




Adipic Acid Production

Bottlenecks and Viability Conditions of the
Conventional, Electrochemical, and Biomass-Based Routes
under European Market and Policy Scenarios

F.A. Hakkaart

MSc Thesis · Mechanical Engineering · Track: Energy, Flow and Process Technology · 2026



Adipic Acid Production

Bottlenecks and Viability Conditions of the Conventional, Electrochemical, and Biomass-Based Routes under European Market and Policy Scenarios

by

F.A. Hakkaart

to obtain the degree of Master of Science
in Mechanical Engineering at the Delft University of Technology,
to be defended publicly on Wednesday 23 September 2026.

Student number:	5279526
Project duration:	February 2026 – September 2026
Thesis committee:	Prof. Dr. W. de Jong, TU Delft, supervisor
	Dr. R. Kortlever, TU Delft, supervisor
	Dr. S.U. Khan, TU Delft, daily supervisor
	Dr. ir. M. Ramdin, TU Delft, committee member

An electronic version of this thesis is available at <https://repository.tudelft.nl/>.

Preface

This thesis is the graduation project of my MSc in Mechanical Engineering, track Energy, Flow and Process Technology. I am grateful to everyone who helped, directly or indirectly, with the work behind it.

I want to thank my supervisors: Wiebren de Jong and Ruud Kortlever, for the questions and input that gave the project its direction, and Shahid Khan as my daily supervisor, for the sparring sessions that helped shape my research. What I appreciated as much as the technical guidance were the conversations that had nothing to do with adipic acid, on geopolitics among other things, which made the meetings something to look forward to. I would also like to thank Mahinder Ramdin for taking the time to be part of the committee. The work also owes a great deal to Edward Yeo, whose process models it builds directly on; his flowsheets gave me a running start and, just as usefully, a set of results to argue with.

I felt welcome in the faculty from the first week, and that is down to the people in it, so my thanks go to the whole electrochemistry group. Thanks too to the friends who were writing their own theses alongside mine, for the second opinions and for the shared frustration when nothing converged. Doing this at the same time as other people made it a far better year than doing it alone would have been.

Finally, to my family and my friends outside the faculty, who, willingly or not, have heard about every aspect of adipic acid for the past eight months, thank you for the support and the patience.

Looking back, it has been a fun, challenging and thought-provoking period, and one I would do again. I hope you enjoy reading this work.

*F.A. Hakkaart
Delft, September 2026*

Abstract

Adipic acid is a large-volume chemical, used mainly for nylon, and about 95 % of world production runs through nitric-acid oxidation of benzene-derived intermediates, a route whose nitrous oxide by-product dominates the product's emissions and whose benzene ties it to a fossil supply chain. Electrochemical oxidation and a biomass route via electrochemical hydrogenation are proposed against both problems, and all three are modelled together in Yeo (2025), the one published baseline that spans them in Aspen Plus at a consistent scale and boundary. This thesis re-evaluates that baseline under the European market and policy conditions it holds fixed, not to rank a fixed winner, but to establish the bottlenecks and viability conditions that decide each route.

The re-evaluation is built on one basis for a 100 kt/a European plant: a full pinch analysis on each flowsheet, cell degradation over service life, carbon priced under the EU ETS, and electricity charged at each Antwerp–Rotterdam–Rhine–Ruhr (ARRRA) country's tariff and grid carbon intensity in place of a single price and an assumed renewable supply. Around that basis every input the published work holds fixed is swept rather than assumed: the conventional route's N_2O destruction efficiency, the carbon price, each country's grid intensity and electricity price, cell performance and degradation, feedstock prices and procurement, load factor, and both biomass endpoints, *trans*-3-hexenedioic acid (*t*3HDA) and adipic acid. Each is varied alone and in combination.

On central inputs the conventional route reaches a levelised cost of adipic acid (LCOA) of 1.87 \$/kg at a cradle-to-gate greenhouse-gas (GHG) intensity of 3.39 kgCO₂e/kg; the electrochemical route 2.19 \$/kg at 3.79 kgCO₂e/kg; and the biomass route 4.03 \$/kg at 2.34 kgCO₂e/kg to *t*3HDA, or 4.34 at 2.56 once catalytically hydrogenated to adipic acid. Against that baseline the electrochemical route's cost premium over the conventional route narrows from 47 to 17 %, and its GHG position reverses, from 54 % below the conventional route to 12 % above. No route is best on both measures.

The binding constraints differ by route. The conventional route turns mainly on the destruction efficiency it sustains and the carbon price on its residual N_2O , the electrochemical route on the carbon intensity and price of its electricity, and the biomass route on the cost of a muconic-acid broth no market yet quotes and on whether a buyer pays for non-fossil origin or for the reactive double bond *t*3HDA retains. Both alternatives also rest on electrochemical cells demonstrated only in the laboratory. Across every case assessed here, abating the conventional process's N_2O costs orders of magnitude less per tonne of CO₂-equivalent than any change of route, which points to unabated capacity, most of it outside Europe, as the larger near-term lever on the adipic-acid sector's GHG emissions.

Contents

List of Figures	vii
List of Tables	viii
Nomenclature	x
1 Introduction	1
1.1 Background	1
1.2 Motivation and Problem Statement	1
1.3 Objective and Scope	2
1.4 Research Questions	3
1.5 Relation to Prior Work	3
2 Literature Review	4
2.1 Market Context and Policy Motivation	4
2.2 Adipic Acid Production Pathways	5
2.2.1 Conventional Adipic Acid Production	5
2.2.2 Electrochemical Adipic Acid Production	7
2.2.3 Biomass-Based Adipic Acid Production	9
2.3 Cross-Route Techno-Economic Comparison	11
2.3.1 Prior Techno-Economic Findings	12
2.3.2 CAPEX Estimation Approaches	13
2.3.3 Scale Assumptions and Reference Capacity	13
2.4 Extensions to Existing Techno-Economic Assessments	13
2.4.1 Carbon Pricing and EU ETS	13
2.4.2 European Industrial Context and Constraints	15
2.4.3 Country-Coupled Electricity Inputs.	15
2.4.4 Biomass Route Product Boundary: <i>t</i> 3HDA versus Adipic Acid	17
2.4.5 Electrochemical Cell Performance and Degradation	17
2.4.6 Feedstock Treatment and Input Values	18
2.4.7 Scenario Framework Summary	19
3 Methodology	20
3.1 Scope and Assessment Framework.	20
3.2 Process Modelling Approach	21
3.2.1 Aspen Plus configuration.	21
3.2.2 Reactor and flowsheet modelling	21
3.2.3 Biomass route extension: <i>t</i> 3HDA to adipic acid.	21
3.2.4 Heat integration.	22
3.2.5 Electrochemical cell model.	24
3.2.6 Cell degradation and electrode replacement	24
3.2.7 Model scope and validation	25
3.3 Economic assessment	25
3.3.1 Computational framework and the OpenPyTEA cost engine.	25
3.3.2 Levelised cost of adipic acid	26
3.3.3 Capital cost (CAPEX)	29
3.3.4 Differences from Yeo (2025)	30

3.4	Environmental assessment	31
3.5	Scenario evaluation	32
3.5.1	Treatment of inputs over plant lifetime.	32
3.5.2	HNO ₃ procurement scenario construction.	33
3.5.3	Capital-cost learning for the electrochemical cells	33
3.5.4	Coherent-future narratives	34
3.5.5	Load-following under dynamic electricity prices.	35
3.5.6	Evaluation modes	36
4	Results	37
4.1	Headline route comparison	37
4.1.1	Levelised cost of adipic acid	37
4.1.2	GHG intensity.	38
4.1.3	Material efficiency	39
4.2	Corrections and extensions to the reference model	40
4.2.1	Cost scope and capital corrections	40
4.2.2	Heat integration.	41
4.2.3	Cathodic hydrogen and the by-product credit	43
4.2.4	Electricity price and grid intensity	43
4.2.5	Cell degradation	43
4.2.6	EU ETS charge and abatement	44
4.2.7	Product boundary: t3HDA versus adipic acid.	44
4.3	Location: price and grid intensity	45
4.4	Procurement scenarios.	46
4.5	Cost and carbon sensitivities	48
4.5.1	EU ETS and abatement	48
4.5.2	The conventional route under complete N ₂ O abatement.	50
4.5.3	The electrochemical cells	51
4.5.4	Capital and discount rate.	52
4.5.5	Tornado analysis	53
4.6	Load-following under dynamic electricity prices.	55
4.7	Coherent-future narratives	57
5	Conclusions	60
5.1	Answers to the research questions	60
5.2	Key bottlenecks and viability conditions	62
5.3	Implications for the wider sector	62
6	Recommendations	63
	Bibliography	65
A	Model inputs and factors	76
A.1	Electricity prices in ARRRA	76
A.2	Grid carbon intensity	78
A.3	Rebased grid carbon intensity trajectories	79
A.4	Carbon price trajectory	79
A.5	Feedstock treatment	80
A.5.1	Benzene	80
A.5.2	Hydrogen	81
A.5.3	HNO ₃	82
A.5.4	MA broth	83
A.5.5	Process chemicals	83
A.6	Utility emission factors	84
A.7	Electrochemical cell performance: the evidence base	84

B	Headline results build-up	87
B.1	Cell-degradation saw-tooth	87
B.2	Cell performance against price and grid intensity	87
B.3	Cathodic hydrogen valuation	88
B.4	Capital accuracy band by equipment group	89
C	Comparison to Yeo (2025)	90
C.1	Input values	90
C.2	Capital-cost decomposition	91
C.3	Sources of the cost difference	91
C.3.1	Cash-flow timing	91
C.3.2	Fixed capital	92
C.3.3	Fixed operating cost	93
D	Heat-exchange networks	95
D.1	Stream properties and stream-list curation	95
D.1.1	Individual film coefficients	95
D.1.2	Crystallisation duty	96
D.1.3	Curation of the extracted stream list	96
D.2	Heat-exchange network grid diagram	96
D.3	Pinch stream lists	97
D.4	Exchanger-by-exchanger detail	99
D.4.1	CON: conventional, HNO_3 oxidation	100
D.4.2	ECO: electrochemical oxidation	100
D.4.3	BIO _{t3HDA} : biomass to the intermediate	101
D.4.4	BIO _{AA} : biomass to adipic acid	101
E	Load-following under dynamic electricity prices	102
E.1	Electricity network-cost decomposition	102
E.2	Threshold optimisation	103
E.3	Effective electricity price against load factor	104
E.4	Scope-2 GHG intensity against load factor	104
E.5	Load-following upper bound: cell oversizing and the intermediate buffer	104
E.5.1	Oversizing to hold output	105
E.5.2	The buffer, and why this is an upper bound	105
E.6	ECO cell start–stop cost	105
F	Aspen Plus flowsheets	107
F.1	Conventional route	107
F.2	Electrochemical route	111
F.3	Biomass route to t3HDA	113
F.4	Biomass route to adipic acid	115
G	Aspen Plus stream tables	117
G.1	Conventional route	118
G.2	Electrochemical route	119
G.3	Biomass route to t3HDA	121
G.4	Biomass route to adipic acid	122
	Acknowledgement of generative AI use	124

List of Figures

2.1	Conventional adipic acid process, simplified PFD	6
2.2	Electrochemical oxidation (ECO) route, simplified PFD	7
2.3	Biomass ECH route, simplified PFD	9
3.1	Biomass-route extension from <i>t</i> 3HDA to adipic acid	23
4.1	LCOA decomposed by cost component, this work vs Yeo (2025)	38
4.2	GHG intensity decomposed by scope, this work vs Yeo (2025)	39
4.3	LCOA against GHG intensity, by route and country	46
4.4	LCOA across each route's feedstock-price band	48
4.5	LCOA against the EU ETS carbon price, by N ₂ O destruction efficiency	50
4.6	Conventional route: one-at-a-time sensitivity	53
4.7	Electrochemical route: one-at-a-time sensitivity	54
4.8	Biomass route: one-at-a-time sensitivity	55
4.9	Effective electricity price against load factor, decomposed	56
4.10	Change in LCOA against load factor	57
4.11	LCOA under the three coherent futures	58
4.12	GHG intensity under the three coherent futures	59
A.1	All-in electricity cost breakdown, ARRR countries	77
A.2	Country grid carbon intensity trajectories	79
A.3	Carbon-price trajectories	80
B.1	ECO specific energy consumption, by stack-life tier	88
C.1	Fixed capital composition, this work against Yeo (2025)	93
D.1	Heat-exchange network, conventional route	97
D.2	Heat-exchange network, electrochemical route	97
D.3	Heat-exchange network, biomass route to <i>t</i> 3HDA	98
D.4	Heat-exchange network, biomass route to adipic acid	98
E.1	Change in scope-2 GHG intensity against load factor	104
F.1	Conventional route, hydrogenation and air oxidation	108
F.2	Conventional route, nitric-acid oxidation and N ₂ O abatement	109
F.3	Conventional route, crystallisation and recovery	110
F.4	Aspen Plus flowsheet of the electrochemical route.	112
F.5	Aspen Plus flowsheet of the biomass route stopping at <i>t</i> 3HDA.	114
F.6	Aspen Plus flowsheet of the biomass route to adipic acid.	116

List of Tables

2.1	Emissions of EU adipic acid plants behind the ETS benchmark	6
2.2	Cross-route baseline performance, cost drivers and uncertainties	12
2.3	Carbon-price scenario parameters	14
2.4	Electrochemical cell-performance tiers	18
2.5	Feedstock treatment and input values	19
2.6	Competitiveness dimensions and indicative ranges	19
3.1	Common basis of design.	20
3.2	Cell-degradation parameters by tier	25
3.3	Construction and start-up schedule	27
3.4	Installed-cost factors by equipment class	29
3.5	Capital and duty changes relative to Yeo (2025)	31
3.6	HNO ₃ procurement scenarios and scope-3 intensities	33
3.7	Coherent-future narratives, by scenario dimension	34
4.1	Central-case results against Yeo (2025)	38
4.2	Cost bridge: Yeo's MSP to this work's LCOA	40
4.3	GHG bridge: Yeo (2025) to this work	40
4.4	Utility duties before and after heat integration	41
4.5	Effect of the designed network on LCOA	42
4.6	Effect of the designed network on GHG intensity	42
4.7	Electricity as a cost and a carbon effect, by route	43
4.8	Electricity price, grid intensity and LCOA by country	45
4.9	GHG intensity on a declining grid	46
4.10	Feedstock supplier swaps, by route	47
4.11	Each route on non-fossil feedstock	48
4.12	The conventional route against N ₂ O destruction efficiency	49
4.13	The electrified routes by cell-performance tier	52
4.14	What cell degradation adds, by durability tier	52
A.1	All-in electricity costs for EHV baseload users, 2024	77
A.2	Grid carbon intensity, ARRR countries and France, 2024	78
A.3	Per-country grid CI trajectory values	79
A.4	Benzene procurement scenarios	80
A.5	Hydrogen procurement scenarios	82
A.6	Component values feeding the HNO ₃ supplier scenarios	83
A.7	Process-chemicals price and scope-3 CI anchors	84
A.8	Utility greenhouse-gas emission factors, per unit of duty (kgCO ₂ e/kWh).	84
B.1	LCOA and GHG intensity by component, per route	87
B.2	The electrified routes over cell tiers, price and grid intensity	88
B.3	Treatments of the surplus cathodic hydrogen	88
B.4	Installed capital by AACE accuracy group	89
B.5	What the capital accuracy band is worth on LCOA	89
C.1	Every input to the headline results, this work against Yeo (2025).	91
C.2	Capital cost by equipment category, this work against Yeo (2025)	92
C.3	Fixed capital, Yeo (2025) to this work, by source	93
C.4	Fixed operating cost, Yeo (2025) to this work, by source	94

D.1	Individual film coefficients by fluid category	96
D.2	Conventional route: pinch stream list	98
D.3	Electrochemical route: pinch stream list	99
D.4	Biomass route to <i>t</i> 3HDA: pinch stream list	99
D.5	Biomass route to adipic acid: pinch stream list	99
D.6	Designed heat-exchange networks, per route	100
D.7	Conventional (HNO ₃ oxidation): final HEN exchangers, sorted by area.	100
D.8	Electrochemical oxidation: final HEN exchangers, sorted by area.	101
D.9	Biomass to <i>t</i> 3HDA: final HEN exchangers, sorted by area.	101
D.10	Biomass to adipic acid: final HEN exchangers, sorted by area.	101
E.1	Per-country electricity network-cost components	102
E.2	Effective electricity price against load factor	104
E.3	Load-following optima under the two operating bounds	106
E.4	ECO start–stop at each country’s cost-optimal load factor	106
G.1	Conventional route: Aspen Plus stream summary	119
G.2	Conventional route: battery-limit stream compositions	119
G.3	Electrochemical route: Aspen Plus stream summary	120
G.4	Electrochemical route: battery-limit stream compositions	121
G.5	Biomass route to <i>t</i> 3HDA: Aspen Plus stream summary	121
G.6	Biomass route to <i>t</i> 3HDA: battery-limit stream compositions	122
G.7	Biomass route to adipic acid: Aspen Plus stream summary	123
G.8	Biomass route to adipic acid: battery-limit stream compositions	123

