserves in a single linear programme, in the spirit of Papavasiliou and Avila (2024). The objective minimises the total expected variable cost,

$$\min \sum_t \left[\sum_g C_g p_{g,t} + V \mathbb{E}[LS_t] + V \eta_t \right], \quad (3.16)$$

where V is the value of lost load, $\mathbb{E}[LS_t]$ is the expected involuntary curtailment in period t , and η_t is the slack on the hard reserve requirement (priced at V as the implicit cap of the vertical reserve-supply curve). The model is subject to the energy balance $\sum_g p_{g,t} + w_t + s_t + LS_t = D_t$ (with dual λ_t , the energy clearing price), the unit-capacity and technical-minimum bounds $\underline{P}_g u_{g,t} \leq p_{g,t} \leq \bar{P}_g u_{g,t}$, the ramp limits

$|p_{g,t} - p_{g,t-1}| \leq RR_g$, the reserve-capability constraints (3.2), and the hard upward-reserve requirement

$$\sum_g r_{g,t} + \eta_t \geq R_t^{\text{req}} \quad \forall t, \quad (3.17)$$

whose dual is the reserve clearing price μ_t . The full LP formulation is collected in Appendix A.

Under this rule, μ_t takes only two qualitative values: it is zero whenever (3.17) is slack, and equals the opportunity cost of the cheapest marginal reserve offer whenever it binds. The transition between these two regimes is discontinuous and fully determined by the administrative parameter R_t^{req} : small changes in net load near the threshold can flip μ_t from zero to a large value (and back). This is the operational manifestation of the knife-edge effect that Joskow and Tirole (2007) identify as the structural weakness of rigid reserve constraints in the reserve-shedding regime of their Proposition 8, the same regime that the welfare-derivative ORDC of Section 3.1 resolves by tracing a continuous height $V \pi_t(r_t) - \hat{M}C_t$ across the same range.

3.4 Piecewise-linear implementation as a soft reserve constraint

For numerical solution, the continuous welfare-derivative curve $\text{ORDC}_t(\cdot)$ of equation (3.11) is approximated by a non-increasing piecewise-linear function with $K = 8$ segments. Define a reference reserve level R_t^{ref} , an intermediate threshold $R_t^{\text{thr}} < R_t^{\text{ref}}$ at which the curve becomes steep, and a cap $\bar{p}$ that limits the maximum scarcity adder. The corresponding shortfall variable $s_t \geq 0$ measures the distance below R_t^{ref} . Replacing the hard requirement (3.17) by the soft constraint

$$\sum_g r_{g,t} + s_t \geq R_t^{\text{ref}} \quad \forall t, \quad (3.18)$$

and adding the scarcity-cost term

$$\sum_t \int_0^{s_t} \text{ORDC}_t(R_t^{\text{ref}} - \xi) d\xi \quad (3.19)$$

to the objective (3.16) converts the welfare-derivative curve into a convex penalty: the more reserves fall short of R_t^{ref} , the larger the per-MW scarcity cost, and the marginal scarcity cost equals the height of the piecewise-linear ORDC at the realised shortfall. Because the integrand is piecewise-linear and non-increasing in r_t , the integral (3.19) is piecewise-quadratic and convex in s_t , so the entire problem remains a convex program and in fact a linear program once the piecewise-quadratic term is disaggregated into segments (Cartuyvels and Papavasiliou, 2022; Papavasiliou and Smeers, 2017). The full slab-disaggregated LP is given in Appendix A.

The dual of (3.18) is again written μ_t , but its behaviour is qualitatively different from the fixed-requirement case. As long as $\sum_g r_{g,t}$ stays in the flat part of the curve, μ_t is small and varies smoothly with the reserve margin; as the margin tightens, μ_t rises along the slope, and reaches the cap $\bar{p}$ only in genuinely scarce states. There is no discontinuous jump between zero and a marginal offer, which is the structural difference with the hard requirement in Section 3.3 (Hogan, 2013).

3.5 Opportunity-cost channel between energy and reserves

A direct consequence of the co-optimised formulation is the arbitrage relation already derived as FOC (3.8). For any flexible unit g that is interior to its capacity and ramp bounds,

$$\lambda_t - C_g = \mu_t \quad (\text{at the dispatch margin}), \quad (3.20)$$

i.e. the unit is indifferent between selling its marginal MW into the energy market (margin $\lambda_t - C_g$) and holding it in reserve (margin μ_t). This is the relation that Papavasiliou and Smeers (2017) derive for the Belgian case study and that Hogan (2013) identifies as the mechanism through which a reserve scarcity adder *simultaneously* lifts the energy price. The practical consequence is that λ_t and μ_t are not independent: under a fixed requirement μ_t is discontinuous, so the energy price inherits a discontinuous scarcity adder; under the ORDC μ_t varies smoothly and the energy price inherits a smooth one. This is the formal sense in which the supervisor's framing of energy and reserves as *complementary products with an opportunity cost between them* enters the model.

3.6 The right marginal price for consumers

An immediate consequence of (3.20) is a sharp statement about the consumer-facing marginal price of one extra MWh. Writing C_t^{en} for the energy-only marginal cost and $\Delta\mu_t \geq 0$ for the induced reserve-price increase, the social marginal cost of consumption is

$$\text{SMC}_t = C_t^{\text{en}} + \Delta\mu_t, \quad (3.21)$$

the operational counterpart of the verbal opportunity-cost argument in Hogan (2013, p. 7): a consumer facing only λ_t does not face the true social marginal cost when reserves are scarce. Under the fixed-requirement baseline $\Delta\mu_t$ is discontinuous at the threshold; under the welfare-derivative ORDC it rises smoothly with the reserve margin and tracks $V \cdot \pi_t(r_t)$ continuously. WP1 reports this through the joint distribution of (λ_t, μ_t) .

3.7 Cost recovery: fixed fee versus energy add-on

Equation (3.21) pins down the *marginal* consumer price; how reserve procurement costs are *recovered* is a separate question. The thesis distinguishes two canonical schemes: *scheme A* (fixed access fee), where the TSO recovers reserve costs through a flat per-period charge $F_t = \mu_t \sum_g r_{g,t}$ collected as a uniform €/MWh uplift F_t/D_t (eq. (3.22)); and *scheme B* (explicit reserve adder), where the consumer is charged the social marginal cost $\lambda_t^{\text{adj}} = \lambda_t + \mu_t$ directly (eq. (3.23)). Scheme A is revenue-neutral by construction but leaves the consumer marginal price at λ_t and thus does not internalise $\Delta\mu_t$. Scheme B transmits the ORDC scarcity signal at the margin but, since served energy is generally larger than reserve volume, its raw revenue $\mu_t D_t$ requires an ex-post rebate or scaling factor for exact budget balance. The comparison in Chapter 4 reports both schemes side by side.

$$F_t = \mu_t \cdot \sum_g r_{g,t}, \quad (3.22)$$

$$\lambda_t^{\text{adj}} = \lambda_t + \mu_t. \quad (3.23)$$

3.8 Scenario matrix and extension

The stylised system is run under three families of scenarios, one for each work package. The *baseline run* (WP1) compares the fixed-requirement and ORDC models on a single representative day at reference parameter values; outputs are dispatch, reserve volumes, the joint distribution of (λ_t, μ_t) , scarcity hours, the reliability proxy $\mathbb{E}[LS_t]$ from (3.4), and total operating cost.

WP2 — Sensitivity to the economic environment. Equation (3.11) identifies three structural ingredients of the welfare-derivative ORDC: net-load uncertainty (which determines $\pi_t(\cdot)$), the share of renewable generation (which drives the spread of ε_t), and the value of lost load V . WP2 varies each on a low–medium–high grid while holding the others at reference. Three predictions follow directly from the derivation: the gap should vanish when uncertainty does (the ORDC reduces to the hard rule), should widen with σ and renewable share (Papavasiliou et al., 2011), and should scale with V in the systemic-collapse tail (Section 3.1.5).

WP3 — ORDC shape extension. WP3 varies the three piecewise-linear parameters of $\text{ORDC}_t(\cdot)$ (the slope between R_t^{thr} and R_t^{ref} , the threshold R_t^{thr} , and the cap $\bar{p}$) on the same stylised system, and traces a cost–reliability frontier in the spirit of Cartuyvels and Papavasiliou (2022). The frontier is expected to be *flat in cost* but *steep in scarcity-price behaviour*, which reinforces the WP1 reporting choice of using tail-price and scarcity-hour metrics rather than averages (Frew et al., 2021).

What is held fixed. Three dimensions are deliberately not varied. The clearing architecture is co-optimised throughout, so the sequential-vs-co-optimised question of Papavasiliou and Avila (2024) is removed. The model is single-stage and deterministic-equivalent, so the back-propagation problem of Papavasiliou (2020) is out of scope. Investment is exogenous, so the flexibility premium of He and Willems (2025) cannot be reimbursed endogenously and is treated as a limit of any single-price reserve auction.

Chapter 4

Results and discussion

This chapter reports the outcomes of the simulations set up in Chapter 3. It is organised in three parts, mirroring the three work packages stated in Chapter 1. Section 4.1 reports the core side-by-side comparison of the fixed reserve requirement and the ORDC on the baseline system (WP1). Section 4.2 reports the sensitivity analyses on net-load uncertainty, renewable penetration and VOLL (WP2). Section 4.3 reports the ORDC shape extension and its cost–reliability frontier (WP3). Section 4.4 interprets the results against the theoretical benchmarks of Chapter 2, and in particular through the Hogan versus Joskow–Tirole lens and the flexibility-premium limit of He and Willems (2025).

4.1 WP1 — Core mechanism comparison

WP1 solves the co-optimised energy-and-reserve LP of Chapter 3 on the Belgium-stylised 24-hour test system: peak load 12 000 MW, a three-block thermal fleet (baseload, mid-merit and peaker), aggregated wind and solar availability, $V = 8\,500$ €/MWh, and a 15-minute response time. The fixed design uses the probabilistic requirement $R_t^{\text{req}} = z_{0.97} \sigma_t$, in the spirit of dynamic dimensioning (De Vos et al., 2019). The ORDC design uses the same level as its reference point $R_t^{\text{ref}} = R_t^{\text{req}}$ and replaces the vertical requirement by the welfare-derivative curve $V\pi_t(r)$ of equation (3.11), discretised into $K = 8$ slabs.

The WP1 outputs are deliberately read in two layers. The baseline instance tests whether the two mechanisms differ on an ordinary day. A stress instance, with all forecast-error standard deviations multiplied by 1.25, tests what happens when the required reserve level presses against the physical reservable capability of the fleet. The detailed tables are exported as `sim/results/wp1/wp1_hourly.csv` and `sim/results/wp1/wp1_stress_hourly.csv`; the headline numbers are exported in `wp1_summary.json` and reproduced in Table 4.1.

4.1.1 Baseline: identical quantities are an economic result

Figure 4.1 shows the dispatch stack under the fixed-requirement model. Baseload covers the night and shoulder hours, mid-merit fills the morning and evening shoulders, and the peaker appears only around the two daily peaks. The ORDC dispatch is visually identical on this instance, so the fixed-dispatch figure is sufficient to document the physical dispatch.




FIGURE 4.1: Hourly dispatch on the baseline instance under the fixed-requirement design. The ORDC dispatch is identical on this calibration because no ORDC shortfall slab is used.