Nomenclature

Abbreviations

Abbrev.	Definition	Abbrev.	Definition
AA	Adipic acid	GHG	Greenhouse gas
AACE	Association for the Advancement of Cost Engineering	GI	Grid-intensity scenario
AC	Activated carbon support	GWP	Global warming potential
AEM	Anion exchange membrane	HEN	Heat-exchange network
AR6	IPCC Sixth Assessment Report	HER	Hydrogen evolution reaction
ARRRA	Antwerp–Rotterdam–Rhine–Ruhr Area	ICC	Indirect Cost Compensation
BAT	Best Available Technique	ICP-OES	Inductively coupled plasma optical emission spectrometry
BEN	Benzene scenario (Fossil / Bio)	IEA	International Energy Agency
BIO	Biomass route (family)	IPCC	Intergovernmental Panel on Climate Change
BIO _{t3HDA}	Biomass route, <i>t3HDA</i> endpoint	IPPC	Integrated Pollution Prevention and Control
BIO _{AA}	Biomass route, AA endpoint	ISBL	Inside battery limits
BREF	Best Available Techniques Reference Document	ISCC	International Sustainability and Carbon Certification
BTX	Benzene, toluene, xylene	KA	Ketone–alcohol oil (CHN/CHA)
CAPEX	Capital expenditure	KPI	Key performance indicator
CBAM	Carbon Border Adjustment Mechanism	LCA	Life cycle assessment
<i>ccMA</i>	<i>cis,cis</i> -muconic acid	LCOA	Levelised cost of adipic acid
CCS	Carbon capture and storage	LDH	Layered double hydroxide
CDM	Clean Development Mechanism	LF	Load factor
CEPCI	Chemical Engineering Plant Cost Index	LMTD	Log mean temperature difference
CHA	Cyclohexanol	LNG	Liquefied natural gas
CHN	Cyclohexanone	LP	Low-pressure (steam)
CI	Carbon intensity	LR	Literature review (Chapter 2)
CIP	Copenhagen Infrastructure Partners	MA	Muconic acid
CN	Combined Nomenclature (EU customs)	MP	Medium-pressure (steam)
CON	Conventional route (HNO ₃ oxidation)	MSP	Minimum selling price (Yeo's term for LCOA)
COP	Coefficient of performance	MSR	Market Stability Reserve
CP	Carbon-price scenario	NPV	Net present value
CRF	Capital recovery factor	NRTL	Non-random two-liquid model
CW	Cooling water	NZE	Net Zero Emissions Scenario
EC	Cell-performance scenario	OPEX	Operating expenditure
ECH	Electrochemical hydrogenation	OSBL	Outside battery limits
ECO	Electrochemical oxidation	PEM	Proton exchange membrane
ED	Electrodialysis	PRIMES	EU energy model behind REF2020
EHV	Extra-high voltage	REF2020	EU Reference Scenario 2020
EI	Emission intensity	RHE	Reversible hydrogen electrode
EI-A	HNO ₃ scenario, EU abated	SEC	Specific energy consumption
EI-B	HNO ₃ scenario, non-EU unabated	SMR	Steam methane reforming

Abbrev.	Definition	Abbrev.	Definition
ENTSO-E	European Network of Transmission System Operators for Electricity	STEPS	Stated Policies Scenario (IEA)
EoL	End of life	<i>t</i> 3HDA	<i>trans</i> -3-hexenedioic acid
EP	Electricity-price scenario	TEA	Techno-economic assessment
EPD	Environmental product declaration	TNO	Dutch Organisation for Applied Scientific Research
ERAA	European Resource Adequacy Assessment (ENTSO-E)	TRL	Technology Readiness Level
ETS	Emissions Trading System	TRY	Titer, rate, yield (fermentation)
FE	Faradaic efficiency	VAT	Value added tax
FLH	Full load hours	WACC	Weighted average cost of capital
FP	Feedstock-price scenario		

Symbols

Symbol	Definition (unit)	Symbol	Definition (unit)
B_{AA}	ETS benchmark, adipic acid (tCO ₂ e/t)	m	Maintenance rate on capital (–)
C_e	Purchased-equipment cost (\$)	M	Molar mass (kg/kmol)
C_{cell}	Cell-block installed capital (\$)	n	Electrons per mole product (–)
$C_{ETS}(t)$	ETS compliance cost (\$/a)	N	Project life (a)
C_{fix}	Total fixed capital (\$)	N_{prod}	Producing years (a)
C_{OPEX}^{fix}	Fixed operating cost (\$/a)	$\Delta \dot{n}_{solid}$	Net solid molar flow (kmol/h)
C_{OPEX}^{var}	Variable operating cost (\$/a)	o_t	Fixed-cost switch, year t (–)
$C_{stack}(t)$	Stack-replacement cost (\$/a)	p_{cap}	Capacity price component (€/MWh)
$\overline{CI}(LF)$	Mean CI of dispatched hours (gCO ₂ e/kWh)	p_{com}	Commodity price component (€/MWh)
$CI_c(t)$	Grid CI, country c , year t (gCO ₂ e/kWh)	$P_{CO_2,s}(t)$	Carbon price, scenario s (€/tCO ₂)
e_{N_2O}	N ₂ O factor, nitric acid (kg/t HNO ₃)	Q	Design annual production (kt/a)
e_{NH_3}	Upstream ammonia CI (kgCO ₂ e/kg)	Q_{cryst}	Crystalliser duty (kJ/h)
$\bar{e}_{GHG}$	Cradle-to-gate GHG intensity (kgCO ₂ e/kg)	r	Real discount rate (–)
e_{scope1}	Scope-1 emission intensity (tCO ₂ e/t)	r_{NH_3/HNO_3}	NH ₃ per unit HNO ₃ (kg/kg)
e_{scope3}	Feedstock emission intensity (kgCO ₂ e/kg)	t	Project year (–)
e_{steam}	Process-steam intensity (kgCO ₂ e/kg)	U	Overall heat-transfer coeff. (kJ/h m ² K)
EI_{HNO_3}	Nitric acid cradle-to-gate CI (kgCO ₂ e/kg)	V_{cell}	Cell voltage (V)
f_{cont}	Contingency factor (–)	$w_{elec}(t)$	Electricity per kg product (kWh/kg)
$f_{D\&E}$	Design-and-engineering factor (–)	$\alpha(t)$	Free allocation fraction (–)
f_{OSBL}	OSBL factor (–)	η	N ₂ O destruction efficiency (–)
F	Faraday constant (C/mol)	η_{mat}	Material efficiency (–)
g_{N_2O}	Unabated N ₂ O factor (kg/kg AA)	$\mu(LF)$	Price ratio, hours run (–)
GWP_{100}	100-year GWP (–)	ν	Product-to-feed mole ratio (–)
h	Individual film coefficient (kJ/h m ² K)	ϕ	Cell share of route electricity (–)
ΔH_{fus}	Enthalpy of fusion (kJ/mol)	ϕ_t	Capital draw, year t (–)
LF	Load factor (–)	ψ_t	Output fraction, year t (–)
$\dot{m}$	Mass flow (kg/h)		

Introduction

1.1. Background

Adipic acid (AA) is a high-volume industrial chemical used primarily as a monomer for nylon-6,6, which dominates demand in most market breakdowns [1, 2]. Global production reached approximately 4.5 Mt in 2025 [3], with a market of approximately \$4.9 billion in 2022 projected to reach \$6.6 billion by 2030 [4]. Because AA is embedded in large-scale polymer supply chains, changes in its production economics or emissions intensity propagate across material systems at scale. The dominant production route relies on benzene-derived intermediates, structurally coupling AA price formation to fossil value chains and introducing exposure to petrochemical market volatility [2, 5].

Beyond price volatility, the structural dependence of conventional AA production on imported fossil feedstocks introduces supply security risks that have gained prominence following geopolitical disruptions to European energy markets. This has reinforced interest in alternative production routes as a means of reducing import exposure, independent of climate policy considerations [6]. Conventional production also carries a substantial emissions burden, both direct and indirect: the oxidation step generates nitrous oxide (N_2O) as an inherent by-product, a greenhouse gas with a 100-year global warming potential of 273 relative to CO_2 [7], while the benzene and nitric acid it consumes carry emissions of their own.

Europe is the focus of this assessment because it is where the regulatory and economic tension surrounding AA production is among the sharpest globally. The EU Emissions Trading System prices the direct emissions of the adipic-acid production process. At the same time, high industrial electricity prices relative to competing regions create a countervailing challenge for electrified production routes. Active industrial policy aimed at retaining energy-intensive chemical production on the continent adds a further dimension to this regulatory environment. The analysis is anchored to the Antwerp–Rotterdam–Rhine–Ruhr Area (ARRRA), Europe’s largest chemical production cluster, accounting for roughly 40 % of EU petrochemical output and hosting the full value chain from feedstock import to polymer production [8].

1.2. Motivation and Problem Statement

In response to these economic and environmental drivers, alternative AA production concepts have been proposed, including electrochemical synthesis routes and biomass-derived pathways. Electrochemical routes target the N_2O problem directly by replacing the nitric-acid oxidation step, while biomass-derived pathways aim to eliminate fossil feedstock dependence and its associated scope-3 emissions entirely.

Assessing their competitiveness rigorously is non-trivial. Reported performance data are often obtained at laboratory scale, and techno-economic assessments rely on assumptions that are not consistently defined across routes. Differences in system boundaries, scale assumptions, and key performance indicator (KPI) definitions can dominate comparative conclusions, so apparent advantages for one route may partly reflect methodological choices rather than genuine process economics.

This thesis addresses eight specific gaps in the existing literature. They are ordered by the route each one decides, the conventional route first, then the electrified routes, then the biomass route, and last the two that bear on all three. (i) *ETS compliance cost and the abatement level it rests on*: to the author’s knowledge, no published techno-economic assessment (TEA) of AA production embeds EU ETS allowance costs into the levelised cost of adipic acid. Nor is that cost assessed together with the plant’s own N_2O destruction efficiency, which sets the direct emissions it is charged on, varied across the levels operating plants actually achieve. (ii) *Feedstock scope-3 emissions*: the role of supplier-choice variation in feedstock scope-3 emissions remains underexplored in prior route comparisons, particularly for HNO_3 (EU-abated versus non-EU unabated, with order-of-magnitude spread depending on N_2O abatement technology), H_2 (grey, blue, and green production types), and benzene (fossil versus bio-based supply). (iii) *Electricity price*: prior TEAs rely on generic electricity price assumptions rather than location-specific values at the appropriate consumption scale, and do not account for the country-specific cost structures and industrial policy instruments that cause effective prices to differ substantially across locations. (iv) *Grid carbon intensity*: prior work assumed fully renewable electricity procurement at constant full-load supply, which does not reflect the grid-mix reality faced by a baseload industrial plant, nor whether following the time-varying electricity price and grid intensity through partial-load operation lowers the levelised cost of adipic acid or scope-2 emissions. (v) *Cell degradation*: prior TEAs assume constant electrochemical cell performance over the plant lifetime, ignoring voltage degradation, Faradaic efficiency decline, and periodic stack replacement, all of which raise levelised costs and emissions relative to initial-performance estimates. (vi) *Product boundary*: existing biomass-route TEAs use inconsistent product endpoints, variously reporting *trans*-3-hexenedioic acid (*t*3HDA) or adipic acid (AA), which complicates direct cross-route comparison of levelised cost. (vii) *Heat integration*: the AA process models adopted as analytical baseline for this thesis implement only partial heat integration; the thermodynamic potential for further utility savings has not been quantified through complete pinch analysis. (viii) *Isolated dimension analysis*: existing route comparisons evaluate scenario dimensions one at a time and miss the correlated variation that real-world scenarios produce. Country profiles link electricity price to grid carbon intensity, and policy futures bundle market, procurement, and policy dimensions into internally coherent narratives.

These gaps shape both the process modelling and the techno-economic assessment: pinch analysis is applied to all three flowsheets, a catalytic hydrogenation reactor extends the biomass route from *t*3HDA to AA, and the remaining gaps are addressed through scenario-based variation of the dimensions above. Several bear on more than one route, so each is placed with the route whose standing it decides most.

1.3. Objective and Scope

The objective of this thesis is not only to assess whether alternative AA production routes can be economically and environmentally competitive, but to identify *which factors determine* their competitiveness and *where the dominant bottlenecks arise*. The study isolates the main technical and system-level drivers of levelised cost of adipic acid (LCOA), greenhouse-gas (GHG) intensity, and material efficiency across three pathways: the conventional nitric-acid oxidation benchmark, electrochemical oxidation (ECO) routes, and biomass-derived (BIO) routes.

The scope covers all three pathways under a consistent system boundary using a scenario-based approach. Routes are compared across production locations in the ARRRRA cluster (Belgium, the Netherlands, Germany), with the plant placed on a 2028–2035 commissioning window that sets the time-varying carbon-policy and grid trajectories the routes are exposed to, rather than serving as a separate comparison axis. The location choice supplies the country-specific delivered electricity cost and grid carbon intensity values and is itself the variation mechanism for these two dimensions; the remaining three tier dimensions vary independently of location: carbon price (sampled across the credible 2030 multi-model projection range), cell performance, and feedstock cost (reflecting each route’s commodity exposure). Three procurement-scenario inputs (hydrogen as grey/blue/green; HNO_3 as EU-abated or non-EU unabated; benzene as fossil or bio-based) carry paired price and scope-3 emission intensity values, capturing supplier-choice exposure across the routes.

Most inputs are held flat at their scenario tier values for the 20-year plant lifetime; cell-performance degradation evolves in every scenario, while grid carbon intensity and a rising net-zero-aligned carbon-price trajectory are tested as time-dependent sensitivity cases; EU ETS free allocation is carried as a flat bracket between full retention and complete withdrawal. Dimensions are varied both independently and in joint combinations, including narrative scenarios that bundle multiple dimensions into coherent futures (for example, “decarbonisation accelerates”

and “decarbonisation lags”), to test whether interaction effects alter the competitive ranking. The scenario values are projections of plausible futures rather than predictions of any single path; the goal is to diagnose drivers and bottlenecks across the credible range, not to deliver a single viability verdict.

1.4. Research Questions

The central research question of this thesis is:

How do conventional, electrochemical, and biomass-based adipic acid production routes compare on levelised cost and cradle-to-gate greenhouse-gas intensity under realistic European market and policy scenarios, and what are the dominant bottlenecks and viability conditions for each route?

This main question is addressed through eight sub-questions, each targeting a specific gap identified in the literature. They are grouped by the route whose standing they decide, the conventional route first, then the electrified routes, then the biomass route, and last the two that bear on all three and test whether the ranking survives the dimensions acting together:

1. How does pricing EU ETS allowance costs, with benchmark-based free allocation carried between full retention and complete withdrawal, affect the levelised cost of adipic acid and relative competitiveness of each production route across the modelled carbon-price scenarios, and how does varying the conventional route’s own N₂O destruction efficiency across the levels achieved in practice, which set the direct emissions that charge falls on, affect both indicators?
2. How does supplier-choice variation in feedstock scope-3 emissions affect each route’s greenhouse-gas intensity and route-comparative competitiveness, when addressed through discrete procurement scenarios for HNO₃ (EU-abated versus non-EU unabated), H₂ (grey/blue/green), and benzene (fossil versus bio-based)?
3. How do industrial electricity costs at the extra-high-voltage consumption scale, including country-specific policy exemptions, affect the economic competitiveness of the electrified routes (electrochemical and biomass-based) across the ARRA cluster?
4. What is the impact of replacing the fully-renewable electricity assumption with actual grid-mix carbon intensities on the greenhouse-gas intensity of the electrochemical and biomass-based routes, and does following the time-varying electricity price and grid intensity through partial-load operation lower the levelised cost of adipic acid or scope-2 emissions?
5. How does electrochemical cell degradation over the plant lifetime affect the levelised cost of adipic acid and greenhouse-gas intensity of the electrified routes relative to constant-performance assumptions?
6. How does evaluating both *trans*-3-hexenedioic acid and adipic acid as product endpoints within a single consistent framework affect the comparative economic and environmental assessment of the biomass route?
7. How does a complete pinch analysis affect the levelised cost of adipic acid and greenhouse-gas intensity of each route relative to partial heat integration?
8. How does joint variation across multiple scenario dimensions affect the competitive ranking, including country-coupled inputs (electricity price linked to grid carbon intensity) and narrative bundles combining policy and market dimensions?

Each sub-question corresponds to the gap of the same number in §1.2, and the Conclusions answer them in this same order.

1.5. Relation to Prior Work

This thesis builds on the process modelling framework developed in the MSc thesis of Edward Yeo, *Adipic Acid Production: Process Modeling of the Benchmark and Two Electrochemical Approaches* [9], which provides a consistent Aspen Plus baseline for the three production routes assessed here. The core process structures and baseline assumptions are retained; the eight gaps identified in §1.2 drive the extensions developed in this work.

Literature Review

This chapter reviews existing research on conventional, electrochemical, and biomass-based adipic acid production routes. The focus is on process architecture, performance metrics, techno-economic framing, and reported scale assumptions, with particular attention to comparability constraints that influence cross-route conclusions. The chapter integrates and evaluates existing work to identify unresolved gaps and structural modelling assumptions that motivate the methodological choices and scenario framework developed and further elaborated in subsequent chapters.

Literature search strategy

The literature for this review was identified through systematic searches of Scopus, Web of Science, and Google Scholar, complemented by industry pricing indices (BusinessAnalytiq, ChemAnalyst, S&P Platts) for current commodity prices, commercial life-cycle assessment (LCA) databases (CarbonCloud, ecoinvent v3) and industry environmental product declarations (EPDs) for emission intensities where peer-reviewed equivalents were unavailable, and EU regulatory texts (EU ETS, CBAM, REF2020) for policy parameters. Search terms targeted adipic acid techno-economic assessment, electrochemical and biomass routes, EU carbon-price projections, European industrial electricity costs, and electrolyser capital cost (CAPEX) scaling; reference lists of key papers and theses (in particular Yeo 2025) were checked for additional sources. Inclusion required quantitative process, cost, or emissions data directly applicable to one of the three production routes, or documented projections of scenario-dimension parameters.

2.1. Market Context and Policy Motivation

This section characterises the market and regulatory conditions that define the competitive context for AA production in Europe, covering supply security, carbon policy, and the price reference used as the TEA benchmark.

The structural vulnerability of conventional AA production to fossil supply disruptions has been illustrated repeatedly in recent years. The energy price shocks that followed geopolitical disruptions to European fossil fuel supply chains from 2022 onward showed how benzene-dependent value chains are exposed to supply-side shocks largely outside the control of individual producers [6, 10]. More recently, the 2026 Strait of Hormuz crisis has further demonstrated this exposure, with crude oil and LNG supply disruptions triggering significant price spikes in European energy and feedstock markets [11]. These episodes reinforce the case for evaluating alternative routes on their resilience to import dependency, independently of climate policy drivers.

Beyond supply security, climate policy shapes the competitive landscape. The CBAM extends EU ETS-equivalent costs to imports of HNO_3 from January 2026, and the full set of regulatory mechanisms (free-allocation phase-out, product benchmark, CBAM scope), the three carbon-price tiers used in the scenario framework, and the supplementary rising-price trajectory are developed in Section 2.4.1.

The scale and location of adipic-acid emissions

Adipic-acid production emits a large quantity of nitrous oxide, and most of it comes from a small number of plants outside Europe. A bottom-up estimate covering the 39 plants identified as operating in 2021 puts global N_2O from adipic acid at about 142 MtCO₂e on the AR6 GWP₁₀₀ of 273, of which about 134 MtCO₂e, or 94 %, is attributed to China, on the assumption that Chinese plants operate without abatement [12]. That assumption is the estimate's principal uncertainty, no plant-level abatement rates for those facilities having been published, but a peer-reviewed national inventory reaching it from Clean Development Mechanism records rather than by assumption arrives at a comparable 128 MtCO₂e of Chinese adipic-acid N_2O for 2020 [13]. China holds roughly half of world capacity [14, 15, 16], so 94 % of the emissions arise on about half the output, and what separates the fleets is the abatement level rather than the chemistry: destruction equipment costs money and returns nothing where the emission carries no price.

Against global production of about 3.07 Mt in 2021, that implies a global average of about 46 kgCO₂e per kg of adipic acid in residual N_2O alone. That sits far above published life-cycle estimates for the conventional route, which span 6.3 to 25.6 kgCO₂e/kg [17, 18] and assume 80 to 85 % N_2O destruction.

Market price and LCOA competitiveness threshold

A defensible TEA requires an external price anchor against which the LCOA of each route can be evaluated. Across 2025 the European adipic-acid price spanned approximately \$1.88–2.04/kg, settling near \$1.91/kg by year-end after a modest fourth-quarter decline (roughly €1.71–1.85/kg at the 1.10 USD/EUR rate held throughout this work, §3.3.3), with Northeast Asian prices significantly lower at around \$1.22/kg, reflecting the structural cost advantage of Chinese producers, who account for approximately half of global AA production capacity [14, 15, 16]. The European premium over global averages reflects import logistics, quality standards, and anti-dumping measures affecting Chinese supply [19]. Prices are historically volatile, having peaked above \$2.50/kg during the 2021–2022 energy and supply-chain crisis before returning to this \$1.88–2.04/kg full-2025 range [14, 15]. European spot rose again in early 2026, but the anchor is held at the 2025 range to remain on the same cost basis as the feedstock prices.

This full-2025 range is adopted as the competitiveness benchmark and held constant across the analysis, on the same 2025 cost basis as the feedstock prices used throughout; for each route, the gap between its LCOA and the band measures the cost reduction or structural change needed to reach parity.

2.2. Adipic Acid Production Pathways

Three production routes are reviewed in this section: the conventional nitric-acid oxidation process, the ECO route, and the biomass-based pathway via muconic acid fermentation. These three pathways were selected because they span the full spectrum from the established commercial benchmark to the most developed low-carbon alternatives, and because sufficient published process and cost data exist to parameterise the simulation models. For each route, a simplified process flow diagram is provided in which the dashed box indicates the system boundary of the Aspen Plus process model; steps shown outside the dashed boundary are included for process context but are treated as purchased inputs in the economic analysis.

2.2.1. Conventional Adipic Acid Production

The dominant industrial route to AA is the nitric-acid oxidation process based on petroleum-derived benzene feedstock, accounting for approximately 95–98 % of global production [20, 21]. Figure 2.1 shows the overall process structure, and the following paragraphs discuss each step and its main cost, energy, and emission drivers.

In the first step, benzene is catalytically hydrogenated to cyclohexane in either liquid-phase or gas-phase configurations, with feed purification required to manage catalyst-deactivating impurities [22, 23]. Cyclohexane is subsequently oxidised with air to produce KA oil (a mixture of cyclohexanol and cyclohexanone) via liquid-phase catalytic oxidation, single-pass conversion being kept low to maintain selectivity with unreacted cyclohexane recycled [24, 25, 21]. In the third step, KA oil is oxidised with concentrated nitric acid to adipic acid, with reported yields in the range of 92–96 % [20, 26]. The product is then recovered and purified by crystallisation-based separation: AA crystallises upon cooling, and purification commonly relies on recrystallisation and solid–liquid separation to meet high purity requirements [27, 22, 28, 29], while the remaining nitric-acid-containing liquors




Figure 2.1: Simplified process flow diagram of the conventional adipic acid production process. The dashed box indicates the Aspen Plus model boundary. Adapted from Yeo [9].

are separated from heavy wastes and recycled to the KA-oil oxidation reactor, closing the acid loop and reducing fresh nitric acid consumption [26, 21].

A key drawback of the nitric-acid oxidation step is the intrinsic formation of NO_x and N_2O as reaction by-products; these off-gases require treatment or catalytic abatement, with large-scale abatement systems reported by major producers [30, 21].

N_2O abatement performance

European producers abate this N_2O almost completely, and have kept improving: the five EU plants averaged about 94 % abatement in 2007/08 [31] and about 99 % by 2016/17. The spread behind that later average runs from 0.32 tCO₂e per tonne of adipic acid for the ten per cent most efficient to a median of 0.76 and a production-weighted mean of 1.14 [32], and the implied abatement with it (Table 2.1).¹ Those five sub-installations are the whole EU industry, one in France, three in Germany and one in Italy [33], all running cyclohexane oxidation with nitric acid [34, 31], so the spread is abatement practice rather than chemistry.

Table 2.1: What EU adipic acid plants actually emit: the distribution behind the EU ETS product benchmark, 2016/17 [32]. These are *direct* emissions inside the ETS boundary, which counts the residual N_2O together with process CO₂ and the fuel burned in the abatement unit itself, so they are not comparable with the published cradle-to-gate figures quoted below, which are roughly an order of magnitude larger in the same units.

Statistic, 2016/17	Direct emissions (tCO ₂ e/t)	Implied N_2O abatement ¹
Ten per cent most efficient	0.32	≈ 99.7 %
Median	0.76	≈ 99.2 %
Mean	0.98	≈ 98.9 %
Production-weighted mean	1.14	≈ 98.7 %
<i>This work, central case</i>	<i>0.89</i>	<i>99.0 %</i>

Those percentages are derived from emission intensities rather than measured. Directly verified figures do exist: for thermal destruction, the route adopted here, the 2006 IPCC guidelines give 96 % and the 2003 EU BREF 94–96 %, while carbon-credit project monitoring up to 2010 reports 99 % for one facility outside Europe and 100 % for the three European plants [34, 35, 36]. Those are measured but older, and cover verification periods of unstated length where the benchmark intensities are whole calendar years, so the sweep below is anchored on the benchmark and bounded above by them.

They reach those levels by different means, which is why the fleet figures are abatement rather than destruction: redundant thermal destruction at Krefeld-Uerdingen, redundant catalytic at the former Ludwigshafen unit, and conversion of the N_2O back to nitric acid at Chalmers [34, 31]. The two German plants added their second unit to buy availability rather than destruction efficiency, and that redundant configuration is the one costed here. What that abatement costs is small next to what it removes: carbon-credit project monitoring reports 0.10–0.40 €/tCO₂e for a single first train [34], far below any carbon price carried in §2.4.1.

How much that abatement premise is worth can be read off the published inventories themselves. Within any single one, the abatement assumed accounts for nearly the whole of the spread in the cradle-to-gate estimates

¹Implied abatement uses this work's scope-1 relation, $81.9(1 - \eta) + 0.073 \text{ tCO}_2\text{e/t}$, with 81.9 the stoichiometric 0.30 kg N_2O per kg of product at the AR6 GWP₁₀₀ of 273. It is indicative: the ETS boundary also carries the abatement unit's own fuel CO₂, so the value reads low, and generation below stoichiometry at 0.27 kg/kg would put the median at 99.1 %.

quoted in §2.1. Corona et al. [17] report 22.6 kgCO₂e/kg for a world-average plant at 85 % against 12.9 for a modern US one at 97 %, the two differing almost only in unabated N₂O, 13.4 against 3.58 kgCO₂e/kg, with identical cyclohexane (1.33) and nitric-acid (4.63) contributions; Aryapratama and Janssen [18] reach the same conclusion from a different inventory, 25.6 kgCO₂e/kg at the ecoinvent default of 80 % against 9.4 at 98 %. Across inventories it is not, the studies resting on different geographical bases and all characterising N₂O at a GWP₁₀₀ of 298 against the AR6 value of 273 used here. They indicate the size of the abatement effect rather than what a European plant would report.

Because achieved performance spans a real range, from the 94–96 % BAT floor to near-complete destruction, this thesis does not fix a single value: it takes 99 % (the modern fleet-median level) as the central case and sweeps abatement across 95–100 % to assess how strongly the conventional route's scope-1 emissions, GHG intensity and ETS position depend on it.

Upstream emissions from nitric acid supply

Concentrated nitric acid is the oxidising agent in the KA-oil oxidation step and is typically sourced from dedicated producers rather than made on-site [26]. Its production falls outside the system boundary and is treated as a scope-3 input; the N₂O emitted at the supplier's plant is that supplier's EU ETS obligation, not the AA plant's. Nitric acid is produced via the Ostwald process, in which catalytic ammonia oxidation generates N₂O as an unavoidable side-product; pre-abatement formation rates vary widely and multi-tier abatement technology can reduce them by up to 99 % [31, 37, 38]. Yeo [9] found HNO₃ emission intensity to be the dominant GHG driver of the conventional route: sourcing from a low-emission supplier reduces AA's CO₂e to approximately 2.3 kg/kg, broadly comparable to electrochemical routes. Within the EU, BREF/IPPC mandates abatement, but the BAT range still spans nearly an order of magnitude [31], making supplier choice a material procurement decision. The two procurement scenarios parameterising this intensity are defined in Section A.5.3.

2.2.2. Electrochemical Adipic Acid Production

ECO has been proposed as an alternative to the conventional nitric-acid oxidation step in adipic acid production. The motivation is twofold: replacing nitric-acid oxidation eliminates the scope-1 N₂O emissions that represent a direct EU ETS liability for the conventional route, and removes the need to purchase HNO₃ entirely, eliminating the scope-3 burden that is the dominant GHG driver of the conventional route [9, 39, 40]. In the electrochemical route, electricity and water are used to oxidise KA oil intermediates to adipic acid under relatively mild conditions. Figure 2.2 illustrates a simplified process flow diagram for the electrochemical adipic acid production route. Similar to the conventional pathway, the upstream section consists of benzene hydrogenation to cyclohexane followed by air oxidation to produce KA oil. In contrast to the nitric-acid oxidation step, the electrochemical process converts these intermediates to adipic acid in an electrocatalytic reactor using an alkaline electrolyte.




Figure 2.2: Simplified process flow diagram of the electrochemical oxidation (ECO) adipic acid production route. The dashed box indicates the Aspen Plus model boundary. Adapted from Yeo [9].

In the electrochemical process, KA oil is fed to a continuous-flow membrane electrolyser in an alkaline electrolyte (typically KOH), where sequential anodic oxidation converts cyclohexanol (CHA) to cyclohexanone (CHN) and subsequently to adipic acid via C–C bond cleavage [9, 39, 40]. Hydrogen evolution at the cathode yields H₂ as a co-product that can offset feedstock cost through internal recycling [41, 40].

Electrocatalyst development has been a major research focus for this pathway. Early studies reported relatively low current densities and moderate yields due to competing oxygen evolution reactions and limited catalytic activity

[9]. More recent work has significantly improved performance, with reported current densities up to 300 mA cm^{-2} at Faradaic efficiencies exceeding 80 % using vanadium-modified nickel layered double hydroxide catalysts [39], and stepwise selectivities of 96.5 % (cyclohexanol→cyclohexanone) followed by 93.6 % (cyclohexanone→adipic acid) through controlled NiOOH active-site engineering [42, 40]. These improvements have motivated renewed interest in the electrochemical route as a potentially viable alternative to the conventional nitric-acid process.

Downstream processing of the electrochemical route differs from the conventional pathway because adipic acid is initially formed as a dissolved species within the alkaline electrolyte. As a result, neutralisation and crystallisation steps are required to recover solid adipic acid product and regenerate the electrolyte solution for recycling within the process [9, 43]. Efficient separation and electrolyte management therefore play an important role in determining the overall techno-economic feasibility of this route.

A further process-level challenge specific to electrochemical systems is progressive performance decay: catalyst deactivation, membrane instability, and rising internal resistance translate into increasing cell voltages and periodic stack-component replacement [44, 45, 46, 47]. Stack performance therefore cannot be assumed constant over the plant lifetime; the durability data and replacement-interval reasoning are developed in Section 2.4.5.

Specific electricity demand of the ECO route

For the ECO route, the direct electrical demand of the electrochemical reactor is the dominant operating expenditure driver and the primary sensitivity parameter for cross-route comparison. The specific electricity consumption scales with the cell voltage V_{cell} and inversely with the Faradaic efficiency FE; the model equation ($n = 6$ electrons per adipic-acid molecule, CHN anodic half-reaction [9]) is given in the methodology (Eq. (3.1)). Under the baseline conditions of the Yeo (2025) process model, the direct cell electrical demand is 2.57 kWh/kg AA [9]; Liu et al. report a comparable value of 2.43 kWh/kg AA under their experimental operating point [39]. This figure is, however, highly sensitive to the assumed Faradaic efficiency and cell voltage: the range of reported experimental conditions spans a factor of three or more in specific energy consumption (see Section 2.2.2), making the electrochemical operating point the primary sensitivity parameter for cross-route comparison.

Downstream purification imposes a non-trivial additional energy demand: processing dilute aqueous product requires substantial separation duties, and electrolyte recovery via electrodialysis adds further electricity consumption [9]. Total ECO route energy intensity therefore substantially exceeds the direct reactor figure, and is sensitive to AA product concentration and electrolyte management strategy, underscoring the importance of heat integration in any realistic assessment.

Electrochemical operating-point scenarios

Cell-performance scenarios test how route LCOA responds to the $(j, \text{FE}, V_{\text{cell}})$ operating point at which industrial-scale cells eventually settle, given the wide range of laboratory-reported values.

Electrochemical route performance is characterised by three coupled figures of merit: current density (j), Faradaic efficiency (FE), and cell voltage (V_{cell}). These parameters are not independent, and their joint values are central to the cost of electrified production pathways [48, 49]. Specific electricity consumption scales with cell voltage and inverse Faradaic efficiency: higher V_{cell} increases energy per unit charge, while lower FE requires more charge per mole of product. Current density controls the electrode area for a given production rate: increasing j reduces stack area (and therefore CAPEX) but raises V_{cell} through ohmic and kinetic overpotentials (raising operating cost, OPEX). Cathode hydrogen co-production is a further economic lever: if internally consumed for benzene hydrogenation, it can meaningfully offset feedstock cost [40]. Yeo's (2025) baseline did not credit this cathodic H_2 co-production; this thesis includes it as an avoided purchase, offsetting the benzene-hydrogenation H_2 demand internally, with any residual demand purchased (methodology §3.3.2).

Reported experimental data for KA-oil electrooxidation span a wide range. Early studies recorded Faradaic efficiencies below 50 % at low current densities, while more recent optimised catalyst systems achieve FE values of 60–80 % at substantially higher current densities [39, 40]. All available performance data originate from laboratory-scale experiments with small-area electrodes and controlled conditions. Whether industrial-scale cells will ultimately perform within, above, or below this range is an open question: scale-up introduces mass-transport non-uniformities and accelerated degradation, but reactor engineering and catalyst development in mature electrochemical processes (water electrolysis, chlor-alkali, adiponitrile production) have historically exceeded early laboratory results. Degradation is treated as a physically distinct, independent parameter (Section 2.4.5).

ECO competitiveness is therefore not a fixed property of the process chemistry, but depends substantially on both the assumed operating point and the electricity system in which the plant operates.

2.2.3. Biomass-Based Adipic Acid Production

Biomass-based pathways address the two fundamental problems of the conventional route: dependence on fossil aromatics as feedstock, and the use of nitric acid (HNO_3) as an oxidant, which generates N_2O as a by-product. Of the biomass-based routes reported in the literature, two have been studied in sufficient depth to support techno-economic assessment: thermocatalytic hydrogenation, which converts purified muconic acid (MA) directly to adipic acid, and electrochemical hydrogenation (ECH), which operates in the unpurified broth and selectively produces *trans*-3-hexenedioic acid (*t*3HDA), a mono-unsaturated intermediate [50, 51]. Both share the same upstream step and the same environmental advantage: MA is produced biologically through fermentation, replacing fossil aromatics with renewable carbon and eliminating nitric acid entirely, thereby removing both the direct (scope-1) N_2O emissions at the AA plant and the upstream (scope-3) N_2O emissions from HNO_3 production at the supplier [52, 53]. The central trade-off between the two routes is that thermocatalytic hydrogenation produces AA directly but requires costly MA purification first, whereas ECH avoids purification but terminates at *t*3HDA rather than AA. This thesis models the biomass route with ECH upgrading to *t*3HDA, illustrated in Figure 2.3. A further extension, converting the crystallised *t*3HDA to AA via a mild catalytic hydrogenation step, is identified as a viable option; the experimental basis and process rationale for this extension are discussed in the sections below, with its techno-economic implications addressed in Section 2.4.4.




Figure 2.3: Simplified process flow diagram of the biomass-based ECH route via muconic acid fermentation and electrochemical hydrogenation to *t*3HDA. The dashed box indicates the Aspen Plus model boundary. Adapted from Yeo [9].

Muconic acid fermentation

The process begins with the biological production of muconic acid through fermentation. Muconic acid is a six-carbon dicarboxylic acid that can exist in several isomeric forms, with the *cis,cis* configuration being the most commonly produced biologically [9]. The concept of producing adipic acid from fermentation-derived muconic acid was first demonstrated by Draths and Frost in the 1990s using genetically engineered microorganisms capable of converting carbohydrates into MA [54]. Various microbial strains have since been investigated, including *E. coli*, *C. glutamicum*, and *P. putida*, each capable of metabolising feedstocks ranging from glucose to lignin-derived aromatics [55, 53, 51].

After fermentation, the broth contains muconic acid along with microbial biomass, salts, and other dissolved species. Filtration and clarification steps are therefore required to remove solids and produce a clarified MA stream suitable for downstream processing [9]. At this point the process diverges: the clarified MA broth can be upgraded via two routes that share the same AA product target but differ in purification strategy and upgrading chemistry, described in the following sections.

Electrochemical hydrogenation to *t*3HDA

In the biomass route as modelled by Yeo [9], muconic acid is reduced electrochemically at the cathode surface using protons from the aqueous medium. The reaction proceeds directly in the fermentation broth, without requiring prior MA purification [56, 57].