The reserve plot explains why. In Figure 4.2, the total reserve procured by the fixed rule and by the ORDC overlaps hour by hour and coincides with $R_t^{\text{req}} = R_t^{\text{ref}}$. This is not a coding failure. With the chosen $z_{0.97}$ requirement and $V = 8500$ €/MWh, the ORDC’s first potentially activated shortfall slabs are still priced above the observed reserve opportunity costs of the thermal fleet. The optimiser therefore prefers to procure the full reference reserve rather than use a slab variable. In this range, the downward-sloping ORDC and the vertical requirement select the same point on the reserve-supply curve.

The same equivalence appears on the dual side. In Figure 4.3, the energy price λ_t and the reserve price μ_t overlap under the two mechanisms. Reserve scarcity is zero in most off-peak hours, reaches 45 €/MWh in the morning block where baseload headroom is the relevant opportunity-cost margin, and reaches 75 €/MWh in the evening peak where the peaker/mid-merit margin is active. The right panel plots the observed price pairs against the three opportunity-cost lines $\lambda_t = \mu_t + C_g$ for the thermal marginal costs $C_g \in \{20, 65, 140\}$ €/MWh. The points sit on or between these lines depending on which unit is marginal for energy and which unit is marginal for reserve, matching the arbitrage logic of equation (3.20).

The planned scarcity-cost decomposition is therefore exactly zero in the baseline case: no load is shed, no fixed-rule slack is used, and no ORDC shortfall slab is activated (Figure 4.4). The two mechanisms also have the same Monte Carlo reliability score, with $\text{LOLE} = 0.747$ h and $\text{EENS} = 215.6$ MWh over the simulated day. The baseline result should therefore be read as an “abundant reserve” benchmark: the ORDC matters less when its inverse demand is never asked to ration reserve scarcity.

4.1.2 Stress case: same reserve, different scarcity price

The difference between a vertical requirement and an inverse demand curve appears when the target reserve approaches the physical reservable capability of the fleet. The knife-edge sweep in Figure 4.5




FIGURE 4.2: Reserve procurement on the baseline instance. Both mechanisms procure exactly the reference requirement in every hour; the ORDC is willing to pay enough at the margin that no shortfall slab is activated.




FIGURE 4.3: Energy and reserve prices on the baseline instance. The two mechanisms overlap hour by hour. The right panel shows the observed (μ_t, λ_t) pairs against the opportunity-cost lines $\lambda_t = \mu_t + C_g$.

varies the peak-hour target from $0.30R^{\text{reservable}}$ to $1.10R^{\text{reservable}}$, holding the rest of the day fixed. The reservable capability is $R^{\text{reservable}} = \Delta \sum_g RR_g \simeq 1750$ MW. Below that level, both mechanisms clear at ordinary thermal opportunity costs. Once the target exceeds the capability, the fixed rule can only use the slack variable, priced at $V = 8500$ €/MWh. The ORDC instead prices the missing MW with the relevant shortfall slab, so the scarcity price remains bounded by the probability-weighted value $V\pi_t(r)$.

The full stress instance makes the same point hour by hour. With σ_t multiplied by 1.25, the requirement rises above the 1750 MW reservable capability from hours 07 to 17. Figure 4.6 shows that both mechanisms procure the same physical reserve in those hours: they hit the capability limit and leave roughly 280–320 MW of reserve shortfall relative to the reference requirement. The difference is not




FIGURE 4.4: Planned scarcity, slack and ORDC-shortfall cost on the baseline instance. The value is zero in every hour for both mechanisms.




FIGURE 4.5: Knife-edge sweep at the peak hour. The fixed-rule reserve price jumps to VOLL when the target crosses the reservable capability; the ORDC price remains bounded by the welfare-derivative curve.

reliability; it is the scarcity price attached to the same physical scarcity.

Figure 4.7 shows the resulting prices. The fixed-rule reserve price saturates at VOLL for eleven hours, whereas the ORDC reserve price remains near 519 €/MWh, the marginal slab price associated with the observed shortfall. The energy price stays similar across mechanisms because the energy balance itself is feasible; the scarcity is in reserve capability, not in energy production. In the right panel, the fixed-rule stress points therefore sit far to the right of the ordinary opportunity-cost lines: the slack variable prices an administrative requirement violation rather than a marginal fuel-cost opportunity.

This pricing difference is also visible in the scarcity-cost decomposition. In Figure 4.8, the fixed rule records 27.74 M€ of planned scarcity/slack cost over the day, while the ORDC records 1.18 M€ of




FIGURE 4.6: Reserve procurement on the stress instance. Both mechanisms hit the same reservable-capability ceiling during the daytime stress block; the fixed rule records this as slack, while the ORDC records it as use of shortfall slabs.




FIGURE 4.7: Energy and reserve prices on the stress instance ($\sigma_t \times 1.25$). The fixed-rule reserve price saturates at VOLL over the stress block; the ORDC replaces that jump with a bounded probability-weighted scarcity price.

ORDC-shortfall cost for the same physical reserve shortfall. The total planned cost therefore rises from 5.19 M€ of energy cost to 32.93 M€ under the fixed rule, but only to 6.37 M€ under the ORDC.

The stress price distributions in Figure 4.9 summarise the point compactly: energy-price distributions remain almost unchanged, while the reserve-price distribution splits sharply between a VOLL mass for the fixed rule and a much lower ORDC mass around 500 €/MWh. Reliability remains equal in the stress case as well (LOLE = 0.996 h and EENS = 406.9 MWh for both mechanisms) because physical reserve is the same. The mechanism difference is therefore a valuation and payment difference, not a hidden reliability improvement.




FIGURE 4.8: Planned scarcity-cost decomposition on the stress instance. The fixed rule prices the unmet reserve requirement at VOLL; the ORDC prices the same reserve shortfall with probability-weighted shortfall slabs.




FIGURE 4.9: Price distributions on the stress instance. Energy prices remain similar; reserve prices reveal the mechanism difference.

4.1.3 Consumer-facing reserve markups, cost recovery and welfare decomposition

Two further WP1 checks complete the picture and are reported in detail in Appendix C.2. First, the consumer-facing reserve markup $\lambda_t + \mu_t$ coincides under the two designs on the baseline and diverges in the stress case between an 8500 €/MWh fixed-rule cliff and a ≈ 519 €/MWh ORDC add-on (Figures C.1–C.2). Cost-recovery scheme A (a flat uplift on served energy) exactly recovers the reserve-capacity payment, while scheme B (a volumetric reserve add-on) transmits the marginal scarcity signal but requires an explicit rebate to balance the budget; the wedge between the two becomes economically large in the stress case (Figures C.3–C.4). Second, decomposing the peak-hour ORDC into $V_{\text{priv}}\pi_t(r)$ and $V_{\text{pub}}\pi_t(r)$ confirms that the simulated curve is the same welfare-derivative object as in Section 3.1.5: the two components share the same π_t shape and add into the total inverse demand curve for reserves

(Figures C.5–C.6).

4.1.4 WP1 summary

Table 4.1 collects the WP1 headline metrics. The upper block is the baseline result: identical dispatch, identical reserve procurement, identical prices, and identical reliability. The lower block is the stress result: identical physical reserve and reliability, but sharply different reserve scarcity valuation.

TABLE 4.1: WP1 headline metrics. Money figures are in €; prices are in €/MWh; LOLE is in hours and EENS in MWh.

	fixed	ORDC
<i>Baseline instance</i>		
Total planned cost	5 114 234	5 114 234
Scarcity / shortfall cost	0	0
Total reserve procured	33 295	33 295
Mean λ_t	43.3	43.3
P95 μ_t	70.5	70.5
LOLE	0.747	0.747
EENS	215.6	215.6
<i>Stress instance ($\sigma_t \times 1.25$)</i>		
Total planned cost	32 929 543	6 374 040
Scarcity / shortfall cost	27 735 010	1 179 506
Reserve shortfall / slab use	3 263	3 263
Mean λ_t	46.5	46.5
P95 μ_t	8 500.0	519.2
LOLE	0.996	0.996
EENS	406.9	406.9
Expected total cost	36 388 392	9 832 888

The WP1 conclusion is therefore precise. The identical baseline fixed/ORDC result is coherent: at the calibrated requirement, the ORDC is willing to procure the full reserve target, so both mechanisms choose the same physical point. The policy-relevant difference appears in scarcity states. A fixed requirement prices unmet reserve at a VOLL cliff; the ORDC prices it at $V\pi_t(r)$, the marginal welfare value of reserve at the realised reserve margin. This is exactly the bridge needed for the thesis question: the same inverse demand curve captures Hogan’s scarcity-pricing logic while allowing the Joskow–Tirole public-good value of reserve to enter through the calibration of V .

4.2 WP2 — Sensitivity to the economic environment

WP2 holds the system topology fixed and varies the four economic ingredients that enter the welfare-derivative ORDC of equation (3.11): the net-load forecast-error volatility σ_t , the renewable-capacity scaling β that tilts the σ_t profile through wind- and solar-side contributions, the value of lost load V , and the public-good share V_{pub}/V of the social $V = V_{\text{priv}} + V_{\text{pub}}$. For each sweep both the fixed-requirement and the welfare-derivative ORDC dispatches are solved on the Belgium-stylised reference instance; cost, tail prices and Monte-Carlo reliability metrics are reported. The four sweeps map one-to-one onto the

four sub-questions of Section 1.3: σ and β probe the quantity-side response to a tighter LOLP environment, V probes the price-side response, and V_{pub}/V probes whether the distinction between Hogan’s private-good ORDC and the Joskow–Tirole welfare-derivative ORDC is operational.

4.2.1 Net-load uncertainty

The first sweep scales σ_t by a factor $\alpha \in \{0.25, 0.5, 0.75, 1.0, 1.25, 1.5, 2.0\}$, applied jointly to the load, wind and solar components of the quadrature sum that builds σ_t . The fixed requirement is re-sized at the same $z_{0.97}$ quantile so it tracks α , and the ORDC reference R_t^{ref} is set to that same requirement. The two designs therefore see the same headline reserve target; what differs is how they price a deviation from it.




FIGURE 4.10: WP2 σ -sweep: planned operating cost. Fixed-requirement and ORDC coincide for $\alpha \leq 1$ and diverge sharply once the joint reserve+capacity constraint binds at $\alpha \geq 1.25$. The fixed-requirement cost grows quadratically with α because every unmet MW of reserve is priced at the full VOLL; the ORDC grows almost linearly because the marginal welfare cost $V\pi_t(r)$ is bounded below the cap.

Three regimes are visible in Figure 4.10. For $\alpha \leq 1$ the fixed-requirement target is below the fleet’s reservable capability, the optimiser meets it exactly and no scarcity slab activates: the two mechanisms are observationally identical, exactly the WP1 baseline finding. At $\alpha = 1.25$ the required reserve breaches the fleet’s reservable cap; the fixed-requirement design absorbs the shortfall through a slack variable priced at $V = 8500$ €/MWh and the cost jumps sixfold, while the ORDC absorbs it through the welfare-derivative slab that is binding at the realised shortfall, with a marginal price $V\pi_t(r) \approx 519$ €/MWh evaluated at the reserve level r at which the LP equilibrates supply and the ORDC heights (see Section 3.4). Beyond $\alpha = 1.5$ the gap widens further; at $\alpha = 2$ the fixed-requirement design pays the equivalent of a ten-day average operating cost in scarcity penalty alone, while the ORDC’s

penalty is roughly an order of magnitude smaller. Tail prices in Figure 4.11 reproduce the same pattern in the clearing price for reserves: μ_t^{p95} jumps to the VOLL cliff under the fixed rule and rises smoothly under the ORDC.