The ECH reaction selectively produces *t*3HDA through mono-hydrogenation of one double bond in muconic acid; further electrochemical reduction to fully saturated AA is barely observed, because electrode conditions

tolerant of broth impurities favour the mono-hydrogenation step [57, 9]. A second ECH step to complete the hydrogenation to AA has been demonstrated only at 18 % Faradaic efficiency in clean model solutions and is not yet validated in fermentation broth [9]. At current technology readiness the biomass route therefore terminates at *t*3HDA, which is the product boundary adopted in the base model of this thesis. A potential extension to adipic acid via a catalytic hydrogenation step is discussed in Section 2.4.4.

*t*3HDA has no established commodity market and is currently available only as a specialty research chemical. Whether a price premium over conventional AA is achievable depends on the value of its retained C=C double bond: unlike fully saturated AA, *t*3HDA can be incorporated into polyamides via thiol-ene chemistry to introduce tailored functionalities (hydrophobicity, flame retardancy, antimicrobial properties), giving access to “bioadvanced” nylon 6,6 grades with enhanced performance [56, 57]. Whether this translates to a commercially realisable price premium has not been demonstrated at scale; the TEA implications are discussed in Section 2.4.4.

Because water in the broth serves as the proton source, hydrogen atoms are generated directly at the cathode surface during the electrochemical reaction, so no external hydrogen gas supply is required. The Faradaic efficiency (FE) of the ECH step, defined as the fraction of charge directed to MA hydrogenation rather than to the competing hydrogen evolution reaction, determines electricity consumption per kilogram of product and carries comparable uncertainty to the ECO route: reported values span roughly 60 % for an air-exposed lead cathode in dilute model solution to near 100 % once dissolved oxygen is purged or the feed is a de-aerated spent broth, and fall below 15 % on metals with low hydrogen overpotential [56, 50, 57]. None of those conditions is the industrial operating point adopted here, which is why the tiers below are set from the efficiency measured at that point rather than from the spread of reported values. An analogous three-tier performance range is therefore applied to both electrochemical routes (ECO and ECH), with route-specific parameter values calibrated separately given the distinct chemistries and operating conditions involved.

Following hydrogenation, the product mixture undergoes crystallisation and solid–liquid separation to recover *t*3HDA. Acidification using sulphuric acid is applied to promote crystallisation and improve product recovery [9]. Crystallisation serves a dual function here: it recovers *t*3HDA as the product and simultaneously removes the biogenic impurities that remain from the fermentation broth. The recovered solid is therefore both the product of the base model and a purified substrate suitable for downstream catalytic hydrogenation, should the process be extended to adipic acid.

Route extension: catalytic hydrogenation of *t*3HDA to adipic acid

Extending the model boundary to AA enables a direct product-equivalent comparison across all three routes, removing the product-boundary mismatch. The crystallisation step already in the base model provides the key enabler: acidification and solid–liquid separation during *t*3HDA recovery remove biogenic broth impurities alongside the product [9, 52], leaving the crystallised solid as an already purified substrate for a downstream catalytic hydrogenation reactor. This is distinct from a one-pot MA-to-AA approach: Capelli et al. showed that buffer salts in a synthetic fermentation broth reduced AA yield from approximately 100 % to 73 % after 180 minutes, confirming that feed purity is a prerequisite for the catalytic step [58]; in the proposed extension that purification is already accomplished by the crystallisation step, avoiding the catalyst poisoning by broth impurities that dominates the cost of the thermocatalytic route [57, 50].

The experimental basis comes from two studies by Capelli et al., both using *trans,trans*-muconic acid (MA) as feed rather than crystallised *t*3HDA directly; an important caveat since neither study demonstrates hydrogenation of *t*3HDA as an isolated substrate. The 2019 study (Pd/AC, 5 wt%, water, 70 °C, 1 bar H₂) achieved full MA conversion within 60 minutes at 14 mmol L⁻¹ substrate concentration, with an AA yield of 95 % [59]. The 2017 study (Pt/C, 5 wt%, 70 °C, 4 bar H₂) achieved 100 % conversion and selectivity within 1 hour, with the catalyst remaining active over at least ten consecutive cycles [58]. Both studies confirm a sequential two-step hydrogenation mechanism with *t*3HDA as the intermediate; the polishing step corresponds only to the second hydrogenation, so a single mole of H₂ per mole of *t*3HDA suffices. The extended route therefore comprises two sequential steps: ECH (electrochemical, *cc*MA → *t*3HDA) followed by catalytic hydrogenation (*t*3HDA → AA), with that second unit representing one additional reactor operation. The substrate concentrations used (14–70 mmol L⁻¹) are low by industrial standards and scale-up to continuous operation has not been demonstrated, but the conditions are mild (70 °C, 1–4 bar H₂, water as solvent), suggesting that the additional equipment cost is modest relative to the main process. The economic implications are discussed in Section 2.4.4.

Because the polishing step changes what leaves the plant, the biomass route is evaluated at both endpoints and carries two labels wherever the endpoint matters; the comparison itself is between three routes throughout, and

each label names the product it delivers. **CON** is the conventional route, nitric acid oxidation of KA oil to adipic acid, and **ECO** its electrochemical counterpart, anodic oxidation of the same KA oil to the same product. **BIO** names the biomass route as a whole, and a subscript gives the endpoint where it matters: **BIO**_{t3HDA} stops at the intermediate, **BIO**_{AA} carries it one step further to adipic acid by hydrogenating that intermediate. The subscript is the product, so a table or figure carrying these labels needs no further note of what each route delivers. Three of the four labels therefore name adipic acid as the product and one does not, which is the asymmetry the product-boundary discussion of §2.4.4 turns on, and the reason the biomass route is reported both ways throughout.

Thermocatalytic hydrogenation to adipic acid

A distinct alternative proceeds via thermocatalytic hydrogenation of muconic acid: purified MA reacts with molecular hydrogen over Pd- or Rh-based catalysts under mild conditions, converting it directly to AA [52, 59]. As noted above, this route requires full MA separation and purification from the fermentation broth before the catalytic step can proceed, as broth impurities irreversibly poison the precious-metal catalyst [57]. This purification is the dominant cost item, contributing up to 60 % of total production cost [50]. The TEA performance of this route relative to the ECH pathway is discussed in Section 2.4.4.

The two routes differ structurally in their primary cost driver. The thermocatalytic route is constrained by MA separation, a largely fixed purification cost insensitive to energy prices or carbon policy. The ECH route is instead governed by fermentation TRY and electricity price, both of which are scenario-sensitive and addressed in the scenario framework of this thesis.

The two routes also differ in their hydrogen requirement. The thermocatalytic route requires purchased H₂ as a reactant; the ECH step does not, though the fraction of charge not directed to MA hydrogenation drives the competing HER, producing H₂ as a by-product that could partially offset electricity costs [50]. This by-product stream is not included in the baseline model of Yeo [9], which treats Faradaic efficiency losses as an electrical efficiency parameter without modelling an explicit H₂ output.

Feedstock options

The fermentation step is feedstock-flexible: muconic acid can be produced from a range of renewable carbon sources, and the choice of feedstock affects both the economics and the carbon intensity of the overall process. Sugar-derived feedstocks (principally glucose from corn starch or sugarcane) were the basis for the first demonstration of biological MA synthesis [54] and remain the most widely studied at laboratory scale. Lignocellulosic feedstocks offer a lower-cost and lower-land-use alternative: lignin is an abundant aromatic by-product of pulping and biorefinery operations, and engineered strains of *Pseudomonas putida* can convert lignin-derived aromatics into muconic acid via the β -ketoadipate pathway [52, 53]. The baseline process uses glucose as the fermentation substrate, consistent with Yeo (2025) [9]; this single-pathway treatment is held throughout because the published ECH performance characterisation [57, 39] is calibrated for glucose-derived broth and does not extend to alternative substrates.

2.3. Cross-Route Techno-Economic Comparison

The three routes reviewed above differ not only in absolute levelised cost of adipic acid but in the *structural nature* of the uncertainty governing whether each cost gap can be closed: treating them as a single ordered ranking is misleading because each pathway's viability hinges on a different lever and responds differently to the same policy or market change. Published TEA studies differ in system boundaries and process assumptions in ways that substantially influence the reported minimum selling price (MSP, the quantity this thesis reports as LCOA) and GHG outcomes, rendering direct cross-route comparison methodologically invalid without a shared framework; the specific methodological gaps motivating this thesis are detailed in Chapter 1.

To enable a valid comparison, all three pathways are evaluated against the same key performance indicators within the same system boundary. This thesis uses the levelised cost of adipic acid (LCOA, \$/kg_{AA}) as the primary economic KPI, kgCO₂e/kg AA as the environmental KPI, and material efficiency (actual-to-theoretical AA yield) as the process performance indicator. The system boundary runs from feedstock conditioning through reaction, separation, and utility supply to the purified AA product gate; downstream polymerisation is excluded. This product-gate boundary follows Yeo (2025) [9] and the system-boundary convention of the international LCA standards ISO 14040/14044 [60, 61]; differences in which unit operations are included within the boundary can otherwise dominate reported MSP and GHG differences.

2.3.1. Prior Techno-Economic Findings

The conventional HNO_3 -oxidation route is well-documented in the process economics literature, with multiple published TEA studies establishing a consistent cost and emissions baseline [28, 62]. By contrast, published TEA studies on the ECO and biomass routes remain sparse, partly because both pathways are still at early technology readiness levels. The most directly relevant baseline spanning all three routes is Yeo (2025), which provides Aspen Plus models under a consistent scale and boundary definition [9]. Table 2.2 summarises the key economic and environmental outputs from this baseline, primary cost drivers, and the uncertainties that motivate the scenario extensions in this thesis.

Table 2.2: Cross-route comparison: baseline economic and environmental performance (Yeo 2025 [9]), primary cost drivers, and key uncertainties for the three adipic acid production pathways. All routes modelled at a single location (Rotterdam, NL) at a fixed baseline electricity price of $\approx \$0.16/\text{kWh}$ ($\$44/\text{GJ}$) with no carbon price; HNO_3 at industrial-average emission factor ($\approx 3.9 \text{ kgCO}_2\text{e/kg HNO}_3$, $\approx 60\%$ N_2O abatement). ² ECO GHG at $50 \text{ gCO}_2\text{e/kWh}$ renewable supply assumption [9].

Route	MSP (\$/kg)	GHG ($\text{kgCO}_2\text{e/kg}$)	Primary cost driver	Key uncertainty / baseline limitation
Conventional (HNO_3)	1.58	4.79	Feedstock cost Benzene OPEX dominates; no carbon price in baseline	HNO_3 emission factor Highly sensitive to supplier abatement status (§A.5.3); Yeo omits ETS cost and supplier variation
Electrochemical (ECO)	2.33	2.20 ²	Electricity + cell performance Electricity dominates OPEX ($\$0.77/\text{kg AA}$); ECO reactor and ED largest CAPEX	Operating point $\times$ electricity price Cell kWh/kg AA spans $\geq 3\times$ across lab operating points; no degradation modelled; grid carbon intensity not ARRRA-anchored
Biomass (<i>t</i> 3HDA)	3.97	1.18 ¹	MA feedstock cost; ECH electricity (secondary) Raw materials dominate at $\$2.05/\text{kg t3HDA}$ at current titres	Fermentation TRY + product boundary ECH Faradaic efficiency $\approx 40\%$ (experimental); fermentation titre unresolved; <i>t</i> 3HDA-to-AA extension not consistently evaluated

¹ MSP and GHG for *t*3HDA as sold product, not AA.

Yeo (2025) does not benchmark against a single price. He notes that comparing an MSP with real market prices is not straightforward given the volatility, places global prices between $\$1.03$ and $\$2.05/\text{kg}$ depending on region, and reports European estimates for Q2 2025 spanning $\$1.4$ – $\$1.8/\text{kg}$, within which his conventional MSP of $\$1.58/\text{kg}$ falls mid-range [9]. His reference is therefore partly a world price, whose lower bound reflects the Chinese supply that prices well below European levels ($\$2.1$); the benchmark adopted here is European only, which is one reason it sits above his. Against his own reference the electrochemical route at $\$2.33/\text{kg}$ and the biomass route at $\$3.97/\text{kg}$ both exceed it, giving premiums of 47 % and 151 % over his conventional figure.

The three premiums are not equivalent in character. Only the conventional route falls within the market range at all, and that is not a margin: a levelised cost of adipic acid is a break-even, so a producer earning its cost of capital prices above it. The electrochemical premium is driven by electricity cost and cell performance. The biomass premium is driven by MA feedstock cost at current fermentation titres, where the purchased broth volume per unit of product is large and is not reduced by lower electricity prices or carbon policy; as titre improves and that share falls, electricity price becomes the secondary driver through the ECH step's own demand. That route also reports an MSP for *t*3HDA rather than adipic acid, preventing a direct AA-to-AA comparison. Each gap therefore closes by a different intervention, not by variations of a single optimisation.

The GHG comparison requires consistent scope treatment. The ECO route value of $2.20 \text{ kgCO}_2\text{e/kg}$ assumes $50 \text{ gCO}_2\text{e/kWh}$ renewable electricity [9], well below both the 2024 EU-27 average of $213 \text{ gCO}_2\text{e/kWh}$ and the grid intensity at the modelled location, Rotterdam (Netherlands, $261 \text{ gCO}_2\text{e/kWh}$ in 2024; Table A.2), so scope-2 emissions are substantially larger even at the specific site Yeo modelled. Yeo (2025) does perform a grid carbon intensity sensitivity from 8 to $256 \text{ gCO}_2\text{e/kWh}$ [9, Sec. 3.4.4], confirming that GHG outcomes are highly sensitive to grid intensity, but the analysis covers a single fixed location and does not combine grid carbon intensity variation with electricity price or carbon price scenarios. The published ECO GHG figure therefore cannot be taken as an ARRRA-representative value; it is a lower bound anchored to an electricity assumption no ARRRA country currently meets. The structured scenario framework in Section 2.4.7 addresses this by varying the primary cost drivers identified in Table 2.2 systematically and in combination.

None of these baseline TEAs has performed a complete pinch analysis; Yeo (2025, §6.1.6) flags this as a limitation [9], addressed quantitatively in the methodology chapter.

2.3.2. CAPEX Estimation Approaches

A standard approach in process-level TEA literature is to size major equipment items from mass and energy balance outputs and estimate costs using built-in cost correlations within simulation tools such as Aspen Plus Economic Analyzer [9]. Capital costs are annualised using standard process economics conventions, with contingency and indirect costs embedded in the total installed cost.

The accuracy of such estimates depends on the level of design detail available. AACE International 18R-97 [63] rates a cost estimate as a whole, on how far the design behind it has been taken: an estimate made when 1–15 % of the design is complete is Class 4, one made at 0–2 % is Class 5. The rating belongs to the estimate, not to any single piece of equipment in it. A process-level TEA based on a complete flowsheet with sized equipment but without detailed engineering is therefore a Class 4 estimate. Class 4's published accuracy is an asymmetric envelope, roughly –15 to –30 % on the low side against +20 to +50 % on the high side, rather than a single symmetric figure; the high side is wider because cost overruns on less-defined equipment are more likely than under-runs. Towler & Sinnott [64] (Ch. 7) reproduce this as a simplified symmetric $\pm 30\%$, and NETL [65] collapses it to a single low/high pair for account-level estimating. Comparable electrochemical TEA studies confirm the same order of uncertainty [66].

2.3.3. Scale Assumptions and Reference Capacity

Scale assumptions differ structurally across the three routes. For the conventional route, reference scales range from the ~100 kt/a adopted in published TEA work [62] to the 400 kt/a nameplate of recently commissioned commercial capacity in the nylon-6,6 chain [67]; conventional unit operations benefit strongly from classical economies of scale. Electrochemical systems scale primarily through modular parallelisation, so stack-level economies of scale are generally weaker than for conventional large petrochemical unit operations, while balance-of-plant equipment may still benefit from scale; this structural difference should be reflected when comparing against large-scale petrochemical benchmarks. Biomass-based pathways face an additional constraint: fermentation at low product titres requires large reactor volumes, and feedstock logistics limit practical plant capacity to approximately 120 kt/a [9, Sec. 3.2.2], in line with the scale studied by Mokwatlo et al. (2024) [51]. A bio-based adipic acid pilot is under development by Toray and PTT Global Chemical, at a stated few thousand tonnes per year by 2030 [68]. Comparative assessments require a consistent functional unit and product-oriented system boundary, per the international LCA standards ISO 14040/14044 [60, 61]; for AA production this means purified AA at the product gate, with downstream polymerisation excluded.

This thesis adopts a plant capacity of **100 kt/a** for all three routes, following the scale selected by Yeo (2025) [9, Secs. 3.2.2, 3.4.2]. This scale falls within the lower end of the conventional route reference range, aligns with recent European capacity exits (BASF closed its ~115 kt/a adipic acid plant at Ludwigshafen in 2025, §2.4.2) [8, 69], and remains comfortably below the biomass logistics constraint. All electricity demand estimates and electricity cost scenario assignments in the following sections are derived from this basis.

2.4. Extensions to Existing Techno-Economic Assessments

The baseline comparison in Section 2.3 identifies several dimensions along which existing TEA studies do not fully characterise route competitiveness. The subsections below take each in turn: EU ETS carbon pricing (price levels, free-allocation phase-out and CBAM scope); the European industrial context and grid-access constraints in the ARRRRA cluster; extra-high-voltage electricity costs and grid carbon intensity; the biomass route product boundary (r3HDA versus AA); and electrochemical cell performance and degradation. Each states the level adopted and the choice made; the underlying price, emission-factor and durability evidence is collected in Appendix A. Sections 2.4.6 and 2.4.7 then consolidate these into the per-feedstock treatment summary and the scenario framework applied in the methodology.

2.4.1. Carbon Pricing and EU ETS

Carbon price is treated as a sensitivity dimension given the structural exposure of the conventional route to rising ETS costs. This subsection establishes the literature evidence base for the three carbon price levels (denoted *CP-Low*, *CP-Central*, and *CP-High*); a similar three-level tier structure is applied to electricity price, grid carbon intensity, feedstock cost, and electrochemical cell performance in the sections that follow.

The EU ETS is the primary carbon-pricing instrument relevant to this thesis. Allowance prices rose from roughly €25/tCO₂ in 2020 to peaks above €90/tCO₂ in 2023 amid Fit-for-55 cap-tightening reforms and the 2022 energy crisis (gas-to-coal switching in the power sector raised allowance demand) [70], then settled into the €70–80/tCO₂ range through 2024–2026 as oversupply tempered demand.

Future carbon-price trajectories depend on EU ETS cap tightening, free-allocation phase-out, Market Stability Reserve (MSR) absorption of surplus allowances, and the pace of industrial decarbonisation [71, 72]. Mainstream 2030 forecasts sit well above that spot range: the models compiled by Ariadne span €84–160/tCO₂, five of the six between 130 and 160 [73], with Pietzcker et al. at €129 [74], the IEA NZE 2030 advanced-economy value [75] and ABN AMRO’s 2030 baseline [71] inside that span. Tier assignments are given in Table 2.3.

Table 2.3: Carbon price scenario parameters (*CP-Low* / *CP-Central* / *CP-High*) sampled from the credible 2030 range established in the literature.

Scenario	Price (€/tCO ₂)	Basis
<i>CP-Low</i>	75	Midpoint of the 2024–2026 EU ETS spot range (€70–80/tCO ₂) [70] and the CIP (2026) Slow Transition trajectory [76]; below the lowest 2030 model estimate [73].
<i>CP-Central</i>	115	IEA Net Zero 2030 advanced-economy value (~USD 130/tCO ₂), the mainstream ambitious-decarbonisation pathway [75].
<i>CP-High</i>	145	Within the upper cluster of 2030 model estimates, €130–160 [73]; ABN AMRO baseline 2030 [71]; CIP (2026) Net Zero 2030 starting value (~€140/tCO ₂) [76].

Beyond 2030, carbon-price trajectories become increasingly uncertain and deviate across peer-reviewed publications and industry scenario reports. CIP (2026) [76] is the most recent industry-scenario report providing explicit post-2030 EU+UK CO₂-price trajectories at the time of writing, built on the Balmore optimisation model with input from Ea Energianalyse, and is adopted here as one trajectory among many for sensitivity analysis rather than as the predicted post-2030 price path. It gives three CO₂-price trajectories bracketing the tier spread used here, with their lowest-ambition scenario (Slow Transition) flat at ~75 €/tCO₂ (matches *CP-Low*) and their highest-ambition scenario (Net Zero) rising 140→210 €/tCO₂ (the 2030 starting value matches *CP-High* at €145, the 2050 endpoint is consistent with the IEA NZE 2050 advanced-economy assumption of ~€225/tCO₂ [75]). The Net Zero trajectory is adopted as the supplementary sensitivity (Section 3.5.1), testing the impact of a rising CO₂-price trajectory under the most ambitious decarbonisation scenario; the trajectory itself, alongside the other two CIP scenarios and the three static thesis tiers, is plotted in §A.4, Figure A.3.

Free allocation and the product benchmark

The EU ETS shields sectors at risk of carbon leakage, where an EU carbon cost could push production to countries applying a lower carbon price or none at all, by issuing free allowances up to a product-specific benchmark set from the most efficient installations; emissions above it must be purchased, emissions below generate net ETS revenue. Commission Implementing Regulation (EU) 2026/1412 of June 2026 sets the adipic acid benchmark at $B_{AA} = 1.40$ tCO₂e/t for 2026–2030 [77], down from 2.12 in 2021–2025 [78]; a plant first producing in 2030 is allocated on the later value, which is used throughout this thesis. Because the update is capped at a 50 % reduction the benchmark lags actual performance: the ten per cent most efficient installations averaged 0.34 tCO₂e/t over 2021–2022 [77], so a compliant plant is allocated well above what the best installations emit. Adipic acid qualifies for that allocation: it is listed as a carbon-leakage sector under NACE 20.14, other organic basic chemicals [79]. How the allocation enters the cost of each route is set out in methodology §3.3.2.

Carbon Border Adjustment Mechanism (CBAM)

The EU’s answer to leakage is shifting from shielding its own producers to charging imports instead. CBAM applies EU ETS-equivalent carbon costs to goods listed in Annex I of Regulation (EU) 2023/956 from 1 January 2026, and free allocation is withdrawn in the sectors it covers as it phases in [80]; the two together equalise the carbon-cost burden between EU and non-EU producers rather than removing it from either.

Which side of Annex I a product sits on therefore decides its allocation as well as its import charge, because Article 10a(1a) of the amended ETS Directive [72] applies the phase-out only to goods on that list. Adipic acid (CN 2917 12 00) is not on it, so it keeps its allocation under Article 10b(1) at 100 % of the benchmark “for the period until 2030” while imported adipic acid bears no EU carbon cost at all, and the free allocation is what compensates an EU producer for that asymmetry.

Neither position is settled across a 2030 plant's producing life, and the two ways it could end are not equivalent. If adipic acid were brought inside CBAM the allocation would be withdrawn and the import charged in the same act, leaving relative positions broadly intact; if Article 10b(1) lapses after 2030 without that extension, the allocation goes and the import charge does not. Inclusion is discussed but not legislated: the Commission's December 2025 review [81] flags adipic acid, the proposal published alongside it extends the mechanism only to steel- and aluminium-intensive downstream goods [82], and the July 2026 ETS review proposal [83] would extend carbon-leakage protection to 2040 instead. The study therefore does not forecast which way it goes: it bounds free allocation between full retention and complete withdrawal and reports what the difference between them is worth (§3.3.2).

By contrast, HNO_3 (CN 2808 00 00) is in CBAM scope under the fertilisers category, so the HNO_3 supplier-choice exposure (EU-abated versus non-EU unabated) is priced directly by CBAM independently of AA's treatment. Benzene (CN 2902 20 00) is outside Annex I like AA, so no border charge attaches to it and its price enters as a market price throughout.

2.4.2. European Industrial Context and Constraints

Why ARRRRA defines the assessment scope

ARRRA hosts approximately 40 % of EU petrochemical output and the full value chain from feedstock to polymer production [8], making it the natural reference cluster for a European AA competitiveness assessment. Recent plant closures within the cluster, notably BASF's ~115 kt/a adipic acid plant at Ludwigshafen (the abated unit of §2.2.1), illustrate that competitive pressure from lower-cost regions has direct consequences for the production system studied here. BASF closed it in 2025 citing weak European demand and higher energy costs than competing regions [8, 69, 84], not a carbon-cost penalty on its own well-abated emissions, which under the EU ETS is a net allowance surplus. This is the distinction the thesis turns on: European adipic acid faces cost pressure even where the plant itself is clean and carbon-advantaged.

Structuring the assessment around three ARRRRA countries rather than a single location serves two purposes. First, it grounds the analysis in a realistic range of input data: EHV effective electricity costs and grid carbon intensity vary considerably across Belgium, Germany, and the Netherlands (Tables A.1 and A.2); using observed country values rather than a generic assumption captures the actual variation a large-scale AA producer faces, directly addressing the location-specificity gap in prior TEA work. Second, treating each country as a scenario makes it possible to assess both the sensitivity of route competitiveness to these variables and the potential of each ARRRRA location for large-scale industrial AA production under each route.

Grid access as a feasibility constraint

A large electrochemical AA plant requires a firm high-voltage grid connection with substantial contracted capacity, and consumer-connection congestion across the ARRRRA region is most acute in the Netherlands, where the Dutch high-voltage grid operator TenneT has closed several provinces to new large consumers and firm contracts are offered only under non-firm Flexible Connection Agreements [85, 86], and Belgium, where the Belgian grid operators Elia and Fluvius have rolled out flexible-access contracts for industrial users in response to a six-fold rise in connection requests [87]; Germany's Rhine-Ruhr area is less constrained at the consumer-connection level, with the binding system-level strain on the north-south transmission corridor instead [88]. For a continuous electrochemical process, non-firm or curtailment-subject contracts reduce effective capacity factor and raise LCOA, so grid access functions as a feasibility constraint rather than just a cost input.

2.4.3. Country-Coupled Electricity Inputs

Renewable procurement and intermittency

Flexible operation of electrified chemical plants under variable renewable electricity is an active research theme [89, 90], and the question pursued here is whether running the flexible electrochemical load at a load factor below unity, following the hourly price, lowers the levelised cost of adipic acid or the GHG intensity. The result therefore turns on the hourly shape of price and carbon intensity rather than on their level, so a fixed-price supply contract would not change it. Supplying such a plant from a dedicated on-site park is not physically precluded, but at roughly 8,000 full-load hours it would need generation oversized to a large multiple of demand and storage to bridge the low-output periods [8, 66]. Luo et al. [91] at Delft model exactly that for an electrified plant in an ARRRRA setting, tying it to an oversized wind and solar park with battery backup that shuts down whenever local generation cannot

sustain its 70 % minimum throughput; their grid-connected reference case, at the wholesale day-ahead price, is cheaper than every dedicated-park case they evaluate. Kashanian et al. [92] do find on-site supply viable for the identical *ccMA* → *t3HDA* pathway, but under US conditions of small distributed facilities, investment tax credits and explicit carbon pricing, which are not those facing a centralised European plant.

Two further things make a dedicated park hard to carry into this work. A larger park buys availability and a higher achievable load factor but commits more capital, and at high oversizing a falling share of the generation reaches the plant at all, so costing it turns on how that capital is split between the chemical plant and the power it cannot absorb, which the process model does not settle. Fixing the plant to one park also removes the freedom to site it where the tariff, the grid carbon intensity or the feedstock logistics are best, which is one of the levers this work sets out to test. Grid coupling is therefore the supply assumption here. For a grid-connected plant, intermittency is not a physical supply constraint but a market signal: the hourly variation in the wholesale price [93] and in the grid's carbon intensity [94], both of which already embody the system-wide penetration of renewables across the European power system [95]. The plant is never forced off by local weather; the operating lever is instead the load factor of the flexible electrochemical fraction, concentrated in the cheapest and cleanest hours, in line with the demand-response and process-flexibility literature [89, 90]. Load-following also means repeatedly starting and stopping the cell, which could shorten its life. Anion-exchange-membrane electrolysis suggests the degradation contribution is minor: over more than 1,500 cold starts of a stack, with a controlled voltage ramp, degradation per start was small enough that the authors conclude AEM electrolysis is highly compatible with regular operational interruptions [96]. The dispatch analysis prices cycling on that measured rate, so it carries the degradation cost of start–stop but not the start-up energy demand the same study quantifies.

How the plant is operated under a variable price is itself a modelling choice. An electrified process can either hold its output fixed by oversizing the flexible unit and buffering an intermediate so the continuous downstream train runs steadily, or flex the whole plant and produce less when power is expensive; the two differ in whether a storage buffer or a lower output absorbs the swing, and which is preferable turns on how far the downstream can turn down. For a stiff downstream the buffer grows large: in cost-optimal power-to-methanol the whole plant is flexed and hydrogen storage kept to a few hours [97], and for a rigid ammonia loop the storage needed to hold synthesis steady approaches 40 % of total capital [98]. Because adipic-acid crystallisation is likewise stiff, this thesis reports both bounds (oversize-and-buffer and produce-less) rather than committing to one; the pricing data and dispatch method are set out in the methodology, and the result with the route comparison.

Electricity prices in ARRA

Electricity cost inputs differ by route and by country. Eurostat's statutory prices for the largest non-household consumption band put the 2024 annual average at €130/MWh in Germany, €110/MWh in the Netherlands and €84/MWh in Belgium [99]. An extra-high-voltage baseload plant of the size modelled here does not pay them: after each country's relief instruments the effective 2024 cost is €46/MWh in Germany, €56/MWh in Belgium and €95/MWh in the Netherlands [100]. The relief is large enough to reorder the three countries, moving Germany from the most expensive on statutory prices to the cheapest, while the Netherlands, which ended its indirect cost compensation scheme in 2023, becomes the most expensive. Carried forward on that basis, the E-Bridge 2030 projections set the three electricity-price tiers used here: EP-Low €49/MWh (Germany), EP-Central €62/MWh (Belgium) and EP-High €92/MWh (Netherlands). Because that relief is national policy rather than a market outcome, and each country's scheme can be withdrawn as the Dutch one was in 2023, a supplementary EP-Statutory tier prices all three countries at the unrelieved Eurostat level, testing what the comparison looks like if the relief goes. All routes are priced on that relieved basis, the conventional route included: although it draws far less power, as an energy-intensive chemicals consumer it qualifies for the same instruments. The tariff components, the relief mechanisms behind them, the 2030 projection and the reasoning behind each value chosen here are set out in §A.1, Table A.1.

Grid carbon intensity

Scope 2 GHG emissions of the electrified routes are driven mainly by the carbon intensity of electricity, which reflects fuel mix and renewable deployment rather than market tariffs and therefore follows a different path from price. On a 2024 lifecycle, consumption-based basis [94] the three ARRA countries are more than a factor of two apart: Belgium 150 gCO₂e/kWh (nuclear and offshore wind, with residual gas), the Netherlands 261 gCO₂e/kWh (gas-dominated dispatch) and Germany 341 gCO₂e/kWh (residual coal after the nuclear exit), against an EU-27 average of 213 gCO₂e/kWh [95]. Those three values are the *GI-Low*, *GI-Central* and *GI-High* tiers. A fourth, supplementary tier, *GI-Renewable*, bounds the analysis from below with France's observed 2024 value of

31 gCO₂e/kWh, grounding the low-carbon case in a real European grid rather than an assumed value. It carries two readings: as a location it is France, and as a supply it is close to what any near-zero-carbon electricity would deliver, so it also measures what coupling the plant to a renewable source would be worth. In the location analysis France brings its own electricity price of €48/MWh with it, so that case is a decarbonised supply on both counts rather than a clean grid at an assumed cost. The three ARRRA tiers all sit well above the 50 gCO₂e/kWh assumed by Yeo (2025) [9], which understates scope-2 emissions wherever in the cluster the plant sits. Because the plant operates for twenty years, a supplementary sensitivity also treats grid intensity as time-dependent rather than frozen at 2024. The country values and their scenario mapping, the projection source and the one country whose projected decline is contested are given in §A.2, Table A.2.

2.4.4. Biomass Route Product Boundary: *t*3HDA versus Adipic Acid

The *t*3HDA-versus-AA product boundary represents an unresolved gap in the biomass-route TEA literature: existing studies evaluate the two endpoints separately rather than as alternatives within a consistent process and cost framework. Matthiesen et al. establish electrochemical selectivity across MA isomers, *t*3HDA, and AA, but do not quantify the cost difference between endpoints at equivalent scale [56, 50]. Prior bio-route TEAs that include AA as a product find that thermocatalytic hydrogenation of purified MA provides a lower MSP than the ECH pathway in the absence of a *t*3HDA market premium [101, 9]; the most recent integrated TEA reports an MSP of approximately \$2.60 kg⁻¹ for the thermocatalytic route under US conditions [51], while *t*3HDA itself has no commodity market reference, with a production cost floor of approximately \$2.00 kg⁻¹ as the only available benchmark [50]. The magnitude and commercial realisability of a *t*3HDA price premium therefore remain the key determinant of which product boundary is economically preferable.

2.4.5. Electrochemical Cell Performance and Degradation

As noted in Section 2.2.2, progressive performance decay raises electricity consumption and eventually necessitates component replacement. Yet most TEA studies assume constant cell performance or represent degradation only through simplified maintenance cost factors, which may underestimate both lifetime electricity demand and the capital reinvestment associated with periodic stack or component replacement.

The durability evidence for both cell types is thin. No published test of either architecture runs beyond about 100 h, against the thousands to tens of thousands of hours quoted for mature electrochemical stacks, and the two tests that exist show measurable decay within that window: Faradaic efficiency falling from 93 to 81 % over 60 h on the ECO cell architecture, and cathodic corrosion with lead leaching on the ECH cell [39, 57, 45]. Degradation must therefore be bounded by analogy, and the two cells need different analogues because they differ in architecture and in dominant mechanism: alkaline and AEM water electrolysis together with chlor-alkali membrane replacement practice for the ECO cell, and industrial electroorganic synthesis on lead cathodes for the ECH cell. Both are an optimistic lower bound, because each cell faces a route-specific mechanism its analogue does not, and the severity of those mechanisms at industrial operating hours is unknown. That is what motivates treating initial cell performance and degradation parametrically rather than fixing either. The durability data and the two analogues are reviewed in §A.7.

Initial-performance tiers

The initial-performance tiers are set out in Table 2.4. No industrial-scale data exist for either cell, so the tiers are a parametric sensitivity spanning each route's reported performance envelope rather than a known operating point. Because Faradaic efficiency falls as current density rises (the parasitic hydrogen-evolution reaction competes at higher current), each route is pinned to its industrially relevant current density and only FE is varied. Specific energy consumption scales with cell voltage divided by Faradaic efficiency, so a rise in cell voltage acts in the same way as the same proportional fall in Faradaic efficiency; the tiers therefore bracket the cell's electricity demand without carrying voltage as a separate dimension. Current density sets the electrode area needed for a given output, so it acts on capital rather than on energy and is treated in the capital cost instead (§2.2.2, Section 3.2.5).

Each tier is an efficiency reported for a real cell rather than an assumed value, and §A.7 gives the catalyst, the current density and the conditions behind each one, including where a tier stands for a target rather than a demonstrated operating point. One feature of the ECH tiers carries into the results and is worth stating here: the charge not reaching *t*3HDA leaves as hydrogen, so Faradaic efficiency and the cathodic hydrogen co-product are inversely coupled and cannot be varied independently. The hydrogen credit therefore shrinks as efficiency rises, and at the

Table 2.4: Electrochemical cell-performance tiers. Current density and cell voltage are pinned on every tier, so the sweep is a pure energy bracket: specific energy consumption scales with $V_{\text{cell}}/\text{FE}$. The catalyst, current density and conditions behind each value are given in §A.7.

Route	Tier	FE	j (mA cm^{-2})	V_{cell} (V)
ECO	Optimistic (target)	0.90	300	1.9*
ECO	Central (demonstrated)	0.80	300	1.9*
ECO	Conservative	0.45	300	1.9*
ECH	Optimistic	0.60	200	5.7
ECH	Central (anchor)	0.40	200	5.7
ECH	Conservative (judgement)	0.30	200	5.7

*ECO V_{cell} is a derived full-cell estimate; the source papers report anode potentials of 1.5–1.9 V vs RHE, not full-cell voltage.

optimistic tier the biomass route to adipic acid can turn from hydrogen-surplus to hydrogen-short. Degradation rate and replacement interval are not measured values but brackets transferred from the analogues, so they sit with the modelling choices in Section 3.2.6.