FIGURE 4.11: WP2 σ -sweep: 95th-percentile reserve clearing price. The fixed-requirement design clears at the VOLL cap as soon as the requirement breaches the reservable capability, illustrating the Joskow–Tirole knife-edge in price (Joskow and Tirole, 2007). The ORDC translates the same physical scarcity into a smooth marginal-welfare price bounded by the curve’s cap.

EENS is identical between the two designs at every α (Figure 4.12). This is a structural property of the stress regime rather than a numerical coincidence: when the required reserve exceeds the total reservable capability of the fleet, both mechanisms procure the maximum the system can deliver, so realised reliability is set by the resource and not by the pricing rule. The cost wedge in Figure 4.10 is therefore purely a wedge in how the unavoidable shortfall is *priced*, not in what is physically procured. This is the empirical analogue of the methodology’s claim that, in capability-constrained operating states, the Hogan and Tirole mechanisms differ only in their price signal, a difference that matters for cost recovery and consumer add-ons but not for the reliability outcome itself.

4.2.2 Renewable share

The renewable-share sweep, scaling installed wind and solar capacity by $\beta \in \{0.5, 0.75, 1.0, 1.25, 1.5\}$ with thermal capacity held fixed, reproduces the same qualitative pattern as the σ sweep: for $\beta \leq 1$ the energy saving from renewable substitution dominates and the two designs are identical; beyond $\beta = 1$ the renewable-driven σ_t pushes the requirement above the reservable cap and the fixed-rule cost jumps to a VOLL-cliff penalty (30.7 M€ at $\beta = 1.25$, 67.2 M€ at $\beta = 1.5$) while the ORDC stays bounded




FIGURE 4.12: WP2 σ -sweep: expected energy not served. The two curves overlap at every α . Realised reliability is set by the fleet's reservable capability, not by the pricing rule; the wedge in Figure 4.10 is therefore a pricing wedge, not a quantity wedge.

by its smooth $V\pi_t(r)$ profile (5.8 and 7.9 M€ respectively). EENS coincides between the two designs at every β and rises from 115 to 722 MWh/day as β grows, so the wedge is again in pricing, not in physical reliability. The detailed figures are in Appendix C.4.

4.2.3 Value of lost load

The VOLL sweep runs $V \in \{3000, 5000, 8500, 12000, 20000\}$ €/MWh at the WP1 stress level $\alpha = 1.25$, so that the reserve constraint binds and V has bite. The qualitative finding mirrors the σ sweep but in price space: fixed-requirement cost grows exactly linearly in V (every unmet MW pays the full V), while the ORDC cost grows sub-linearly because V enters only through the curve heights $V\pi_t(r)$, which are bounded by the survival probability π_t . At $V = 20000$ €/MWh the realised tail price under the ORDC is roughly 1200 €/MWh, a sixteen-fold compression relative to the fixed-requirement clearing price. EENS is invariant to V , because in the capability-bound regime raising V does not unlock additional reserves. The detailed sweep figures are in Appendix C.5.

4.2.4 Public-good share of VOLL

The fourth sweep keeps the social $V = 8500$ €/MWh fixed and varies the public-good share $V_{\text{pub}}/V \in \{0, 0.1, 0.2, 0.4, 0.6\}$, with $V = V_{\text{priv}} + V_{\text{pub}}$ following the decomposition of equation (3.14). Three designs are compared:

- the fixed-requirement design, whose unmet-reserve slack is priced at the full social V ;
- the welfare-derivative ORDC built from the full social V , i.e. the Joskow–Tirole-aware reading in which the public-good externality is internalised in the curve;
- a Hogan-style ORDC built from V_{priv} alone, i.e. the curve a decentralised market would produce when the public-good externality is not internalised.




FIGURE 4.13: WP2 public-good share sweep: 95th-percentile reserve price. Fixed-requirement and Tirole-ORDC are by construction invariant in V_{pub}/V because they price scarcity at the full V (or at $V\pi_t$ derived from it). The Hogan-ORDC flattens as V_{pub} grows because its heights scale with $V_{\text{priv}} = V(1 - V_{\text{pub}}/V)$; at $V_{\text{pub}}/V = 0.6$ the Hogan tail price is below 200 €/MWh, less than half the Tirole price.

Two facts are visible in Figure 4.13. First, both the fixed-requirement and the Tirole-ORDC mechanism produce clearing prices that are flat in V_{pub}/V . This is the operational counterpart of the algebra in Section 3.7: the welfare-derivative ORDC absorbs the entire V into its heights $V\pi_t(r)$, so the private/public split shifts the labelling of V but not the clearing rule. Second, the Hogan-ORDC drops monotonically: a decentralised market that values the reserve only at private willingness-to-pay V_{priv} will undervalue scarcity by exactly $V_{\text{pub}}\pi_t(r)$. The same wedge is visible in operating cost (Figure 4.14), where the Hogan-ORDC line drops below the Tirole-ORDC line by $\approx 12\%$ at $V_{\text{pub}}/V = 0.6$.

The economic interpretation is the answer the thesis gives to sub-question (iv) of Section 1.3. The welfare-derivative ORDC is invariant to the private/public split of V *when the regulator calibrates the curve to the full social V* . The Hogan reading is recovered as the special case in which only V_{priv} is




FIGURE 4.14: WP2 public-good share sweep: planned operating cost. The Hogan-ORDC reports the lowest cost because it is the cheapest-to-clear design, but, as Figure 4.15 shows, that is a property of how the design measures cost, not of the social cost it delivers.

internalised; if a regulator implements that reading and the true social V contains a non-trivial public-good component, the resulting ORDC is too flat and the scarcity price is biased downward by exactly $V_{pub}\pi_t(r)$. In capacity-constrained operating states this bias does *not* translate into more EENS (the fleet has nothing more to give), but it does translate into a mispriced add-on on the consumer-facing reserve component (cf. Section 3.7), so the operational answer is that the Hogan and Joskow–Tirole readings are not interchangeable once cost recovery is taken into account, even when the form of the curve is the same.

4.2.5 WP2 summary

Three robust patterns emerge across the four sweeps. (a) The welfare-derivative ORDC and the fixed-requirement design coincide on quantity in every operating state; their wedge is in *price*, not in physical reliability. EENS is identical within every sweep, because the system reservable capability binds at the stress level that makes scarcity bite. (b) The pricing wedge scales monotonically with each of σ , β and V , and is bounded by the cap of the ORDC, so the operating-cost ratio between the two designs grows roughly linearly in V and quadratically in σ . (c) The public-good share of V does not enter the form of the welfare-derivative ORDC at all, but it shifts the calibration of V that goes into the curve; ignoring V_{pub} produces a Hogan-style ORDC that under-prices scarcity by exactly $V_{pub}\pi_t(r)$. WP2 therefore confirms, on the simulated Belgian instance, the analytical prediction of Chapter 3 that the Hogan and the Joskow–Tirole readings of the ORDC differ in calibration, not in form.




FIGURE 4.15: WP2 public-good share sweep: planned + realised expected cost (planned cost plus social VOLL applied to realised EENS). When the Monte-Carlo loss is priced at the *social* V , the Hogan-ORDC’s apparent saving in Figure 4.14 narrows. The remaining difference reflects the slack-price wedge (the part of scarcity the Hogan design under-prices), not a difference in physical reliability, since EENS is identical across the three designs at every V_{pub}/V .

4.3 WP3 — ORDC shape extension

WP3 leaves the welfare-derivative form $V\pi_t(r) - \widehat{MC}_t$ of WP1/WP2 to one side and traces the cost–reliability frontier across a 45-point grid of piecewise-linear ORDC shapes $(R^{\text{ref}}, R^{\text{thr}}, \bar{p})$, in the spirit of Cartuyvels and Papavasiliou (2022). The grid spans $R^{\text{ref}} \in \{0.85, 1.00, 1.15\} R^{\text{req}}$, $R^{\text{thr}} \in \{0.02, 0.10, 0.25\} R^{\text{ref}}$ and $\bar{p} \in \{0.05, 0.10, 0.20, 0.50, 1.00\} V$, and each variant is solved on the WP2-stress instance (σ scaled by 1.25). The detailed grid, frontier figures and robustness checks are reported in Appendix C.6; this section states the three headline findings.

(a) Pareto-dominance of the ORDC family. The fixed-requirement design lies far outside the empirical cost–reliability scatter: any reasonable piecewise-linear ORDC shape Pareto-dominates the fixed rule at equal reliability, by roughly a factor of three on cost (Figure C.12). This consolidates the WP1/WP2 pricing wedge into a single cross-design comparison.

(b) Calibration levers separate cleanly. Within the ORDC family the reference level R^{ref} is the dominant determinant of EENS, the cap $\bar{p}$ is the dominant determinant of cost, and the threshold R^{thr} has a third-order effect on either (Figures C.13 and C.14). This is the operational counterpart of the calibra-

tion logic of Cartuyvels and Papavasiliou (2022) on a fully-fledged Belgian system-operation simulator: aggregate cost varies modestly across shapes while scarcity-price distributions move much more.

(c) The welfare-derivative ORDC carries an analytical premium. The welfare-derivative ORDC of WP1/WP2 is *not* on the empirical Pareto front: a flatter linear ramp on the same R^{ref} delivers the same EENS at $\approx 7.8\%$ lower clearing cost ($\approx 0.77 \text{ M€}/\text{day}$ on this calibration). The thesis’s preference for the welfare-derivative form is therefore a preference for an *analytical* property (the explicit $V = V_{\text{priv}} + V_{\text{pub}}$ decomposition and the Hogan/Joskow–Tirole reconciliation of Section 3.1.5), not for a numerical cost optimum. The empirical cost premium it carries is the price of that analytical clarity, and is consistent with He and Willems (2025): a shape-only optimisation within the single-price family does not pin down the technology mix anyway, because that mix is set by the option-value structure of a multi-stage market.

4.4 Discussion

This section reads the WP1–WP3 results against the theoretical benchmarks of Chapter 2 along four points; the operational synthesis is consolidated in Chapter 5.

(a) The Hogan LOLP–VOLL adder is recovered exactly. The welfare-derivative ORDC of equation (3.11) carries height $V\pi_t(r) - \hat{M}C_t$, which is precisely the LOLP · VOLL scarcity adder of Hogan (2013) and the closed form of Cartuyvels and Papavasiliou (2022, Eq. 1). The WP1 baseline confirms it operationally: the price pairs (μ_t, λ_t) in Figure 4.3 sit on the $\lambda_t = \mu_t + C_g$ opportunity-cost lines of Papavasiliou and Smeers (2017, eq. 7). The verbal smoothing property of Hogan (2013, p. 7) also shows up quantitatively in WP2: the fixed-requirement μ_t^{p95} jumps from zero to VOLL as σ crosses the reservable cap, while the ORDC’s μ_t^{p95} rises smoothly along $V\pi_t(r)$.

(b) The Joskow–Tirole public-good content lives in the calibration of V . The WP2 public-good sweep shows that the welfare-derivative ORDC built from the social $V = V_{\text{priv}} + V_{\text{pub}}$ is invariant to the V_{pub}/V split, while the Hogan-style ORDC built from V_{priv} alone under-prices reserve scarcity by exactly $V_{\text{pub}}\pi_t(r)$. This is the empirical version of the Joskow and Tirole (2007, p. 81) public-good argument: in capability-constrained states EENS is unchanged because the system is at its physical limit, so the wedge lives entirely in the consumer-facing reserve adder. The Hogan and Joskow–Tirole readings are not interchangeable once cost recovery is taken into account, even when the curve form is the same.