2.4.6. Feedstock Treatment and Input Values

Feedstock costs dominate operating expenditure. The CON and ECO routes depend on benzene-derived intermediates, coupling them to crude-oil and naphtha volatility [10, 6, 2]; the biomass route is exposed instead to MA broth, decoupled from fossil markets but with its own supply and cost uncertainty [102, 101]. Hydrogen is purchased by the conventional route, while on the electrified routes part or all of the supply is made in the electrochemical cell itself.

Each feedstock's price and scope-3 CI is treated either as a discrete procurement scenario (a supplier or pathway choice) or as a continuous sensitivity range (a low/central/high sweep), and the two can co-exist on one feedstock: HNO_3 uses procurement scenarios for CI alongside a sensitivity-range price, and benzene's two procurement scenarios each carry an internal price tier. H_2 , HNO_3 and benzene are evaluated as procurement scenarios, MA broth as a sensitivity range under its single glucose-derived pathway. Three process chemicals are carried alongside them: potassium hydroxide for the ECO cell's electrolyte, and sulphuric acid and calcium hydroxide to crystallise the product and neutralise the wastewater on the biomass route; these three are carried as sensitivity ranges on price, their intensities held.

Feedstock prices are recent observed European levels, each central cross-checked against independent indices to confirm it is not shock-distorted, but the vintage differs by feedstock and each is stated with its value in §A.5. Nitric acid, potassium hydroxide and sulphuric acid are 2026 levels from the BusinessAnalytics index [103], and calcium hydroxide a 2025 one. Benzene is carried on its 2024–2025 average instead, because the 2026 Strait-of-Hormuz crisis inflates the spot, which sits in the upper band rather than at the central. Hydrogen is a 2024 anchor. MA broth, having no index at all, is derived from Mokwatlo (2024) [51]: that study's minimum selling price for muconic acid, less the cost of crystallising it out of the broth, consistent with the approach of Yeo (2025) [9]. How wide each band is set depends on what moves that price rather than on a common tolerance: fossil benzene spans its own observed range, from the 2015–2019 trough to the April 2022 peak; bio-benzene is triangulated from the proxy prices reported for it; nitric acid runs asymmetrically from the oversupplied global floor to the 2022 gas-crisis peak; hydrogen is anchored per production type and bracketed by the IEA projection below and Dutch project costs above; the MA broth carries a flat $\pm 30\%$ for want of any market at all; and the three process chemicals carry the price ranges reported for them. Each band and the reasoning behind it is given in §A.5. The framework mechanics are set out in the methodology.

Bio-attributed benzene needs reading with more care than a price band conveys. It is modelled here on the catalytic fast-pyrolysis route, at TRL 6 the highest of the nine renewable-benzene routes Miller et al. assess [104], so both its price and its intensity are that route's values rather than a generic bio-benzene figure. There is no public contract price for it, so the band is triangulated from proxy prices reported in the literature rather than read from an index, and its central \$1.77/kg is a reported minimum premium over fossil rather than a mid-market quote; the price is therefore uncertain rather than established. Its $1.08 \text{ kgCO}_2\text{e/kg}$ intensity is derived by applying one reported reduction, measured on a single feedstock, to the fossil benzene value, and is therefore uncertain. The other main uncertainty is supply, which at present scale would not be sufficient to feed a 100 kt/a adipic-acid plant. The derivation and the full notes on interpretation are in §A.5.1.

Table 2.5: Feedstock treatment and input values, on 2024–2026 observed European market levels. Every value is derived in §A.5, one subsection per feedstock in the order of this table, with its sources and the reasoning behind each band.

Feedstock	Scenario / tier	Price (\$/kg) low / central / high	Scope-3 CI (kgCO ₂ e/kg) per tier, or low / central / high
Benzene procurement scenarios (2)	<i>BEN-Fossil</i> naphtha cracker	0.39 / 0.88 / 1.65	1.87
	<i>BEN-Bio</i> ¹ catalytic fast pyrolysis	0.77 / 1.77 / 2.31	1.08
	<i>Grey SMR</i>	3.20	10
H₂ ² procurement scenarios (3)	<i>Blue SMR + CCS</i>	3.50	3
	<i>Green electrolysis</i>	6.65	1
HNO₃ price range; EI scenarios (2)	<i>EI-A</i> EU abated	0.22 / 0.30 / 0.50	0.93 (EI-A)
	<i>EI-B</i> non-EU unabated	one band, both tiers	3.99 (EI-B) per kg 100 % HNO ₃
MA broth sensitivity range	single glucose-derived path- way; no public spot market	1.30 / 1.86 / 2.42 ±30 %	0.25 / 0.38 / 0.55 per kg MA content, net of biogenic carbon
Process chemicals price sensitivity ranges	<i>KOH</i>	0.62 / 0.69 / 0.83	1.48
	<i>H₂SO₄</i>	0.13 / 0.19 / 0.35	0.14
	<i>Ca(OH)₂</i>	0.12 / 0.14 / 0.20	0.94

¹ Bio-attributed benzene has no public contract price and, at present, no supply at this scale; its band and the notes on how it should be read are in §A.5.1.

² Hydrogen values are 2024 NW-European production costs rather than transaction prices, so all three tiers are lower bounds on what a plant would pay; the brackets either side and the notes on how they should be read are in §A.5.2.

2.4.7. Scenario Framework Summary

Table 2.6 consolidates the literature-supported ranges for the five tier dimensions swept low/central/high in this thesis (carbon price, electricity price, grid carbon intensity (CI), feedstock-cost joint sensitivity, electrochemical cell performance), alongside the sixth dimension set out above, the discrete procurement choice for hydrogen, HNO₃ and benzene (Table 2.5, §2.4.6).

Table 2.6: Literature-supported competitiveness dimensions and indicative ranges for comparative TEA of adipic acid production routes. Electrochemical cell performance applies to electrochemical and biomass routes only. The columns are ordered by the direction each tier moves the result rather than by the value of the input itself, which is why the optimistic cell tier sits under *Low*. Every tier is derived in Appendix A, one section per dimension, with its sources and the reasoning behind each range.

Dimension	Low	Central	High
Carbon price N ₂ O and CBAM exposure	<i>CP-Low</i> €75/tCO ₂ , 2024–2026 spot mid-point	<i>CP-Central</i> €115/tCO ₂ , IEA NZE 2030	<i>CP-High</i> €145/tCO ₂ , upper cluster of 2030 estimates
Electricity price ¹ primary OPEX of the electrified routes	<i>EP-Low</i> €49/MWh, Germany	<i>EP-Central</i> €62/MWh, Belgium	<i>EP-High</i> €92/MWh, Netherlands
Grid CO₂ intensity ² scope-2 GHG of the electrified routes	<i>GI-Low</i> 150 gCO ₂ e/kWh, Belgium	<i>GI-Central</i> 261 gCO ₂ e/kWh, Netherlands	<i>GI-High</i> 341 gCO ₂ e/kWh, Germany
Feedstock cost every feedstock swept together (Table 2.5)	<i>FP-Low</i>	<i>FP-Central</i>	<i>FP-High</i>
Electrochemical cell performance <i>j</i> , FE, <i>V</i> _{cell} and degradation (Table 2.4)	<i>EC-Optimistic</i> FE: ECO 0.90, ECH 0.60	<i>EC-Central</i> FE: ECO 0.80, ECH 0.40	<i>EC-Conservative</i> FE: ECO 0.45, ECH 0.30

¹ E-Bridge (2024) 2030 all-in EHV projections [100]; all routes are priced on this EHV basis. The supplementary *EP-Statutory* tier prices them at the unrelieved Eurostat band IG level instead (§A.1).

² A supplementary *GI-Renewable* tier takes France's 2024 intensity of 31 gCO₂e/kWh as a decarbonised-grid bound (Table A.2); in the location analysis it carries France's electricity price with it.

Beyond the five tier dimensions and procurement choice, supplementary checks probe additional inputs' impact without being folded into the main sweep: EU ETS free allocation, bounded between full carbon-leakage protection and none; country-specific grid decarbonisation pace; the rising post-2030 CO₂-price trajectory; the Class 4 accuracy band on the capital estimate; and the cost of the electrochemical cells themselves, swept between a pilot-scale and a matured estimate.

Together, the scenario framework and these supplementary variables map each route's levelised cost and cradle-to-gate greenhouse-gas intensity under realistic European market and policy scenarios, identifying the dominant bottlenecks and viability conditions for each route.

Methodology

This chapter describes the methodology for comparing the three adipic acid production routes (conventional HNO_3 oxidation, CON; electrochemical oxidation, ECO; electrochemical hydrogenation, BIO) under consistent techno-economic and environmental assumptions. A Python-based evaluation framework reads mass and energy balances from the Aspen Plus process models and computes LCOA and GHG intensity across the full evaluation space: route, country, commissioning year, and the input dimensions identified in the literature review.

3.1. Scope and Assessment Framework

The three routes are compared under a common basis of design (Table 3.1), consistent with Yeo (2025) [9] and standard TEA convention [64]. All routes are evaluated over the same 20 production years and at the same 8 % real discount rate, so that differences in the levelised cost arise from the routes' own capital, feedstock and energy requirements rather than from route-specific financing or amortisation assumptions. Plants are sited in three ARRA countries (BE, NL, DE); France is included as a reference case in the grid-intensity analysis because its near-nuclear baseline represents a structurally different electricity system. The headline plant is commissioned in 2030, and the first production year is then shifted over 2028 to 2035 as a sensitivity. That window spans the EU ETS Phase 4 (2021–2030) to Phase 5 (2031–) transition, so it tests how far the results depend on the free allocation, whose statutory basis runs only to 2030, and on the national grid-intensity decline over the producing life. Product purity is fixed at ≥ 99.5 wt%, the minimum specification for polymer-grade adipic acid in nylon-6,6 production [26, 22], consistent with Yeo (2025). This ensures KPI differences reflect route chemistry rather than downstream finishing choices. For the biomass route, an alternative endpoint at *t*3HDA is evaluated in parallel.

Table 3.1: Common basis of design.

Parameter	Value
Production capacity	100 kt/a
Product purity (AA or <i>t</i> 3HDA)	≥ 99.5 wt%
Annual operating hours	8 000 h/yr
Plant lifetime	20 production years
Commissioning (first production)	2030 headline; 2028–2035 sensitivity
Real discount rate	8 %
Locations	Belgium, Netherlands, Germany
Cost basis	2024 US dollars (CEPCI 800)
System boundary	Feedstock conditioning through product gate
Functional unit	1 kg purified AA (or <i>t</i> 3HDA)

The 100 kt/a capacity is the nominal design basis on which the flowsheets were sized; it is not imposed on the converged models, which settle at 97.5, 101.9, 98.8 and 94.1 kt/a on the conventional, electrochemical and two biomass routes respectively (Appendix G). Every per-kilogram figure reported in this work is taken on each route's own converged rate rather than on a common 100 kt/a, so the routes are compared per kilogram of product and not per plant.

The methodology has four parts. Section 3.2 describes how the process models are built, including the biomass-route extension to adipic acid. Sections 3.3 and 3.4 define the headline base case, the first giving the cost engine and the levelised cost of adipic acid and the second the cradle-to-gate GHG intensity, both at central input values with cell degradation applied. Section 3.5 then sets out the additional parameters evaluated on top of that base case, the procurement scenarios, lifetime input trajectories, coherent-future narratives, and the load-following sensitivity, keeping the headline result distinct from the parameters swept around it.

3.2. Process Modelling Approach

The process simulations build on the Aspen Plus models of Yeo (2025) [9]. This thesis retains those flowsheets and extends the process models with: (i) heat-integration refinements; and (ii) a catalytic hydrogenation reactor (*t*3HDA $\rightarrow$ AA) for the biomass route. All other extensions are implemented in the Python evaluation layer (Section 3.5) rather than in the Aspen models, including the packaged boiler that raises steam from the biomass routes' surplus cathodic hydrogen (§3.3.2). The converged flowsheets are reproduced in Appendix F and their stream tables in Appendix G.

3.2.1. Aspen Plus configuration

Simulations use Aspen Plus V14, inheriting the property-method framework and all component data of Yeo (2025) [9]. The conventional and ECO routes use **NRTL-RK** (NRTL activity coefficients for the liquid phase, Redlich–Kwong equation of state for the vapour phase, Henry's law for dissolved gases), suitable for the polar non-ideal aqueous–organic mixtures at the moderate pressures in these flowsheets. The biomass route uses **ELECNRTL** because the fermentation broth and electrochemical hydrogenation environment contain dissolved ionic species (NH_4^+ , SO_4^{2-} , organic-acid anions including the muconate dianion) whose activity coefficients and speciation must be tracked explicitly; a charge-balance block enforces electroneutrality. All binary interaction parameters are inherited from Yeo (2025). For *trans*-3-hexenedioic acid (*t*3HDA), published thermodynamic data are limited; properties are estimated using the Joback group-contribution method, with enthalpy of solution approximated from adipic acid in the absence of a direct measurement.

3.2.2. Reactor and flowsheet modelling

The CON and ECO flowsheets are inherited from Yeo (2025) [9] without structural modification; the ECO cell replaces the HNO_3 reactor, with downstream separation (electrodialysis, crystallisation, electrolyte recovery) also following Yeo (2025). The BIO route adds one structural extension: a catalytic hydrogenation reactor (*t*3HDA $\rightarrow$ AA) that enables dual-endpoint evaluation (AA alongside *t*3HDA), operating under Capelli conditions (Pd/AC, 70 °C, 1 bar H_2) [59].

Each reactor is modelled at the level of detail needed to close mass and energy balance while keeping performance parameters accessible to the scenario framework. Conventional catalytic reactors (benzene hydrogenation, KA-oil air oxidation, HNO_3 oxidation) use the reactor types and yield specifications of Yeo (2025). The electrochemical units (ECO cells for cyclohexanone oxidation, ECH cells for muconic acid hydrogenation) and the catalytic hydrogenation reactor are modelled as **RStoic blocks**: single-pass yield, Faradaic efficiency, and cell voltage are imposed as parameters, and the electrical duty for the electrochemical units is calculated from Faraday's law (Section 2.2.2). This decouples reactor performance from the flowsheet, so parameter ranges and their time evolution can be varied in the scenario framework without changing the Aspen models.

Fermentation is outside the system boundary; muconic-acid broth enters as a purchased feedstock [51]. The ECH cell operates directly in the broth without prior muconic-acid purification [57]; residual impurities (sugars, proteins, salts) affect electrode degradation (Section 2.4.5) but do not require separate modelling because the mass balance is imposed via RStoic conversion specifications.

3.2.3. Biomass route extension: *t*3HDA to adipic acid

Yeo (2025) ends the biomass route at *trans*-3-hexenedioic acid (*t*3HDA), so that route yields *t*3HDA rather than adipic acid and no flowsheet exists for the conversion. To evaluate adipic acid on the same basis as the other routes, this thesis models the conversion explicitly as a catalytic hydrogenation step followed by a dedicated product-recovery train, shown in Figure 3.1.

The purified *l*3HDA leaves the upstream recovery train as a solid, is redissolved in deionised water and heated to reaction temperature. That purification is retained rather than stripped out, because catalytic hydrogenation over a precious-metal catalyst is poisoned by the sulphur compounds and salts in the fermentation broth: the polishing step cannot run in the raw broth the way the upstream electrochemical *cis,cis*-muconate step does. The point matters in the model as much as in the plant, since an RStoic reactor at fixed conversion converts whatever it is given, so a stripped-down flowsheet would converge and look sound while hiding a catalyst-poisoning problem.

Hydrogen is drawn from the plant header and split at H2-SPLIT into a reactor feed of 90.20 kmol/h and a sale stream of 45.50 kmol/h, the reactor feed being compressed at CMPR4 on the way in. That compressor's electricity is tagged to the route total (§3.3.3).

The two streams meet in R-POLISH, where *l*3HDA is hydrogenated to adipic acid over a Pd/AC catalyst under Capelli conditions (5 % Pd/AC, 70 °C, 1 bar H₂) [59]. It is modelled as an RStoic block at fixed conversion (§3.2.2).

The polishing step is the 1:1 hydrogenation *l*3HDA + H₂ → AA. In the model it is specified instead as the relabelling step T3A(S) → AA; because both sides are then the same component, the heat of reaction cannot be derived from them and is entered directly as $\Delta H_{\text{rxn}} = -129 \text{ kJ/mol}$ against T3A(S) [59]. The relabelling itself is forced by a property-data constraint inherited with the flowsheet: the *l*3HDA and adipic-acid pseudo-components (AA, AA(S), T3A(S), T3A(AQ), T3A⁻, T3A²⁻) all point at the same ADIPIC-ACID databank entry and therefore carry the same molar mass, $M = 146.14 \text{ g/mol}$. Writing the reaction with hydrogen would therefore leave mass in and mass out unequal, short by the 2.016 g/mol the hydrogen carries, so it is written as a relabelling, which closes exactly.

Writing the reaction this way leaves the hydrogen unconsumed, so it leaves in the reactor outlet and would carry through the crystallisers into the mother liquor as dissolved gas. A flash immediately downstream (FAKE-H2) removes it as H2-RXN and passes the liquid on to the recovery train. Nothing stands at this point in the real process: the reacting fraction is taken up inside the reactor itself, and a genuine excess would be recovered or recycled rather than vented.¹ FAKE-H2 is therefore treated as a modelling artefact alongside the charge-balance and property-correction blocks: its duty is excluded from the heat-integration analysis on the same basis as the conventional route's bridge blocks (§3.2.4), and it carries no equipment cost. Hydrogen consumption for the cost and emission accounting is likewise computed stoichiometrically, one mole H₂ per mole adipic acid, and enforced at the H2-SPLIT ratio rather than read from an Aspen stream.

A further 0.5 % of the fed *l*3HDA arrives as the dissolved, deprotonated form (T3A⁻/T3A²⁻), which a reaction written on the neutral species leaves unconverted. These ions leave with the mother liquor rather than co-crystallising, so they lower the conversion accounting by about half a per cent without contaminating the product.