(c) The knife-edge of Joskow and Tirole (2007) shows up structurally. Proposition 8 of Joskow and Tirole (2007, p. 82) predicts that under a fixed reserve rule the reserve clearing price swings “from $(1 + r_i)s$ to $v/(1 + r_i)$ ” for arbitrarily small operator deviations. The WP1 knife-edge sweep (Figure 4.5) reproduces this discontinuity, and the WP2 σ -sweep extends it: the fixed-rule μ_t^{p95} saturates at V once $\alpha \geq 1.25$ while the ORDC’s tail price grows smoothly in α , β and V .

(d) Single-price reserve auctions cannot deliver the flexibility premium. As He and Willems (2025) prove, no single-price reserve-demand rule (ORDC included) reimburses the option value $v_{\theta, \omega}$ that flex-

ible units earn from *waiting* to commit. WP3 confirms one operational manifestation: the welfare-derivative ORDC sits inside the Pareto frontier of the shape grid (Figure C.12), and a flatter linear ramp delivers the same EENS at $\approx 7.8\%$ lower clearing cost on the stress instance. The thesis’s preference for the welfare-derivative form is therefore a preference for an analytical property (the $V = V_{\text{priv}} + V_{\text{pub}}$ decomposition), not a numerical optimum; a shape-only optimisation within the single-price family does not pin down the technology mix because that mix is set by the option-value structure of a multi-stage market the simulator does not represent. Two boundary conditions follow: the model is single-stage and deterministic-equivalent, so the back-propagation problem of Papavasiliou (2020) is out of scope; and investment is exogenous, so the flexibility premium is not priced. The conclusions in Chapter 5 therefore restate the results as statements about the *operational logic* of reserve valuation, not as policy recommendations.

Chapter 5

Conclusion

This thesis compared a fixed reserve requirement with an Operating Reserve Demand Curve inside a common, stylised and Belgium-inspired co-optimisation framework. The methodology (Chapter 3) derived the ORDC used in the model as an inverse demand function for reserves obtained from the derivative of expected welfare with respect to the reserve margin, and showed that its two first-order conditions, $\lambda_t = C_{gm} + \mu_t$ and $\mu_t = V \pi_t(r_t) - \hat{M}C_t$, recover both the Hogan LOLP · VOLL pricing rule and the Joskow–Tirole welfare optimum in a single equation. The results (Chapter 4) report the three work packages on a 24-hour Belgium-stylised instance and on a $\sigma_t \times 1.25$ stress instance.

5.1 Answers to the four sub-questions

The four sub-questions of Section 1.3 are answered, on the model used here, as follows.

(i) *Energy and reserves as complementary products with an opportunity cost between them.* The arbitrage relation $\lambda_t - C_g = \mu_t$ of Papavasiliou and Smeers (2017, eq. 7) holds in both designs, but its content differs: under the fixed requirement μ_t is discontinuous and transmits a discontinuous signal into λ_t ; under the ORDC μ_t varies smoothly and the energy price inherits a smoothed scarcity adder.

(ii) *The right marginal price for consumers.* Equation (3.21) shows that the social marginal cost of one extra MWh is $\lambda_t + \Delta\mu_t$. The fixed-requirement baseline transmits $\Delta\mu_t$ as a discontinuous near-VOLL spike at the threshold; the ORDC transmits it as a smooth markup proportional to $V \cdot \pi_t(r_t)$.

(iii) *Cost recovery.* Scheme A (fixed access fee (3.22)) is budget-balanced by construction but leaves the consumer marginal price at λ_t , so demand-side incentives are not internalised. Scheme B (per-MWh add-on (3.23)) transmits the marginal signal but is not mechanically revenue-equivalent and needs a rebate or scaling rule.

(iv) *The socially optimal reserve level, with two values per MW.* The welfare-derivative ORDC prices both the private-good value $V_{\text{priv}} \pi_t(r_t)$ that Hogan captures and the public-good value $V_{\text{pub}} \pi_t(r_t)$ that Joskow–Tirole stress and that a pure private-VOLL calibration omits. The optimal r_t^* reflects their sum; the WP2 sweep on V_{pub}/V shows that a Hogan-only calibration under-prices reserve scarcity by exactly $V_{\text{pub}} \pi_t(r_t)$.

5.2 Answers to the work packages

WP1 (core mechanism comparison). On the baseline system, the fixed requirement and the welfare-derivative ORDC produce the same dispatch, reserve volume, prices and reliability. This is a coherent calibration result, not a failure: at the chosen requirement the ORDC is willing to procure the full reference reserve, so no shortfall slab is activated. On the $\sigma_t \times 1.25$ stress instance the two mechanisms procure the same physical reserve and reach the same reliability, but value the same shortfall very differently: the fixed rule prices it at VOLL, raising planned cost to 32.93 M€, while the ORDC prices it at the probability-weighted scarcity value, raising it to only 6.37 M€.

WP2 (sensitivity to the economic environment). The four WP2 sweeps confirm three robust patterns. (a) EENS is identical between the two designs at every grid point on the stylised system, because the binding constraint on physical reserve is the fleet’s reservable cap, not the pricing rule. (b) The cost wedge scales with each of σ , β and V and is bounded above by the ORDC cap: the fixed rule’s μ_t^{p95} saturates at V once the cap is breached (the operational counterpart of the Joskow and Tirole (2007) knife-edge), while the ORDC’s tail price grows smoothly along $V\pi_t(r)$. (c) The public-good share enters only the calibration of V , not the form of the curve; a Hogan-style ORDC under-prices scarcity by $V_{\text{pub}}\pi_t(r)$, which confirms the Joskow and Tirole (2007, p. 81) prediction that a private-WTP calibration omits the systemic-collapse externality.

WP3 (ORDC shape extension). The 45-point shape grid on the WP2 stress instance traces the cost–reliability frontier across the $(R^{\text{ref}}, R^{\text{thr}}, \bar{p})$ family. (a) Any reasonable ORDC shape Pareto-dominates the fixed-requirement design at equal reliability by roughly a factor of three on cost. (b) Inside the shape family, R^{ref} sets EENS, the cap $\bar{p}$ sets cost, and R^{thr} has a third-order effect on either, an empirical reading of the calibration logic of Cartuyvels and Papavasiliou (2022). (c) The welfare-derivative ORDC sits inside the empirical Pareto front: a flatter linear ramp on the same R^{ref} achieves the same EENS at $\approx 7.8\%$ lower cost. The preference for the welfare-derivative form is therefore a preference for an *analytical* property (the explicit $V = V_{\text{priv}} + V_{\text{pub}}$ decomposition and the Hogan/Joskow–Tirole reconciliation) rather than a numerical optimum.

Hogan versus Joskow–Tirole, revisited. The two foundational papers do not engage with each other directly. The welfare-derivative ORDC of Chapter 3 treats them as two readings of one inverse demand function: its height $V \cdot \pi_t(r_t) - \hat{M}C_t$ recovers the LOLP · VOLL form of Hogan (2013) and Cartuyvels and Papavasiliou (2022) and Papavasiliou and Smeers (2017), while the public-good content emphasised by Joskow and Tirole (2007, p. 81) is absorbed into the calibration of $V = V_{\text{priv}} + V_{\text{pub}}$ and into the tail of π_t . The thesis’s qualified answer to the overarching reconciliation question of Section 1.2 is therefore that the two views are *compatible at the level of equation form but distinguishable at the level of calibration*, and that the calibration of V and of the tail of $\pi_t(\cdot)$ is where the residual distinction between the two readings actually lives.

5.3 Limitations

The model is deterministic and stylised, uses a single upward reserve product, does not endogenise investment, and does not represent the real-time reserve-capacity market whose absence in Europe is identified in (CREG, 2021; Elia, 2020; Papavasiliou, 2020) as the main obstacle to ORDC back-propagation. It also abstracts from the technology-specific flexibility premium that He and Willems (2025) show no single-price reserve-auction design can reimburse without a complementary call-option menu. The conclusions are therefore insights about the *operational logic* of reserve valuation, not policy recommendations for Belgium or the EU, and in particular not a complete answer on long-run investment adequacy.

5.4 Future work

Natural extensions include: a stochastic two-stage version of the model; a dynamic, temperature- or forecast-error-conditioned $\pi_t(\cdot)$ in the spirit of (Lavin et al., 2020); an explicit real-time reserve- capacity market to study back-propagation (Papavasiliou, 2020); a multi-product ORDC (Hogan and Pope, 2019); and an investment-layer extension embedding the call-option menu of He and Willems (2025) on top of the welfare-derivative ORDC, so that the short-run reserve price and the long-run flexibility premium can be studied jointly.

BIBLIOGRAPHY

- Bajo-Buenestado, R. (Feb. 1, 2021). “Operating reserve demand curve, scarcity pricing and intermittent generation: Lessons from the Texas ERCOT experience”. In: *Energy Policy* 149, p. 112057.
- Cartuyvels, J. and A. Papavasiliou (2022). “Calibration of Operating Reserve Demand Curves Using a System Operation Simulator”. In: *IEEE Transactions on Power Systems*.
- CREG (2020). *Note on Frequently Asked Questions regarding the scarcity pricing mechanism*. Commission for Electricity and Gas Regulation.
- (2021). *Study on the implementation of a scarcity pricing mechanism in Belgium*. F2203EN. Commission for Electricity and Gas Regulation.
- De Vos, K., N. Stevens, O. Devolder, A. Papavasiliou, B. Hebb, and J. Matthys-Donnadieu (Jan. 2019). “Dynamic dimensioning approach for operating reserves: Proof of concept in Belgium”. In: *Energy Policy* 124, pp. 272–285.
- Dowell, J., G. Hawker, K. Bell, and S. Gill (July 1, 2016). “A Review of probabilistic methods for defining reserve requirements”. In: *2016 IEEE Power and Energy Society General Meeting (PESGM)*, pp. 1–5.
- Elia (Dec. 7, 2020). *Final report on Elia’s findings regarding the design of a scarcity pricing mechanism for implementation in Belgium*. Elia System Operator.
- European Commission (Nov. 23, 2017). *Commission Regulation (EU) 2017/2195 of 23 November 2017 establishing a guideline on electricity balancing*. Consolidated text 2021-03-15.
- Frew, B., G. Brinkman, P. Denholm, V. Narwade, G. Stephen, A. Bloom, and J. Lau (Sept. 2021). “Impact of operating reserve rules on electricity prices with high penetrations of renewable energy”. In: *Energy Policy* 156, p. 112443.
- He, D. and B. Willems (Oct. 29, 2025). “Flexibility in Power System: Market Design Matters”. Working paper.
- Hogan, W. W. (2005). *On an “Energy Only” Electricity Market Design for Resource Adequacy*. URL: https://www.researchgate.net/publication/252934713_On_an_Energy_Only_Electricity_Market_Design_for_Resource_Adequacy.
- (2013). *Electricity Scarcity Pricing Through Operating Reserves: An ERCOT Window of Opportunity*. URL: <https://www.hks.harvard.edu/centers/mrcbg/programs/growthpolicy/electricity-scarcity-pricing-through-operating-reserves-ercot>.
- Hogan, W. W. and S. L. Pope (2019). “PJM Reserve Markets: Operating Reserve Demand Curve Enhancements”. In: *Market Design Report*.

- Joskow, P. and J. Tirole (2007). “Reliability and Competitive Electricity Markets”. In: *The RAND Journal of Economics* 38.1, pp. 60–84.
- Lavin, L., S. Murphy, B. Sergi, and J. Apt (Dec. 1, 2020). “Dynamic operating reserve procurement improves scarcity pricing in PJM”. In: *Energy Policy* 147, p. 111857.
- MISO (Nov. 2024). *Updated Shortage Pricing White Paper: Value of Lost Load, Operating Reserve Demand Curve, and Pricing VOLL Circuit Breaker*. Midcontinent Independent System Operator.
- Ovaere, M., E. Heylen, S. Proost, G. Deconinck, and D. Van Hertem (Sept. 1, 2019). “How detailed value of lost load data impact power system reliability decisions”. In: *Energy Policy* 132, pp. 1064–1075.
- Papavasiliou, A. (Dec. 1, 2020). “Scarcity pricing and the missing European market for real-time reserve capacity”. In: *The Electricity Journal* 33.10, p. 106863.
- Papavasiliou, A. and D. Avila (2024). “Welfare Benefits of Co-Optimising Energy and Reserves”. In: *Research and policy report*.
- Papavasiliou, A., S. Oren, and R. O’Neill (Nov. 1, 2011). “Reserve Requirements for Wind Power Integration: A Scenario-Based Stochastic Programming Framework”. In: *IEEE Transactions on Power Systems* 26.4, pp. 2197–2206.
- Papavasiliou, A. and Y. Smeers (2017). “Remuneration of flexibility using operating reserve demand curves: a case study of Belgium”. In: *The Energy Journal* 38, pp. 105–135.
- Potomac Economics (2025). *2024 State of the Market Report for the ERCOT Electricity Markets*. Independent Market Monitor for ERCOT.
- Zhou, Z. and A. Botterud (Jan. 2014). “Dynamic Scheduling of Operating Reserves in Co-Optimized Electricity Markets With Wind Power”. In: *IEEE Transactions on Power Systems* 29.1, pp. 160–171.

Appendices

Appendix A

Model formulation in detail

This appendix collects, in one place, the full mathematical formulation of the co-optimised energy-and-reserve dispatch model used to compare the fixed reserve requirement of Section 3.3 with the welfare-derivative ORDC of Sections 3.1–3.4. The notation follows Chapter 3; cross-references are given to the equations introduced there.

A.1 Sets and indices

$\mathcal{G}$	set of thermal unit groups, indexed by g (baseload, mid-merit, peaker)
t	$\in$ hourly periods, $T = 24$
$\{1, \dots, T\}$	
k	$\in$ ORDC slabs, $K = 8$ in the implementation
$\{1, \dots, K\}$	

A.2 Parameters

C_g	marginal energy cost of unit g (€/MWh)
$\bar{P}_g$	installed capacity of unit g (MW)
$\underline{P}_g$	technical minimum of unit g (MW)
RR_g	ramp rate of unit g (MW/h)
Δ	reserve response time (h), 0.25 h (15-min product) in the baseline
D_t	inelastic load in period t (MW)
W_t, S_t	wind and solar availability in period t (MW), price-takers in the LP
V	social value of lost load (€/MWh), $V = V_{\text{priv}} + V_{\text{pub}}$
σ_t	standard deviation of net-load forecast error ε_t in period t
R_t^{req}	fixed reserve requirement in period t (MW), used in the baseline
R_t^{ref}	reference reserve level used by the ORDC in period t
R_t^{thr}	ORDC threshold below which the cap price applies
$b_{t,k}^{\text{ORDC}}$	breakpoint k of the ORDC slab discretisation in period t
$\bar{p}$	ORDC cap price (€/MWh)
$p_{t,k}^{\text{slab}}$	marginal price of slab k in period t (€/MWh)

A.3 Decision variables

$p_{g,t} \geq 0$	energy dispatch of unit g in period t (MW)
$r_{g,t} \geq 0$	upward reserve held by unit g in period t (MW)
$w_t, s_t \geq 0$	realised wind and solar dispatch (MW), bounded by W_t, S_t
$LS_t \geq 0$	load shed in period t (MW)
$\eta_t \geq 0$	reserve slack in the fixed-requirement formulation only
$s_{t,k}^{\text{slab}} \geq 0$	shortfall on ORDC slab k in period t (ORDC mode only)

A.4 Common constraints (both modes)

Energy balance.

$$\sum_{g \in \mathcal{G}} p_{g,t} + w_t + s_t + LS_t = D_t, \quad \forall t, \quad (\text{A.1})$$

with dual λ_t (the energy clearing price).

Capacity and technical-minimum bounds.

$$p_{g,t} + r_{g,t} \leq \bar{P}_g, \quad \forall g, t, \quad (\text{A.2})$$

$$p_{g,t} \geq \underline{P}_g u_{g,t}, \quad \forall g, t, \quad (\text{A.3})$$

with $u_{g,t} \equiv 1$ in the LP relaxation (the unit-commitment extension is left to future work).

Ramp limits.

$$|p_{g,t} - p_{g,t-1}| \leq RR_g, \quad \forall g, t \geq 2, \quad (\text{A.4})$$

with $p_{g,0}$ given as the initial dispatch from Section 3.2.

Reserve capability (eq. (3.2)).

$$r_{g,t} \leq \Delta RR_g, \quad \forall g, t. \quad (\text{A.5})$$

Renewable upper bounds.

$$w_t \leq W_t, \quad s_t \leq S_t, \quad \forall t. \quad (\text{A.6})$$

A.5 Mode A — fixed reserve requirement (Section 3.3)

Reserve demand.

$$\sum_{g \in \mathcal{G}} r_{g,t} + \eta_t \geq R_t^{\text{req}}, \quad \forall t, \quad (\text{A.7})$$

with dual μ_t (the reserve clearing price). The slack η_t is priced at the cap of the implicit vertical curve, V , in the objective.

Objective.

$$\min \sum_t \left[\sum_g C_g p_{g,t} + V L S_t + V \eta_t \right]. \quad (\text{A.8})$$

A.6 Mode B — welfare-derivative ORDC (Sections 3.1–3.4)

The continuous welfare-derivative curve $\text{ORDC}_t(r) = V \pi_t(r) - \hat{\text{MC}}_t$ of equation (3.11) is approximated by $K = 8$ slabs $\{(b_{t,k}^{\text{ORDC}}, p_{t,k}^{\text{slab}})\}_{k=1}^K$ between R_t^{thr} and R_t^{ref} , with slab heights $p_{t,k}^{\text{slab}} = V \pi_t(\tilde{r}_{t,k})$ evaluated at the slab midpoint $\tilde{r}_{t,k} = (b_{t,k-1}^{\text{ORDC}} + b_{t,k}^{\text{ORDC}})/2$ and sorted non-decreasingly as r falls.

Reserve demand (eq. (3.18)).

$$\sum_{g \in \mathcal{G}} r_{g,t} + \sum_{k=1}^K s_{t,k}^{\text{slab}} \geq R_t^{\text{ref}}, \quad \forall t, \quad (\text{A.9})$$

with dual μ_t .

Slab capacities.

$$0 \leq s_{t,k}^{\text{slab}} \leq b_{t,k-1}^{\text{ORDC}} - b_{t,k}^{\text{ORDC}}, \quad \forall t, k. \quad (\text{A.10})$$

Objective (eq. (3.19)).

$$\min \sum_t \left[\sum_g C_g p_{g,t} + V L S_t + \sum_{k=1}^K p_{t,k}^{\text{slab}} s_{t,k}^{\text{slab}} \right]. \quad (\text{A.11})$$

The integrand of the scarcity-cost term is piecewise-linear and non-increasing in r_t , so the resulting integral $\int_0^{s_t} \text{ORDC}_t(R_t^{\text{ref}} - \xi) d\xi$ is convex and, after disaggregation into the K slab variables, the entire programme is a linear programme.

A.7 Equilibrium / interpretive identities

At the optimum of either programme, the standard KKT stationarity conditions of the co-optimised LP imply, for any flexible unit g at neither its capacity nor its technical minimum and whose ramp is not binding,

$$\lambda_t - C_g = \mu_t, \quad (\text{A.12})$$

the opportunity-cost arbitrage relation (3.20) of Section 3.5; this is the version of the identity stated in Papavasiliou and Smeers (2017, eq. 7).

The consumer-facing social marginal cost of one extra MWh of consumption is, by equation (3.21),

$$\text{SMC}_t = C_t^{\text{en}} + \Delta\mu_t, \quad (\text{A.13})$$

with $\Delta\mu_t = \mu_t$ in the LP, since the reserve component of the consumer add-on is exactly the reserve dual.

The welfare-derivative ORDC factorises by the additive split $V = V_{\text{priv}} + V_{\text{pub}}$ as equation (3.15),

$$\text{ORDC}_t(r_t) = V_{\text{priv}} \pi_t(r_t) + V_{\text{pub}} \pi_t(r_t) - \hat{\text{MC}}_t. \quad (\text{A.14})$$

Both components share the same shortfall probability $\pi_t(r_t)$. In the implementation $V_{\text{priv}} = 0.8V$, $V_{\text{pub}} = 0.2V$ in the baseline; the WP2 public-good sweep varies the split between $V_{\text{pub}}/V \in \{0, 0.1, 0.2, 0.4, 0.6\}$.

A.8 Numerical implementation

The programme is built and solved in Python using the PuLP modelling layer with the open-source CBC solver shipped via the PULP_CBC_CMD backend. All MW figures are in MW, all costs in €/MWh, all times in hours. The code that materialises the formulation is in `sim/lsm_sim/dispatch.py`; the ORDC discretisation is in `sim/lsm_sim/ordc.py`; the system parameters and profiles are in `sim/lsm_sim/system.py`; and the realised-reliability Monte Carlo is in `sim/lsm_sim/evaluation.py`. The work-package drivers `run_wp1.py`, `run_wp2.py` and `run_wp3.py` reproduce the figures and tables of Chapter 4 end-to-end.

Appendix B

Input data and scenario tables

This appendix lists the numerical values that parameterise the stylised Belgium-inspired test system used throughout Chapters 3–4. All values are chosen to reproduce the broad shape of the Belgian system as documented in Papavasiliou and Smeers (2017) and Cartuyvels and Papavasiliou (2022), while keeping the model small enough to be reproduced end-to-end in a master’s thesis codebase.

B.1 Thermal fleet

The thermal fleet is reduced to three unit groups. The capacities roughly aggregate $\simeq 5.5$ GW of low-marginal-cost baseload (the $\simeq 5$ GW of Belgian nuclear at the time of the case study, plus the small coal-equivalent block), $\simeq 4.5$ GW of CCGT mid-merit, and $\simeq 2.5$ GW of OCGT/diesel peaker capacity. Ramp rates are calibrated so that only the fast units can meaningfully provide upward reserve over a 15-minute response window, in the spirit of Papavasiliou and Smeers (2017, p. 109–112). Reserve capability per unit is the product of the response time $\Delta = 0.25$ h and the ramp rate RR_g .

TABLE B.1: Stylised thermal fleet parameters used in WP1–WP3.

Unit	$\bar{P}_g$ (MW)	$\underline{P}_g$ (MW)	C_g (€/MWh)	RR_g (MW/h)	ΔRR_g (MW)
baseload	6 000	3 000	20	600	150
midmerit	4 500	600	65	2 400	600
peaker	2 500	0	140	4 000	1 000
Total	13 000	—	—	—	1 750

The total reservable capability of the fleet is therefore $\sum_g \Delta RR_g = 1\,750$ MW, the binding ceiling that appears as the dashed reference line in the WP1 knife-edge sweep (Figure 4.5) and as the cap that triggers the price wedge in the WP2 σ -sweep (Figure 4.10).

B.2 Demand and renewable profiles

Demand is a stylised winter-weekday Belgian load profile normalised to a peak of $D^{\text{peak}} = 12\,000$ MW, with a two-peak shape (morning around 09:00, evening around 19:00) and a night floor at $0.55 D^{\text{peak}}$.