Adipic acid is then recovered by two-stage cooling crystallisation rather than evaporation, chosen because it runs on electricity and so suits the biomass route's low-steam character. The hydrogen-free liquid is chilled from 70 to 17 °C and crystallised, then cooled further to 2 °C and crystallised again, both stages sharing one centrifuge (AA-CEN) that separates the adipic-acid solid (AA-SOLID) from the remaining solution, the mother liquor (ML-AA). The sub-ambient duty is met by a dedicated chiller and costed as refrigeration electricity, kept separate from the process cooling-water and chilled-water services (§3.2.4). The recovered solid is the adipic-acid endpoint's product; the route is also reported at the *l*3HDA endpoint, so the two products are compared on the same flowsheet.

3.2.4. Heat integration

Each route's flowsheet is heat-integrated to identify utility-saving opportunities and quantify the resulting hot and cold utility demands, whose effect on the levelised cost of adipic acid is reported in the results (§4.2.2). The networks are designed by the pinch design method of Linnhoff & Hindmarsh [105], performed using Aspen Energy Analyzer V14 with a minimum approach temperature (ΔT_{min}) of 10 °C. Towler & Sinnott [64] (Ch. 3) note that practical ΔT_{min} values for heat exchangers lie between 5 and 30 °C, with the economic optimum typically flat over the 10–30 °C range; 10 °C is chosen here as the lower bound of that optimum.

Cooling water is specified at 25 → 30 °C, the standard cooling-tower-design value applicable across the ARRA petrochemical cluster [64] (Ch. 3). This is more conservative than Yeo's (2025, §3.2.3) [9] Rotterdam-specific

¹H2-RXN is accordingly not a single physical quantity. H2-SPLIT delivers 90.196 kmol/h of hydrogen to a reactor converting 84.084 kmol/h of *l*3HDA, so 84.084 kmol/h of what the flash removes is hydrogen the real reaction would consume and only 6.112 kmol/h is a genuine excess; a further 0.050 kmol/h stays dissolved and leaves with the mother liquor, leaving 90.146 kmol/h in H2-RXN.




Figure 3.1: Biomass-route extension from *t*3HDA to adipic acid, built in Aspen Plus V14 on the inherited Yeo (2025) biomass model. Blue and red mark the streams that cross the flowsheet boundary, inlets and outlets respectively. FAKE-H2 is a modelling device rather than a physical unit; the flowsheet and that simplification are described in the text.

10 → 15 °C assumption based on direct North-Sea cooling: under the 15 °C-warmer inlet used here, the lowest temperature any hot process stream can be cooled to with plain CW is 35 °C (with $\Delta T_{\min} = 10$ °C), so any stream that needs to reach below this temperature must be served by chilled water at 5 → 10 °C, which is substantially more electricity-intensive to deliver than plain CW. Forcing more cold-utility duty onto chilled water therefore lowers the calculated heat-integration savings, but the reported number is then achievable at any ARRRR site rather than only at Rotterdam-coast plants with direct seawater intake. Hot utilities (LP steam at 124 °C, MP steam at 174 °C, with steam-generation credits) follow Yeo (2025).

This conservatism is not neutral between routes, and its main effect is on carbon rather than on cost. Refrigeration is delivered as electricity and therefore carries the grid intensity, so specifying the warmer cooling water converts part of what would have been a scope-2 saving into a scope-2 penalty. The penalty scales with a route's sub-ambient duty, whereas the saving it offsets scales with how little heat recovery the inherited flowsheet already had; the two are unrelated, so the same absolute penalty is negligible against a large steam saving and decisive against a small one. It falls hardest on the electrified routes, which combine little remaining recovery with substantial sub-ambient duty, and lightest on the conventional route. The consequences per route are quantified with the results (§4.2.2, Table 4.6), where the choice is also re-examined against Yeo's colder basis so its contribution can be separated from that of heat integration itself.

Cooling water is priced as a utility at \$0.003 per kWh of duty following Yeo (2025) [9], a charge dominated by makeup water and treatment chemicals. Its circulation-pump and cooling-tower-fan power is charged separately as scope-2 electricity, at 0.034 kWh_e per kWh_{th} on the country grid intensity, derived from the 1–2 kWh per 1000 gal circulated given by Towler & Sinnott [64] (Ch. 3); the quoted range spans 0.023–0.045.²

Chilled water is delivered below the pinch, where heat must be lifted from below ambient by a compressor rather than rejected to ambient by pumps and fans, so its electricity intensity is about six times the cooling-water figure above. Delivered by a vapour-compression chiller whose compressor work dominates, it is not priced as a \$/GJ utility: its duty is converted to refrigeration electricity at a coefficient of performance of 5.0, equivalently 0.70 kW per ton of refrigeration for a water-cooled centrifugal machine, and charged on the country electricity supply, carrying both cost and grid carbon there. ASHRAE 90.1-2019 Table 6.8.1-3 [106] sets a full-load minimum of 6.29 for this size class, which falls to about 5.8 once corrected to the temperatures used here; 5.0 is taken as a conservative operating value against what is a rating for new equipment at its design point.

Stream properties

Two stream properties are set rather than left to the tool. Exchanger area, and through it exchanger capital, follows from the overall heat transfer coefficient of each match, and Aspen Energy Analyzer applies one stock film coefficient to every process stream, the gas-service value of Towler & Sinnott [64], which is an order of magnitude too low for the aqueous service that dominates all three flowsheets and would oversize almost every exchanger. Each stream is therefore assigned its own coefficient by fluid category. This changes area and so capital, but not the utility targets, which are set by duty alone. Second, the crystallisers exchange latent heat, which the inherited property model does not carry: what the blocks report instead is a liquor-specific heat-of-mixing term that disagrees between routes. Each crystalliser is therefore entered as its own isothermal stream carrying the enthalpy of fusion of adipic acid, $\Delta H_{\text{fus}} = 34.85$ kJ/mol [107], on the net solid formed across the block. The category values,

²At their 10 °C cooling-range worked example, 1 kW of cooling duty circulates 181.2 thousand gal per year over 8000 h, so 1.5 kWh per 1000 gal gives 271.8 kWh_e against 8000 kWh_{th}, or 0.034.

the per-match consistency check, the crystallisation duty definition and the rules by which the extracted stream list is curated are set out in §D.1.

Yeo (2025, §6.1.6) [9] integrated his flowsheets only partially, raising medium-pressure steam from the two reactor exotherms and placing three process-process exchangers by inspection, and left systematic pinch analysis out of scope. The analysis here is a retrofit rather than a grassroots design: it starts from that existing network and adds to it, so what is reported is the saving a systematic design recovers beyond what he already had. It is re-run wherever a route-specific extension adds streams, notably the BIO polishing reactor.

The resulting stream list for each route is given in §D.3, and the utility demands and crystallisation duties it produces are reported with the results (§4.2.2).

3.2.5. Electrochemical cell model

The electrochemical cells set the specific electricity demand per unit product, which feeds both the energy cost in the levelised cost of adipic acid and the scope-2 GHG intensity. From Faraday's law,

$$w_{\text{elec}} = \frac{V_{\text{cell}} n F}{FE \cdot M \cdot 3600} \quad [\text{kWh kg}^{-1}], \quad (3.1)$$

with cell voltage V_{cell} , electrons transferred n , Faraday constant F , Faradaic efficiency FE and product molar mass M . Each route is pinned to its own operating point, so V_{cell} and j are held fixed and only FE varies across the tiers of Table 2.4.

Each tier sets two things, because one current carries both: the energy above, and the hydrogen co-product. The routes differ on the second, because the product is made at a different electrode. ECO makes adipic acid by anodic oxidation, so the cathode is free to reduce water and every electron evolves hydrogen, a yield going as $1/FE$; the biomass cathode makes the product itself and hydrogen evolution competes with it, so only the $(1 - FE)$ share of charge that misses the product evolves, going as $(1 - FE)/FE$. Either way it is netted against the route's own demand, a shortfall bought at the scenario price and carried in scope-3, a surplus burned for steam and carried in scope-2 (§3.3.2).

Sweeping FE across the tiers is how the study shows what a different energy demand is worth to the electrified routes, on both levelised cost of adipic acid and emissions. It is not the same as sweeping cell voltage, even though the two enter Eq. (3.1) only through their ratio: voltage would move the energy alone, while efficiency moves the hydrogen with it. Current density does not enter Eq. (3.1): it sets electrode area and so the stack capital, which is why j is pinned here and its effect carried in the CAPEX dimension instead.

3.2.6. Cell degradation and electrode replacement

Progressive cell decay, its mechanisms, and the evidence available for each cell type are set out in Section 2.4.5. Two points from there govern the treatment here. Yeo holds cell performance constant over the plant life, so degradation enters neither his lifetime electricity demand nor his capital reinvestment. And no published test on either cell runs much beyond 100 h, so the parameters below are bracketing estimates transferred from the mature-stack analogues rather than measured service lives. They are carried to show how far each route's standing depends on cell durability, not to predict how long a cell will last.

The two cells are modelled differently because they fail differently. The ECO cell degrades by a rising cell voltage at constant Faradaic efficiency, so it carries both a rising specific energy and periodic membrane-electrode replacement: end of life is taken at a 10 % cell-voltage rise, the PEM-electrolyser convention [108], which fixes the degradation rate once the stack lifetime is chosen. Because the rise is linear in time, the levelised specific energy over an interval sits half way up it, 5 % above the beginning-of-life value, and that penalty is the same for every stack lifetime; what the tier changes is how often the replacement is bought. Replacement is charged against the *bare stack*, the membrane-electrode assembly and frames rather than the installed system, at 0.35 of its capital and swept across 0.25–0.50 [109, 108]. The BIO cell instead corrodes at essentially constant voltage, so it has no energy channel and only its cathode is replaced, at intervals of about 7, 4 and 2 years for the optimistic, central and conservative tiers, derived from lead-cathode corrosion rates of 0.35–0.9 mm yr⁻¹ in the adiponitrile analogue [110, 111].

For both routes a replacement is booked as a capital outlay in the year it occurs and discounted within the levelised-cost cash flow, matching the stack-replacement term of Eq. (3.2). The central tier is carried in the headline

case, and because the underlying durability tests are short, the optimistic-to-conservative range and the cathode-cost bracket are swept as sensitivities rather than fixed. Disabling degradation and replacement recovers Yeo's constant-performance basis, which is how the reproduction of Section 3.3.1 is run. Per-tier parameters are collected in Table 3.2.

Table 3.2: Cell-degradation parameters by tier. No long-term durability data exist for either cell; the tiers are bracketing estimates anchored to mature-process analogues, not measurements. Mechanisms and sources are given in the text.

Route	Tier	Replacement interval	Voltage rise at EoL ¹	Cost / replacement
ECO	Optimistic	~5 yr (40 000 h)	+10 %	\$1.7–3.4 M ²
ECO	Central	~2.5 yr (20 000 h)	+10 %	\$1.7–3.4 M ²
ECO	Conservative	~0.6 yr (5 000 h)	+10 %	\$1.7–3.4 M ²
ECH	Optimistic	7 yr	—	\$2.5–12 M ³
ECH	Central	4 yr	—	\$2.5–12 M ³
ECH	Conservative	2 yr	—	\$2.5–12 M ³

¹ The +10 % rise is the PEM convention of 4.8 mV/khr over a 40 000 h stack life at 1.9 V/cell [108]; the tier sets the interval, not the rise. ² ECO bare stack 209 \$/kW; the central 0.35 is the DOE replacement charge of 11 % of installed capital (≈ 15 % of direct, or 33 % of stack at a 45 % stack share) [108], and the 0.25–0.50 range is the spread in how much of the stack sits in the replaceable catalyst-coated membrane [109]. ³ ECH cathode is an indicative bracket (fabricated-lead floor to a judgement ceiling); the headline uses its mid value.

3.2.7. Model scope and validation

The process models resolve each route to a converged mass and energy balance, from feedstock to product, which is the level at which the downstream assessment consumes them. Limitations that bear on interpreting the results rather than on the models are taken up in the Conclusions (Chapter 6).

Two simplifications bound what can be asked of them. The electrochemical cells are imposed rather than resolved (§3.2.2), so the model evaluates the economic consequences of specified operating points but cannot predict how performance changes with electrode area or operating conditions. And the Aspen flowsheets are adopted from Yeo (2025) without structural modification, apart from the heat integration of §3.2.4 and the BIO polishing reactor, with the hydrogen-fired boiler costed outside the flowsheet (§3.3.2); his engineering exclusions carry over, pumping energy and capital on the grounds that pumping work is small against gas compression (Yeo, 2025, §3.3), and wastewater as a bulk disposal stream at \$2.2/tonne. The BIO route is evaluated on a glucose-based fermentation pathway only, for the reason given in the literature review (LR §2.2.3); a robustness check on that assumption is reported with the results.

Validation is against operating data where it exists and internal where it does not. Parameters are anchored to plant measurements where they can be, notably the N₂O abatement capital and availability (§3.4) and the electrode-replacement intervals from chlor-alkali and adiponitrile practice (§3.2.6). The flowsheets admit no such comparison, being inherited, generic in the conventional case and without operating counterpart in the electrified ones, so there the checks are internal: mass and energy balance closure on each flowsheet, separation duties against standard design heuristics, and the reported utility totals against the Aspen calculation reports. That last check found the two ECO feed compressors missing from the reported electricity total; the corrected duty is used throughout and its cost effect is given with the other corrections in §3.3.3. For the cells no public benchmark exists at all, so their parameter uncertainty is propagated through the scenario framework instead.

3.3. Economic assessment

This section sets out how the levelised cost of adipic acid is computed: the cost engine, the capital and operating build-up, and the levelisation itself, all evaluated at the central input values. The parameters swept on top of this base case are set out in Section 3.5.

3.3.1. Computational framework and the OpenPyTEA cost engine

The evaluation is implemented as a Python framework that reads the converged mass and energy balances from the Aspen Plus models (§3.2.1) and returns the LCOA (Eq. (3.2)) and cradle-to-gate GHG intensity (Eq. (3.7)) for every point in the evaluation space. Its economic core is **OpenPyTEA**, an open-source techno-economic assessment toolkit developed at TU Delft [112] (MIT licence, v2.1.0), which returns the CAPEX, OPEX and levelised cost of product from the Towler & Sinnott correlations [64] (§3.3.3). Costing on a published, openly

documented engine rather than a bespoke spreadsheet is a deliberate choice: it makes the cost basis reproducible, and comparable with other techno-economic assessments built on the same standard correlations rather than on in-house factors.

Four layers are added on top of it for this work: the EU ETS term C_{ETS} (Eq. (3.3)), which turns the levelised cost into the policy-inclusive LCOA; the cradle-to-gate GHG intensity (§3.4); the electrochemical-cell model, whose specific electricity demand w_{elec} (Eq. (3.1), §3.2.5) enters both the energy cost and the scope-2 GHG; and the scenario machinery (§3.5.1).

The same equipment is then costed two ways in parallel. Every result reported in this thesis uses the OpenPyTEA costing: its correlation library wherever a correlation exists, and direct purchased costs for the units it does not cover (§3.3.3). Alongside it, the factorial method of Yeo (2025) [9] is retained as a reference costing, used only to reproduce his published figures as a cross-check on the inherited cost data. It runs under his own economic convention, which differs from the one used here: he converts the total capital into a fixed annual charge and divides it, with the annual operating cost, by the annual production, carrying no construction period and producing at design rate from the first year. The levelisation used here (Eq. (3.2)) is the general case of that same calculation, with capital drawn over a build period and output ramping to the design rate. His convention is kept unchanged there, so the two sets of figures are never mixed. What the cross-check surfaced, and what else differs from his basis, is set out in §3.3.4.

3.3.2. Levelised cost of adipic acid

The levelised cost of adipic acid (LCOA) is the constant product price that yields a zero net present value over the plant lifetime, capturing all capital, operating, and policy costs in a single comparable metric. It is therefore the constant price yielding $\text{NPV} = 0$ at the 8 % real discount rate over the 22-year project life, and embeds a return on capital at that rate in the same way a levelised cost of electricity does. It is also inclusive of the plant's EU ETS position: the carbon cost enters the levelisation as an annual operating term rather than being added afterwards, which is where the free-allocation result of §4.5.1 sits. Yeo (2025) [9] reports the same quantity under the name minimum selling price; his figures keep his term throughout, and because both are $\text{NPV} = 0$ break-evens they are directly comparable. Its capital term is the total installed cost built up in §3.3.3. Its operating term is the annual OPEX of feedstock, electricity, utilities, waste disposal, maintenance and labour, the electricity scaling with both the degradation-driven energy demand and the country-specific EHV price. For the electrochemical routes, stack or electrode replacement enters as periodic reinvestment at interval τ_{stack} (§3.2.6), and the EU ETS position enters through the C_{ETS} term below.

The analysis is pre-tax: corporate tax and depreciation schedules are omitted because they are route-independent at this level of comparison, while working capital is retained as OpenPyTEA computes it. The cell-CAPEX learning curve is not a headline input; it enters only as a narrative lever for the electrochemical cells (§3.5.3).

Levelised cost

The LCOA is the levelised cost: the plant's whole-life cost divided by its whole-life output, both discounted to present value, using the construction and start-up schedule of Table 3.3 – the capital draw ϕ_t , the fixed-cost switch o_t , and the capacity ψ_t :

$$\text{LCOA} = \frac{\sum_{t=1}^N \frac{\phi_t C_{\text{fix}} + o_t C_{\text{OPEX}}^{\text{fix}} + \psi_t C_{\text{OPEX}}^{\text{var}}(t) + C_{\text{ETS}}(t) + C_{\text{stack}}(t)}{(1+r)^t}}{\sum_{t=1}^N \frac{\psi_t Q}{(1+r)^t}} \quad (3.2)$$

where t is the project year ($t = 1 \dots N$), $N = 22$ the project life (two construction years and 20 producing years, with the capital draw and production-ramp schedule of Towler & Sinnott (Table 9.2) set out below), $r = 8\%$ the real discount rate, C_{fix} the total fixed capital, Q the design-basis annual production volume, and $C_{\text{OPEX}}^{\text{fix}}$, $C_{\text{OPEX}}^{\text{var}}$, C_{ETS} , C_{stack} the annual fixed operating, variable operating, ETS, and stack-replacement costs. The variable term carries a year index because the electrochemical cells' specific energy rises as they degrade, so their electricity demand is not constant over the plant life (§3.2.6). The schedule sequences ϕ_t , o_t and ψ_t and the conventions they carry are given with Table 3.3.