Wind availability uses 4000 MW of installed capacity, a 30 % capacity factor and a small sinusoidal envelope; solar uses 6000 MW of installed capacity and a deterministic sun-elevation bell with peak capacity factor 0.65. The exact analytic forms are in `sim/lsm_sim/system.py`; the resulting hourly values are reproduced in Table B.2.

TABLE B.2: Hourly load and renewable availability used in WP1 (rounded to the nearest MW).

t	D_t	W_t	S_t	t	D_t	W_t	S_t
0	6607	1527	0	12	9114	756	3900
1	6630	1695	0	13	8089	809	3702
2	6701	1645	0	14	7532	711	3143
3	6888	1339	0	15	7618	921	2319
4	7295	1771	0	16	8454	833	1379
5	8029	1641	0	17	9905	819	514
6	9102	1380	0	18	11372	797	0
7	10333	1518	514	19	12000	995	0
8	11346	1374	1379	20	11364	1154	0
9	11741	1252	2319	21	9874	1206	0
10	11346	1091	3143	22	8352	1624	0
11	10335	1096	3702	23	7330	1659	0

B.3 Forecast-error distribution and reserve sizing

The additive net-load forecast error $\varepsilon_t \sim \mathcal{N}(0, \sigma_t^2)$ has time-varying standard deviation built as the quadrature sum of three contributions, in line with the standard reserve-sizing literature (De Vos et al., 2019; Dowell et al., 2016; Papavasiliou et al., 2011):

$$\sigma_t = \sqrt{(\sigma^L D_t)^2 + (\sigma^W \bar{W} \mathbf{1}[W_t > 0])^2 + (\sigma^S \bar{S} \mathbf{1}[S_t > 1])^2}, \quad (\text{B.1})$$

with the reference fractions $\sigma^L = 2\%$ (load), $\sigma^W = 15\%$ (wind capacity) and $\sigma^S = 10\%$ (solar capacity), and installed capacities $\bar{W} = 4000$ MW, $\bar{S} = 6000$ MW. The fixed reserve requirement used as the WP1 baseline and as the ORDC reference reserve is

$$R_t^{\text{req}} = z_{0.97} \sigma_t, \quad z_{0.97} \simeq 1.881, \quad (\text{B.2})$$

i.e. the probabilistic, dynamic-dimensioning rule of De Vos et al. (2019).

B.4 ORDC discretisation

The continuous welfare-derivative curve $\text{ORDC}_t(r) = V\pi_t(r) - \hat{M}C_t$ is approximated by $K = 8$ slabs of equal width on $[R_t^{\text{thr}}, R_t^{\text{ref}}]$ with $R_t^{\text{thr}} = 0.05 R_t^{\text{ref}}$ in WP1 and WP2. Slab heights are evaluated at slab midpoints and sorted non-decreasingly as r falls. The cap price is set to $V\pi_t(R_t^{\text{thr}})$, capped at V itself by construction. The ORDC heights $V\pi_t(r)$ at the peak hour of the baseline run are shown in Figure C.6; the private/public decomposition at the same hour is shown in Figure C.5.

B.5 Reference values and scenario grids

TABLE B.3: Reference values of the economic parameters and the WP2 sensitivity grids.

Parameter	Reference	WP2 grid
V (€/MWh)	8 500	$\{3\,000, 5\,000, 8\,500, 12\,000, 20\,000\}$ (run at $\alpha = 1.25$)
V_{pub}/V	0.20	$\{0.0, 0.1, 0.2, 0.4, 0.6\}$ (at $\alpha = 1.25$)
σ -multiplier α	1.00	$\{0.25, 0.5, 0.75, 1.0, 1.25, 1.5, 2.0\}$
RES scaling β	1.00	$\{0.5, 0.75, 1.0, 1.25, 1.5\}$
Response time Δ (h)	0.25	—
Quantile for R^{req}	0.97	—

TABLE B.4: WP3 ORDC shape grid (45 variants, run on the $\alpha = 1.25$ stress instance).

Parameter	Grid
$R^{\text{ref}}/R^{\text{req}}$	$\{0.85, 1.00, 1.15\}$
$R^{\text{thr}}/R^{\text{ref}}$	$\{0.02, 0.10, 0.25\}$
$\bar{p}/V$	$\{0.05, 0.10, 0.20, 0.50, 1.00\}$

B.6 Realised-reliability evaluation

The deterministic-equivalent dispatch is scored by Monte Carlo on $N = 5\,000$ draws (or $N = 2\,500$ in the WP2/WP3 sweeps for runtime reasons) of $\varepsilon_t \sim \mathcal{N}(0, \sigma_t^2)$. For each draw, the realised reserve shortfall is $\max(\varepsilon_t - \sum_g r_{g,t}, 0)$, and the realised load shed is the same shortfall priced at V . The reliability metrics $\text{LOLE} = \sum_t \Pr(\varepsilon_t > \sum_g r_{g,t})$ and $\text{EENS} = \sum_t \mathbb{E}[\max(\varepsilon_t - \sum_g r_{g,t}, 0)]$ are sample averages over the N draws. The per-draw daily load-shed distribution is summarised by the mean and the 95th percentile; both are reported in `wp1_summary.json` for the baseline and stress instances.

Appendix C

Supplementary results and derivations

This appendix collects supplementary material that supports Chapters 3–4 but is not required to follow the headline findings. It is organised in three sections. Section C.1 contains the detailed side-by-side comparison of the welfare-derivative first-order conditions of Chapter 3 with the optimality conditions of Hogan (2013) and Joskow and Tirole (2007). Section C.2 reports the detailed WP1 consumer-side and welfare-decomposition analyses (cost-recovery schemes and the private/public split of V). Section C.6 reports the full WP3 ORDC-shape extension, including the cost–reliability frontier across the 45-point shape grid.

C.1 Detailed comparison of the first-order conditions with Hogan (2013) and Joskow and Tirole (2007)

This section expands the brief comparison of Section 3.1 into the full mapping between equation (3.11) and the optimality conditions of Hogan (2013) and Joskow and Tirole (2007).

Comparison with Hogan (2013)

Hogan (2013, p. 6) states the LOLP · VOLL result as economic intuition, not as the conclusion of a formal optimisation:

“Whenever there is involuntary load curtailment and the system has just the minimum of contingency operating reserves, then any increment of reserves would correspondingly reduce the load curtailment. Hence the price of operating reserves should be set at the value of lost load. At any other level of operating reserves... the value of an increment of operating reserves would be the same VOLL multiplied by the probability that net load would increase enough in the coming interval to reduce reserves to the minimum level where load would be curtailed... Hence the incremental value of operating reserves would be the analogous to the product of the loss of load probability (LOLP) and VOLL, or LOLP · VOLL.”

In the thesis’s notation, Hogan’s verbal rule is exactly (3.11). At the load-shedding boundary ($\pi_t \rightarrow 1$), $\text{ORDC}_t \rightarrow V - \hat{M}C_t$, i.e. VOLL net of avoided activation cost; at any other reserve level, the marginal

value is $V \cdot \pi_t(r_t) - \hat{MC}_t$, which is exactly Hogan’s “LOLP-VOLL”. The corresponding energy-price form (3.13) coincides with the closed-form $\lambda = MC_{gm} + (VOLL - \widehat{MC}) \cdot LOLP(R_\Delta)$ of Cartuyvels and Papavasiliou (2022) and Papavasiliou and Smeers (2017, eq. 10), modulo the commitment-vs-activation interpretation of $\hat{MC}$ discussed in Section 3.1.

The substantive difference is therefore the *provenance*, not the *content*. Hogan (2013) states the rule as economic intuition and *does not take the derivative of $L_t(r_t)$* : he treats r_t as exogenous and asks what price reserves should clear at. The thesis takes the opposite route: r_t is treated as endogenous, L_t is differentiated, and the resulting FOC (3.10) both pins down r_t^* and *equates the marginal value of reserve to its marginal cost*, an optimality condition Hogan does not state.

Comparison with Joskow and Tirole (2007)

Joskow and Tirole (2007, Section 4(a) and Proposition 8) *do* take the welfare derivative explicitly, but they do not draw the resulting object as a downward-sloping demand curve. Their reserve coefficient r_i is reserve held as a proportion of nominal supply, and their FOC for the optimal high-reserve regime is

$$h_i v / (1 + r_H)^2 = s \quad (C.1)$$

(Joskow and Tirole, 2007, eq. 10, p. 79): marginal benefit of reserve equals marginal cost. Translating $h_i \leftrightarrow \pi_t(\cdot)$, $v \leftrightarrow V$, and $s \leftrightarrow \hat{MC}_t + \mu_t$, (C.1) becomes

$$V \pi_t(r_H^*) = \hat{MC}_t + \mu_t, \quad (C.2)$$

which is *exactly* the thesis’s FOC (3.10). The functional form of the marginal benefit differs (J-T use a specific parametric form, the thesis uses the more general $V \pi_t(\cdot)$), but the economic content is identical.

Joskow and Tirole (2007, Proposition 8, p. 82) then characterise three regimes as forecast demand grows. They translate directly into three sub-regions of (3.10)–(3.11):

Off-peak ($r_t = r_H$). Capacity is slack, $\pi_t(r_H)$ is small, $\mu_t \approx 0$, and the ORDC sits in its low-height range. Consumers pay $(1 + r_H)s$ for one MWh of served load (Joskow and Tirole, 2007, p. 81), the analogue of $\lambda_t = C_{gm}$.

Reserve-shedding ($r_t \in (r_L, r_H)$). Capacity binds, generation is rationed across reserve and load, π_t is intermediate, and J-T report a quasi-rent $h_i v / (1 + r_i)^2 - s$, exactly the welfare-derivative reserve price $V \pi_t(r_t) - \hat{MC}_t$. Crucially, J-T point out that here “the supply curve and the total demand curve are vertical and identical”, so a “small mistake... creates wild swings in the market price (from $(1 + r_i)s$ to $v / (1 + r_i)$ conditionally on being served)” (Joskow and Tirole, 2007, p. 81). The ORDC (3.11) replaces this vertical band with a smooth downward-sloping curve and *eliminates the knife-edge by construction*.

Load-shedding ($r_t = r_L$). $\pi_t(r_L) \rightarrow 1$, load is curtailed, and the reserve price approaches $V - \hat{MC}_t$. The thesis’s ORDC reaches its cap; this is Hogan’s “price-of-reserves-equals-VOLL” boundary.

The substantive difference is therefore not in the optimality condition, which is the same FOC, but in what the analyst does with it. J-T use (C.2) to pin down the regime boundaries r_L and r_H , and leave the

intermediate regime as a knife-edge in price; the thesis takes the same FOC, draws the marginal welfare function as a downward-sloping curve, and prices the intermediate regime continuously along that curve.

Hogan vs. Joskow–Tirole, restated. The ORDC (3.11) is therefore the welfare-derivative one obtains by doing the derivative Hogan did not do, on the welfare function Joskow and Tirole did write down, and smoothing the regime they leave as a knife-edge.

C.2 Detailed WP1 consumer-side and welfare-decomposition analyses

This section reports two WP1 sub-analyses that the main text (Section 4.1) summarises in passing: the consumer-facing reserve markup and cost-recovery schemes (Section C.2.1), and the welfare-decomposition check that the simulated ORDC is the same object as the welfare-derivative ORDC of equation (3.11) (Section C.2.2).

C.2.1 Consumer-facing reserve markups and cost recovery

The consumer-facing reserve markup in the WP1 implementation is the unscaled per-MWh add-on μ_t of Section 3.6, shown as $SMC_t = \lambda_t + \mu_t$ in Figure C.1. On the baseline instance, the fixed rule and ORDC again coincide: the markup is zero in most hours, 45 €/MWh in the morning reserve-scarcity block, and 75 €/MWh in the evening peak.




FIGURE C.1: Consumer-facing reserve markup on the baseline instance. The unscaled add-on μ_t is identical under the two mechanisms because the underlying reserve prices are identical.

On the stress instance (Figure C.2), the fixed-rule add-on reaches VOLL during the stress block. With this calibration, that is 8500 €/MWh, while the ORDC add-on remains around 519 €/MWh. This is the consumer side of the same mechanism difference: a fixed requirement turns an unmet reserve target into a VOLL-priced cliff, while the ORDC turns it into a probability-weighted scarcity adder.

The cost-recovery figures separate the accounting question from the marginal-incentive question. In the baseline case (Figure C.3), scheme A exactly recovers the reserve-capacity payment $\mu_t \sum_g r_{g,t}$




FIGURE C.2: Consumer-facing reserve markup on the stress instance. The fixed-rule add-on inherits the VOLL cliff; the ORDC add-on remains tied to $V\pi_t(r)$.

through an average uplift F_t/D_t ; this totals 478 184 € over the day. Scheme B passes the marginal add-on μ_t through on served energy. It gives the stronger marginal signal, but the raw volumetric revenue is not equal to the reserve-capacity payment unless it is paired with a rebate, lump-sum transfer or scaling factor.

In the stress case (Figure C.4), this distinction becomes economically large. Fixed-rule reserve payments at $\mu_t = V$ would imply very high capacity payments and very high volumetric revenues if passed through raw. The ORDC compresses both numbers because its reserve price is tied to $V\pi_t(r)$ rather than to the full VOLL. This is the concrete cost-recovery implication of the mechanism comparison: a fixed requirement can be made budget-balanced through a flat uplift, but it gives a poor marginal signal in scarcity; an ORDC-compatible add-on gives a better marginal signal, but its budget balance has to be designed explicitly.

C.2.2 The welfare interpretation of the ORDC

WP1 checks that the ORDC used in the simulations is the same object that links the Hogan and Joskow–Tirole readings in the methodology. Figure C.5 decomposes the peak-hour ORDC into $V_{\text{priv}}\pi_t(r)$ and $V_{\text{pub}}\pi_t(r)$ using $V_{\text{pub}}/V = 0.20$. The components share the same loss probability $\pi_t(r)$ and add vertically into the total curve. Operationally, only their sum affects dispatch; conceptually, the split records how much of the reserve value is read as private interruption value and how much is read as public-good systemic reliability value.

Figure C.6 shows the same peak-hour curve without the decomposition. The step shape is the LP approximation to the smooth welfare derivative; scarcity prices rise as reserve r falls, and the cap near $R_t^{\text{thr}} = 0.05R_t^{\text{ref}}$ is about 3 940 €/MWh in this hour. The peak- σ hour produces the same qualitative curve and is left in the results folder as an audit figure.




FIGURE C.3: Cost recovery on the baseline instance. Scheme A is an average uplift that exactly recovers reserve-capacity payments; scheme B is an unscaled volumetric add-on that transmits the marginal reserve scarcity signal.

C.3 The flexibility-premium critique: full exposition

This section reproduces the full discussion of the He and Willems (2025) flexibility-premium critique that is summarised as a scope limit in Section 2.7 of Chapter 2.

He and Willems (2025) develops a three-stage peak-load-pricing model in which uncertainty is resolved progressively: at $T = 1$ a generator chooses capacity K_θ at investment cost I_θ for each technology θ ; at $T = 2$, once an interim signal ω is observed, the generator commits energy and reserve quantities $(Q_\omega, R_\omega^U, R_\omega^D)$ at commitment costs (I_θ^U, I_θ^D) ; at $T = 3$, once the final demand state ξ is realised, reserves are activated at operating costs (c_θ^U, c_θ^D) that may differ from the energy marginal cost c_θ . This three-stage structure is the specific extension of Joskow and Tirole (2007): it preserves the public-good motivation for reserves but adds the *commitment* margin and the *asymmetric adjustment-cost* structure that the original two-state Joskow–Tirole framework did not contain.

The technology-specific flexibility premium. The central object identified by He and Willems (2025) is the technology-specific flexibility premium $v_{\theta,\omega} \equiv \lambda_{\theta,\omega} - \lambda_{\theta,\omega}^{\text{nr}}$, i.e. the additional value a flexible




FIGURE C.4: Cost recovery on the stress instance. The fixed-rule VOLL spike dominates both the reserve-capacity payment and the raw volumetric add-on; the ORDC keeps both at the probability-weighted scarcity level.

technology generates over an otherwise identical non-responsive technology with the same marginal cost. In a state-contingent real-time pricing setting, this premium can be expressed as the maximum of put or call options on the spot price (their Lemma 2). The economic intuition is the same as the supervisor’s framing of energy and reserves as *complementary products*: a flexible unit’s reserve revenue is not just the opportunity cost of selling its marginal MW into the energy market, it also includes the option value of *waiting* to commit until the state is observed.

Why standard reserve auctions miss this premium. He and Willems (2025) show that, in the absence of an efficient real-time market, every market-based reserve-auction design they consider fails to reimburse $v_{\theta,\omega}$. Forward-only markets pay $\max\{p_{\omega}^F - c_{\theta}, 0\}$, which depends only on the expected spot price and the marginal cost; the option premium collapses to zero. Uniform-price reserve auctions (Chao–Wilson style; Papavasiliou and Smeers, 2017), integrated energy–reserve co-optimisation (Papavasiliou and Avila, 2024, Oren–Sioshansi style) and non-linear scoring rules all benchmark reserve payments on opportunity cost rather than option value, and therefore *systematically under-invest in flexible technology*. He and Willems (2025) prove that the only market-based design that achieves the first best is one




FIGURE C.5: Peak-hour ORDC decomposed into private-good and public-good value components. Both components use the same shortage-probability shape and add into the total inverse demand curve for reserves.

in which reserves are settled through a *menu of call-option contracts* indexed by strike prices that match the adjustment costs of the available technologies (their Proposition 4).

Implications for this thesis. Three implications follow directly. First, the comparison of a fixed reserve requirement with an ORDC carried out in this thesis is a comparison *within* the class of single-price reserve-demand rules that He and Willems (2025) shows to be structurally incomplete for long-run investment in flexible technology. The thesis therefore operates at the operational, short-run layer, and any conclusion about investment incentives is flagged as outside scope. Second, the flexibility-premium argument *reinforces* the Joskow–Tirole reading of Section 2.1: a pure LOLP · VOLL curve omits part of the social value of reserves; here that missing part is the technology-specific option value on activation, and in the stylised model of Chapter 3 it can only be partially absorbed into the calibration of $\pi_t(\cdot)$ and V . Third, the same paper *legitimises* the choice of co-optimisation as the baseline clearing architecture in Section 3.3: integrated clearing removes a first-order source of mis-pricing (the opportunity-cost error of sequential clearing) even when, on its own, it cannot deliver the investment optimum. The reconciliation between Hogan and Joskow–Tirole that the welfare-derivative ORDC of Chapter 3 achieves *at the level of operational pricing* is therefore not contradicted but rather *bounded above* by He and Willems (2025): a single-price curve, however calibrated, cannot deliver the flexibility premium that a state-contingent option menu would.




FIGURE C.6: Peak-hour welfare-derivative ORDC, discretised into eight slabs between R_t^{ref} and $R_t^{\text{thr}} = 0.05R_t^{\text{ref}}$.

C.4 WP2 renewable-share sweep — detailed results

This section reproduces the detailed figures for the WP2 renewable-share sweep summarised in Section 4.2.2.

The renewable-share sweep scales installed wind and solar capacity by $\beta \in \{0.5, 0.75, 1.0, 1.25, 1.5\}$. The thermal fleet, the load and the load-side error fraction are kept fixed so that the sweep isolates the contribution of wind- and solar-side forecast errors to σ_t and, indirectly, to the reserve requirement.

The two effects in Figure C.7 are visible separately. Energy cost decreases monotonically with β (6.2 M€ at $\beta = 0.5$ down to 4.4 M€ at $\beta = 1.5$). The reserve cost is zero up to $\beta = 1$, then jumps once the renewable-driven σ_t pushes the requirement above the reservable cap. Beyond that point the same wedge as in the σ sweep opens up: the fixed-requirement design pays a VOLL-cliff penalty (30.7 M€ at $\beta = 1.25$, 67.2 M€ at $\beta = 1.5$) while the ORDC limits the scarcity penalty to its smooth $V\pi_t(r)$ profile (5.8 M€ and 7.9 M€ respectively). Realised EENS, shown in Figure C.8, again coincides between the two designs and rises with β from 115 to 722 MWh/day, confirming that the wedge in cost is a wedge in pricing rather than in physical reliability.

The economic reading is that, on this calibration, a renewable-share increase that is not accompanied by additional flexible capability translates into reserve scarcity, not into energy scarcity. The fixed-requirement design transmits this as a discontinuous reserve price, the ORDC as a continuous one. The quantity outcome (the EENS curve) is the same.

C.5 WP2 VOLL sweep — detailed results

This section reproduces the detailed figures for the WP2 VOLL sweep summarised in Section 4.2.3.




FIGURE C.7: WP2 renewable-share sweep: planned operating cost. Energy cost falls with β as renewable energy displaces thermal output, while the reserve requirement rises with the wind- and solar-side error contributions. For $\beta \leq 1$ the energy saving dominates; beyond $\beta = 1$ the joint reserve+capability constraint binds and the fixed-requirement design pays the same VOLL-cliff scarcity penalty as in the σ sweep.

The VOLL sweep is run at the WP1 stress level $\alpha = 1.25$ (so the reserve constraint binds and V has bite) and varies $V \in \{3000, 5000, 8500, 12000, 20000\}$ €/MWh. Below the stress level, the constraint is slack and V does not enter the operating cost, a prediction the WP1 baseline confirms.

The two prices in Figure C.9 are visually parallel to the two designs in Figure 4.10 (and to the corresponding renewable-share result in Figure C.7), and they carry the same economic content. The fixed-requirement scarcity adder is V itself; the ORDC scarcity adder is $V\pi_t(r)$ where $\pi_t(r)$ is the survival probability of the reserve shortfall. At $V = 20000$ €/MWh, π_t at the worst-case slab is below 6%, so the realised tail price under the ORDC is roughly 1 200 €/MWh, a sixteen-fold compression of the price signal relative to the fixed-requirement clearing price of 20 000 €/MWh (Figure C.10).

EENS is invariant to V (Figure C.11). The reason is the same capability-binding mechanism as in the σ sweep: with the fleet at its reservable cap, raising V does not unlock additional reserves. The interpretation matches the theory of Section 3.1: in capability-constrained states the ORDC trades cost-recovery efficiency for the welfare value of reserve, while the fixed requirement collects the full VOLL on every unmet MW even though the unmet MW cannot be physically served.




FIGURE C.8: WP2 renewable-share sweep: expected energy not served. EENS scales with β because the additional wind- and solar-side error widens the residual-load distribution.

C.6 WP3 — ORDC shape extension

WP3 leaves the welfare-derivative form $V\pi_t(r) - \widehat{MC}_t$ of WP1/WP2 to one side and tests how the operating outcome depends on the three *shape* parameters of a more general piecewise-linear ORDC (the reference level R^{ref} , the threshold R^{thr} and the cap $\bar{p}$), following the calibration logic of Cartuyvels and Papavasiliou (2022). The shape ORDC is constructed by `build_ordc_shape`: a linear ramp from 0 at R^{ref} up to $\bar{p}$ at R^{thr} , discretised into eight slabs. The grid spans $R^{\text{ref}} \in \{0.85, 1.00, 1.15\} R^{\text{req}}$, $R^{\text{thr}} \in \{0.02, 0.10, 0.25\} R^{\text{ref}}$ and $\bar{p} \in \{0.05, 0.10, 0.20, 0.50, 1.00\} V$, giving 45 shape variants. Each is solved on the same WP2-stress instance (σ scaled by 1.25) so that the reserve constraint binds and the shape parameters separate.

Two natural reference points anchor the comparison: the fixed- requirement design at R^{req} (the pre-ORDC benchmark), and the welfare-derivative ORDC of WP1/WP2 (the Joskow–Tirole-aware reference shape). Both are run on the same stress instance and overlaid on the cost–reliability scatter so the reader can see whether any shape variant Pareto-dominates the welfare-derivative anchor.

C.6.1 Cost–reliability frontier

Three patterns are visible in Figure C.12. (a) The shape grid resolves into three horizontal bands of EENS, one per R^{ref} level. The band structure is sharp because the system is reservable-capability-bound at the stress level (the same mechanism diagnosed in WP2): once R^{ref} is set, the amount of reserve the fleet can deliver is fixed, so the realised EENS is essentially determined by R^{ref} and only weakly




FIGURE C.9: WP2 VOLL sweep at $\alpha = 1.25$: planned operating cost. Fixed-requirement cost is exactly linear in V (slope = EENS · slack hours) because every unmet MW pays the full V . The ORDC cost rises sub-linearly because V only enters through the curve heights $V\pi_t(r)$, which are bounded by the survival probability π_t .

modulated by the shape inside the curve. Figure C.13 colours the same scatter by R^{ref} to make this explicit.

(b) The Pareto-efficient subset is the south-west envelope: $R^{\text{ref}} = 1.00$ at the lowest-cap-price corner of its band ($\approx 8.9\text{M€}$, EENS ≈ 402) connects to the $R^{\text{ref}} = 1.15$ band, which is then traced from low cost / EENS ≈ 355 to slightly higher cost / EENS ≈ 338 as $\bar{p}$ rises. Inside the $R^{\text{ref}} = 1.15$ band the cap-price increase from $0.05V$ to $0.50V$ buys a 5% reduction in EENS at the price of a 60% increase in cost; raising the cap further to $1.00V$ buys no additional reliability and adds another 60% in cost (it is dominated by the $\bar{p} = 0.50V$ variants on the same band). $R^{\text{ref}} = 0.85$ is entirely Pareto-dominated: lowering R^{ref} below R^{req} saves a small amount of energy cost (the system carries less reserve) but pays for it many times over in EENS. The threshold R^{thr} has near-zero effect on EENS within each band and only a marginal effect on cost (visible as the small spread of points at each cap colour in Figure C.13); the curve's marginal prices in the $[R^{\text{thr}}, R^{\text{ref}}]$ region are what the LP actually trades against, and only the height of the cap region (which is rarely active here) responds to R^{thr} .

(c) The two reference designs sit very differently on the plot. The fixed-requirement design clears at the same EENS as $R^{\text{ref}} = 1.00$ (EENS ≈ 402) but reports a total cost of 36.3M€ , more than three times the entire shape grid. This is the WP1/WP2 finding once again: the fixed rule prices unmet reserve at the full V , the shape ORDC at $\bar{p} \leq V$. The welfare-derivative ORDC sits at $(9.83\text{M€}, 401.6)$, inside the $R^{\text{ref}} = 1.00$ band and visibly off the Pareto front: at the same EENS, the cheapest matching shape




FIGURE C.10: WP2 VOLL sweep: 95th-percentile reserve price. The fixed-requirement clearing price is identically equal to V once the constraint binds; the ORDC clearing price rises in proportion to $V\pi_t(R^{\text{thr}})$, much more slowly.

variant ($R^{\text{ref}} = 1.00$, $R^{\text{thr}} = 0.02R^{\text{ref}}$, $\bar{p} = 0.20V$) clears at 9.06M€, a saving of $\approx 7.8\%$ relative to the welfare-derivative point. A flatter $\bar{p} = 0.05V$ ramp clears at 8.87M€ but only at the higher EENS ≈ 424 . The reason for the wedge is that the welfare-derivative curve $V\pi_t(r)$ has higher per-MW heights than a linear ramp at the same R^{ref} and $\bar{p}$: the survival function π_t rises sharply as r falls below R^{ref} , so the welfare-derivative ORDC prices intermediate slabs more aggressively. In a capability-bound regime where the marginal slabs are activated *anyway*, that aggressive pricing translates into higher scarcity revenue without buying additional physical reserve.

C.6.2 Reading the result

The WP3 frontier provides three operational answers to the calibration question of Section 3.1.5. First, the shape grid spans the full cost-reliability space the simulator can resolve, and the fixed-requirement design lies far outside it: any reasonable ORDC shape Pareto-dominates the fixed rule on cost at equal reliability. Second, within the shape grid the dominant lever is R^{ref} , not $\bar{p}$ or R^{thr} . This is the empirical reading the thesis reaches on the stylised system: the reference reserve level acts as a primary calibration lever for the realised reliability, while the cap $\bar{p}$ sets the intensity of the scarcity-price signal but, on a system that is already capability-bound, has only a minor reliability effect. The qualitative pattern (aggregate cost varies modestly across shapes while reserve and energy price distributions move visibly) is consistent with the calibration evidence reported in Cartuyvels and Papavasiliou (2022, Tables III–IV) on a fully-fledged Belgian system-operation simulator. Third, the welfare-derivative ORDC of WP1/WP2 is




FIGURE C.11: WP2 VOLL sweep: expected energy not served. Invariant in V , because in the binding-capability regime V enters the operating problem only through the price of unmet reserve, not through its quantity.

not on the empirical Pareto front of the shape grid: a flatter linear ramp with the same R^{ref} delivers identical reliability at lower clearing cost. The welfare-derivative form is preferred for the reasons developed in Chapter 3: it carries the explicit welfare-decomposition $V = V_{\text{priv}} + V_{\text{pub}}$ and makes the cost-recovery scheme A/scheme B distinction transparent. The empirical cost-recovery wedge it carries is nevertheless real and quantifiable on this calibration at roughly 0.77 M€/day at the stress level ($\approx 7.8\%$ of total operating cost) when EENS is held constant. A regulator that values the analytical clarity of the welfare-derivative form must accept this premium; one that prefers the cost-minimising piecewise-linear shape must accept that the resulting curve is no longer the marginal-welfare ORDC of equation (3.11) and that its private/public decomposition is no longer model-implied.

C.6.3 Robustness with respect to the shape parameters

Figure C.14 confirms the reading visually. The right panel is colour-segmented by EENS band, the left panel is colour-segmented along each band by $\bar{p}$, and the middle panel shows essentially overlapping points across the three R^{thr} values. The robustness is convenient for calibration: a regulator who fixes R^{ref} at the operationally accepted reserve target need only choose $\bar{p}$ to sit on the desired point of the cost-reliability frontier; R^{thr} can be set on technical grounds (e.g. as the smallest reserve level at which the activation procedure remains reliable) without disturbing the economic outcome.




FIGURE C.12: WP3 cost–reliability frontier. Coloured dots are the 45 shape variants of the WP3 grid; colour encodes the cap price $\bar{p}$. Black-circled points are the Pareto-non-dominated subset (the WP3 frontier proper). The blue diamond is the welfare-derivative ORDC and the red cross the fixed-requirement design, both run on the same stress instance. The fixed-requirement design lies far off-frontier on the cost axis; the welfare-derivative ORDC sits inside the frontier at the EENS ≈ 402 MWh/day band.

C.6.4 WP3 summary

WP3 confirms three points. (a) Any reasonable piecewise-linear ORDC Pareto-dominates the fixed-requirement design at equal reliability, by a factor of 3 on cost on the WP2 stress instance, a consistent extension of the WP1/WP2 reading. (b) Inside the family of piecewise-linear ORDCs, the reference level R^{ref} is the dominant determinant of EENS, the cap $\bar{p}$ is the dominant determinant of cost, and the threshold R^{thr} has a third-order effect on either. This is the operational counterpart of the calibration logic of (Cartuyvels and Papavasiliou, 2022). (c) The welfare-derivative ORDC of WP1/WP2 is not the cost-minimising shape; a flatter linear ramp on the same R^{ref} achieves the same EENS at $\approx 7.8\%$ lower clearing cost. The methodology’s preference for the welfare-derivative form is therefore a preference for an *analytical* property (the explicit private/public decomposition), not for a numerical optimum, and the empirical cost premium it carries is the price of that analytical clarity. WP3 thus pins down precisely what the welfare-derivative form costs in operational terms and is the empirical answer to the methodological question of which ORDC shape a regulator should pick.




FIGURE C.13: WP3 shape grid grouped by reference level R^{ref} . Each R^{ref} level produces a near-flat EENS band with cost varying along the band as a function of the cap and threshold. The Pareto frontier of Figure C.12 therefore consists almost entirely of $R^{ref} = 1.15 R^{req}$ variants.




FIGURE C.14: WP3 cost–reliability scatter coloured separately by the cap price $\bar{p}$, the threshold R^{thr} and the reference R^{ref} . Confirms that R^{ref} is the dominant lever (right panel separates the EENS bands cleanly), the cap $\bar{p}$ is the dominant cost lever within each band (left panel) and R^{thr} has near-zero effect on either dimension (middle panel).

Résumé :

Ce mémoire compare deux mécanismes d'approvisionnement des réserves opérationnelles dans les marchés de l'électricité : une exigence fixe de réserve et une *Operating Reserve Demand Curve* (ORDC). Un modèle stylisé de co-optimisation énergie-réserve, inspiré du contexte belge, est développé ; l'ORDC est construite comme une fonction de demande inverse pour les réserves, dérivée de la valeur attendue de la fiabilité, et relie la tarification des réserves à la probabilité de perte de charge et à la valeur de l'énergie non servie.

Les résultats montrent que, dans un système disposant de réserves suffisantes, les deux mécanismes donnent des résultats presque identiques. En revanche, dans un contexte plus tendu, l'exigence fixe conduit à des pics de prix très élevés, tandis que l'ORDC produit un signal de rareté plus progressif et pondéré par le risque. Les deux mécanismes atteignent un niveau de fiabilité similaire mais avec des signaux de prix et des coûts très différents : le choix du mécanisme influence surtout la formation des prix et la répartition des coûts, plus que la fiabilité physique du système.

Abstract:

This thesis compares two operating-reserve procurement mechanisms in electricity markets: a fixed reserve requirement and an Operating Reserve Demand Curve (ORDC). Using a stylised, Belgium-inspired energy-and-reserve co-optimisation model, the ORDC is built as an inverse demand function derived from the expected value of reliability. In a well-supplied system the two mechanisms behave almost identically; under stress the fixed rule produces very high price spikes, while the ORDC delivers a smoother, risk-weighted scarcity signal. Both can reach a similar reliability level, but with very different prices and costs.

UNIVERSITÉ CATHOLIQUE DE LOUVAIN

Louvain School of Management

Place des Doyens, 1 bte L2.01.01, 1348 Louvain-la-Neuve

Boulevard Emile Devreux 6, 6000 Charleroi, Belgique

Chaussée de Binche 151, 7000 Mons, Belgique

www.uclouvain.be/lsm