The 8 % real discount rate reflects the chemical sector’s cost of capital: KPMG’s 2024 survey reports a nominal post-tax WACC of 8.5 % for chemicals and pharmaceuticals and 9.2 % for chemicals alone [113], so 8 % real sits at the sector’s upper end once inflation is removed, a mildly conservative choice. It is applied uniformly across routes so the LCOA reflects each route’s own capital and energy rather than route-specific financing. Yeo’s 15 % is retained as a sensitivity (§3.3.4).

Table 3.3: Construction and start-up schedule [64] (Ch. 9, Table 9.2), giving the sequences ϕ_t , o_t and ψ_t over which Eq. (3.2) is discounted. The conventions they carry are given in the text.

Project year	1	2	3	4	5–22
Capital draw ϕ_t	0.30	0.60	0.10	0	0
Fixed-cost switch o_t	0	0	1	1	1
Capacity ψ_t	0	0	0.30	0.80	1

Construction and start-up schedule

A plant is built over several years and reaches full output only gradually, so capital is committed well before any revenue arrives, and discounting weights that early spending heavily. The cash flow therefore follows the typical construction and start-up schedule of Towler & Sinnott [64] (Ch. 9, Table 9.2), set out in Table 3.3: fixed capital is drawn 30/60/10 % over the first three years, and production begins in the third year at 30 % of design basis, reaches 80 % in the fourth, and holds at the design rate from the fifth. Two years therefore carry capital but no output, and 20 produce, though the 30 % and 80 % start-up years leave 19.1 years of full-rate output over the plant life. For the headline case the first production year is 2030, so construction runs over 2028–2029 and the plant produces over 2030–2049; the commissioning years of Table 3.1 are these first-production years, and the sensitivity shifts the whole schedule with them.

Three conventions follow. Fixed operating costs – labour, maintenance, insurance – do not scale with output, so the switch o_t holds them off through the two construction years and turns them on in full at first production. This is the one adjustment to OpenPyTEA’s stock cash flow, which charges fixed operating costs in every project year. Variable costs do scale with output through ψ_t , but the same factor multiplies the output in the denominator of the levelised cost, so it cancels from the per-kg variable cost and the ramp loads only the capital and fixed-cost terms. The one place the ramp does not cancel is a per-kg term that itself varies from year to year: the scope-2 GHG intensity under a declining grid (§4.3), where each year’s intensity is weighted by ψ_t when levelised, so the two low-output start-up years, which fall on the dirtiest early grid, carry less weight than the full-rate years. Working capital is drawn in the final construction year and released in the last, contributing only its financing cost.

This staged timing is more detailed than Yeo (2025) [9], whose annualised-capital method assumes full output from the first year with no construction period, so the absolute levelised costs of adipic acid reported here sit above what his flat timing would give. All three routes share the same schedule, so the comparison between them is unaffected; the size of this timing effect, isolated by re-running the routes without the ramp, is assessed in the headline comparison (§4.1).

EU ETS cost

EU ETS costs enter through the C_{ETS} term:

$$C_{\text{ETS}}(t) = [e_{\text{scope1}} - \alpha(t) \cdot B_{\text{AA}}] \cdot P_{\text{CO}_2,s}(t) \cdot Q \quad (3.3)$$

where e_{scope1} is the verified scope-1 emission intensity (tCO₂e/t product), $B_{\text{AA}} = 1.40$ tCO₂e/t the EU ETS product benchmark for the 2026–2030 allocation period, published in June 2026 [77] and the value a plant first producing in 2030 is allocated on (LR §2.4.1), $P_{\text{CO}_2,s}(t)$ the carbon price under scenario s (€/tCO₂, held flat over the plant lifetime), $\alpha(t)$ the free allocation fraction (share of the benchmark received at no cost; $\alpha = 1$ means full free allocation up to B_{AA} , $\alpha = 0$ means every emitted tonne must be purchased), and Q the annual production volume, here in tonnes per year to match the emission terms. The bracketed term is positive when scope-1 emissions exceed the free allocation (net ETS cost) and negative when emissions fall below (net ETS revenue from sellable surplus allowances).

Free allocation

Adipic acid’s free allocation is legislated only as far as 2030, and what follows is open (LR §2.4.1), so it is carried as a bound rather than as a forecast. The headline retains it in full at $\alpha = 1$, on the same convention used for

every other policy input here of holding the value that stands today; the bound withdraws it entirely at $\alpha = 0$, which is the most the change could be worth. The credit is benchmark-based, so it moves all three routes by the same amount: the bound is on the level at which they sit, not on the distances between them. Both ends hold α flat across the whole producing life, which no realised outcome would do: the earliest either route to withdrawal can take effect is 2031, part-way through the twenty producing years of a plant first producing in 2030, and a withdrawal phased in over several years rather than applied at once lands further in still. A realised levelised fraction therefore falls between the two ends. Levelised cost of adipic acid is linear in α , so any such case is read from the bounds rather than run separately.

Waste disposal

The waste-disposal term covers process effluent and by-product disposal, principally wastewater together with the glutaric-acid and dilute-adipic-acid by-product streams. Each stream is priced at a per-kilogram disposal rate after Towler & Sinnott [64] and applied as Yeo (2025) [9] models it;³ multiplied by the route stream flows these give 8.15, 8.96, 6.60 and 6.01 M\$/yr for CON, ECO, BIO_{t3HDA} and BIO_{AA} respectively, each charged as an annual operating cost and reported as its own component of the levelised cost of adipic acid. The biomass figures differ from Yeo's published 5.07 and 5.50 M\$/yr because his totals net off a by-product-disposal credit for the surplus cathodic hydrogen, which is withdrawn here and replaced by a steam credit for the same stream (*Cathodic hydrogen*, below). CON and ECO carry no such surplus and are unchanged.

Cathodic hydrogen

The electrochemical cells evolve hydrogen at the cathode, and each route uses it rather than selling it. In the ECO route it offsets the benzene-hydrogenation feedstock cost as an avoided purchase, with any residual demand bought at the prevailing market price.

Both biomass routes evolve 273.6 kg/h of cathodic hydrogen at the ECH cell (99.7 mol% H₂, 30 bar), and neither buys any. The route to adipic acid feeds its reactor from the cathode, the Aspen flowsheet sending 181.8 kg/h to the reactor, which consumes 169.5, and leaving a 91.7 kg/h surplus; the route to t3HDA hydrogenates electro-chemically at the cathode rather than with molecular hydrogen, so it consumes none of the evolved stream and all 273.6 kg/h is surplus. These are the values at the central FE 0.40. Hydrogenating muconic acid and the parasitic hydrogen-evolution reaction both draw two-electron cathode current, so the by-product scales as $(1 - \text{FE})/\text{FE}$; the surplus, its steam credit and the boiler are recomputed at each cell-performance tier by scaling this 273.6 kg/h from the FE 0.40 anchor. At the optimistic tier the route to adipic acid evolves less than its reactor needs and buys the shortfall as makeup, no longer self-sufficient. The surplus raises low-pressure steam on site: at a lower heating value of 120 MJ/kg [114] and a mid-range 85 % boiler efficiency (Towler & Sinnott [64], Ch. 3) it yields 7.8 and 2.6 MW on the routes to t3HDA and to adipic acid, meeting 76 % and 25 % of their steam demand. A packaged boiler raises it, costed through the same correlation as the rest of the plant (§3.3.3); that costing is conservative on two counts, set out with the rest of the valuation in §B.3. Burning the surplus on site also removes a load the reference flowsheet carries: it compresses the surplus hydrogen through a three-stage train for export (the compressed stream leaves at the system boundary), which the on-site steam use makes unnecessary. Dropping that export compression lowers the biomass routes' electricity by 536 kW at the t3HDA endpoint and 180 kW at the adipic-acid one, which keeps only the compressor feeding its polishing reactor, and removes those compressors from the fixed capital.

Yeo credits the same surplus, but only in cash: his stream costing gives hydrogen a $-\$0.70/\text{kg}$ disposal price at a zero emission factor, embedded in the inherited waste figure. Steam displacement is the alternative, of similar cash value, so his credit is withdrawn from the waste line wherever the steam credit applies; the steam is valued at the plant's own purchase price, since it never exceeds demand. The difference is in the carbon, not the cash: burning hydrogen yields only water, so displaced steam removes a real emission the disposal credit leaves untouched, giving the biomass routes a scope-2 reduction the inherited treatment does not. Per-route effects are reported in §4.2.3. On-site steam is the conservative credit: merchant sale would be worth more per kilogram, and why it is not valued here is set out in §B.3.

³Wastewater is charged at \$2.2/tonne; the glutaric-acid and dilute-adipic-acid streams at \$0.34/kg, costed as the alkali required to neutralise them, following Yeo's stream costing.

3.3.3. Capital cost (CAPEX)

Equipment costing

Each unit's purchased cost is obtained from its size parameter S (area, mass, power, or volume, taken from the Aspen models) through the Towler & Sinnott correlation

$$C_e = a + b S^n, \quad (3.4)$$

with coefficients (a, b, n) from Table 7.2 of [64] (US Gulf Coast, January 2010 basis). Two groups fall outside this correlation route. The electrochemical cells and the electrodialysis unit have no Towler correlation, being too specialised to cost from a standard curve. The reactors, crystallisers, centrifuges, coolers and turbine do, but are not re-derived from a size parameter here; this work retains the purchased values recorded in Yeo (2025)'s workbooks [9]. Both groups enter at their bare purchased cost (Yeo's workbook column B); installation is then applied uniformly in the build-up below rather than inherited from Yeo's Lang-factor scheme [9], which installs the electrochemical cells and electrodialysis unit on a separate "packaged" factor of 3.1 against the 4.4 used for fluids equipment. Costs are escalated to the 2024 reference year with the Chemical Engineering Plant Cost Index [115, 116], following the cost-escalation method of Towler & Sinnott [64] (Ch. 7, Eq. 7.14),

$$C_{2024} = C_{2010} \frac{\text{CEPCI}_{2024}}{\text{CEPCI}_{2010}}, \quad \text{CEPCI}_{2024} = 800, \quad \text{CEPCI}_{2010} = 532.9. \quad (3.5)$$

Capital build-up

OpenPyTEA assembles the fixed capital factorially from the purchased costs following Towler & Sinnott's factorial method [64]. Each item's purchased cost $C_{e,i}$ is raised to an installed cost by a field factor $f_{\text{inst},i}$ covering piping, instrumentation, electrical, civil and erection work; summed, these give the inside-battery-limits cost $\text{ISBL} = \sum_i C_{e,i} (1 + f_{\text{inst},i})$ on the US Gulf Coast basis. The field factor is not a single value but Towler's equipment-class set (Table 7.5) [64]: about 2.2 for fluid and mixed equipment and 1.5 for solids in carbon steel, so installed costs run near 3.2 and 2.5 times purchased. Off-site ($f_{\text{OSBL}} = 0.40$), design-and-engineering

Table 3.4: Installed-cost factors (purchased $\rightarrow$ installed) by equipment class: Yeo's Lang factors against this work's Towler per-item factors.

Equipment class	Yeo (Lang)	This work (Towler)
Fluids (HX, reactors, columns, compressors, turbines)	4.4	3.2
Solids (centrifuges, crystallisers)	3.5	2.5
Packaged (coolers, electrodialysis)	3.1	3.2

($f_{\text{D\&E}} = 0.25$) and contingency ($f_{\text{cont}} = 0.10$) factors are then applied,

$$C_{\text{fix}} = (\text{ISBL} + \text{OSBL})(1 + f_{\text{D\&E}} + f_{\text{cont}}), \quad \text{OSBL} = f_{\text{OSBL}} \text{ISBL}. \quad (3.6)$$

giving $C_{\text{fix}} = 1.89 \text{ISBL}$. The contingency is held at Towler's 10 % minimum on all routes. Towler allows up to 50 % where the technology is uncertain, which would apply to the electrochemical cells; that risk is carried instead as an explicit uncertainty band on the capital estimate rather than folded into the point estimate, so the central capital of the electrified routes sits at the lower, not the conservative, end of Towler's range. Working capital follows the inherited convention of Yeo (2025) [9] at 15 % of total capital, drawn at start-up and recovered at end of life, with its interest carried in fixed OPEX. The location factor applies to CAPEX only: the country dimension (§3.5.1) sets electricity price and grid intensity, not the plant-cost location.

Capital uncertainty: AACE equipment weighting

The study is a single Class-4 estimate under AACE 18R-97 [63]: a complete flowsheet with sized equipment and no detailed engineering. Its accuracy envelope is resolved across three equipment groups by how completely each is defined. The correlation-costed items, sized from the converged flowsheet, carry $-15/+30\%$: the heat exchangers, vessels, columns and compressors. The equipment costed directly at its recorded purchased price carries $-20/+40\%$: the reactors, centrifuges, crystallisers, coolers and turbine. The electrochemical cells, concept designs with no test beyond about 100 h, carry $-30/+50\%$, the widest of the three; the electrodialysis unit is weighted with them, having no Towler correlation either (§3.3.3). Each route's overall band is these three weighted by its own shares of inside-battery-limits capital, so the band is both asymmetric and route-specific. The off-site, engineering and contingency factors are uniform multipliers on that base and therefore cannot move it, and the N_2O abatement facility, which enters outside the factorial build-up, is held fixed while the band is applied; the resulting bands and their effect on the levelised cost of adipic acid are reported with the results (§4.5.4). It is an uncertainty *on* the estimate and not an allowance added to it: the contingency stays at the 10 % above.

N₂O abatement capital

The conventional route's N₂O decomposition reactors carry no capital in Yeo's estimate, and so no fixed operating cost either, since that is charged on the capital base; the abatement train is absent from his cost build-up entirely. Here it is costed, priced from operating-plant data rather than through the equipment correlations. The figure is an installed cost for the complete thermal-decomposition facility, the reactor together with its burner, refractory, high-temperature heat recovery and controls, rather than a bare reactor vessel, so it is added directly to the fixed capital of Eq. (3.6), outside the factorial build-up rather than through it. Being an installed field cost, it is not multiplied by the installation and indirect factors of the build-up; it does, however, carry the plant's maintenance and insurance loading at the standard rates, since a running abatement facility needs the same upkeep as any other installed equipment. A redundant two-train configuration is costed rather than a single unit: one thermal decomposition train destroys almost all tail-gas N₂O while running, but is available only 85–90 % of the year, so its realised annual abatement would sit near that level however well the train itself destroys. The 99 % central case is therefore costed with a second train. Each facility is taken at roughly 10 M€, 20 M€ for the pair [31, 34]. The LANXESS reference plant is of the same order of magnitude as the design basis here, about 88 kt/a of adipic acid [117] against 100 kt/a, so the facility cost transfers without a scale correction.

That is a 2010-vintage European installed cost, so it is escalated on a euro basis rather than with the US-derived CEPCI: the euro-area capital-goods producer price index rose about a quarter over 2010–2024 [118], converted at the 2024 ECB average of 1.10 USD/EUR [119]. The figure already covers the tail-gas heat exchangers and expander costed separately in the inherited flowsheet, so that overlap, about \$1.9 M, is subtracted, leaving a net addition of about \$25 M to conventional fixed capital. It is an indication rather than a precise estimate.⁴

The abatement facility's cost is expressed per tonne of CO₂-equivalent it removes and analysed against two benchmarks: the abatement cost operating plants report, as a plausibility check on the facility figure taken here, and the prevailing carbon price, to establish the rationale for abating at all. On the same basis the redundant two-train configuration the 99 % case assumes is set against a single train, taken indicatively at 90 % realised abatement: what the second train buys is the difference between a level a single train could plausibly sustain and the 99 % central case, not a level it would necessarily fall to.

3.3.4. Differences from Yeo (2025)

The cost basis departs from Yeo (2025) in method, in corrections to his cost data, and in one addition of scope. Equipment costing itself is common ground: it follows the same Towler & Sinnott correlations he used (Eq. (3.4)), applied here through OpenPyTEA rather than his own spreadsheet, which is what allows the two to be compared item by item.

Two corrections restore items Yeo mis-costed in fixed capital. First, Yeo's centrifugal-compressor constants sit a factor of ten below the Towler & Sinnott Table 7.2 values ($a, b = 58,000, 2,000$ against 580,000, 20,000), while his exponent and validity range reproduce the published row exactly ($n = 0.6, 75\text{--}30,000\text{ kW}$), so the inherited machine costs are an order of magnitude too low; they are restored to the published correlation. His gas-turbine row carries the same erroneous constants, Table 7.2 having no gas-turbine correlation to draw on. Second, the N₂O decomposition reactors (R5, R6) are modelled on the CON flowsheet to close the tail-gas balance but left uncoded, which would leave N₂O abatement as a free benefit; their capital is added here, costed as set out in §3.3.3.

One addition of scope also enters fixed capital: the BIO route gains a catalytic hydrogenation reactor and adipic-acid recovery train not present in Yeo (§3.2.3), whose equipment runs through the same Towler build-up.

A further correction concerns duty rather than capital: in the ECO flowsheet the two feed compressors compute shaft work (1,626 and 1,587 kW) but were never assigned to the electricity utility, so the reported utility total of 61,060 kW omitted them; the corrected draw is 64,273 kW, used throughout. Table 3.5 collects the capital and duty items and their impact.

Three further differences are of overall method rather than of individual items. Installed capital is built from each item's bare purchased cost through Towler's per-item field factors (§3.3.3) rather than the class-wide Lang factors Yeo applied, so the two reach installed cost by different routes. The cash flow is levelised over a construction and start-up schedule rather than as an annualised capital charge (§3.3.1), and the discount rate is 8 % against the

⁴A convert-to-dollars-then-CEPCI treatment (the 2010 exchange rate escalated with CEPCI) would give about \$38 M rather than \$25 M; the euro basis is preferred because this is a European field cost. The escalation basis is the main source of that range.

15 % he applied, with his value kept as a sensitivity (§3.3.2). All three are differences of convention rather than corrections, and they set the basis on which the two sets of figures are compared.

All of the above concerns the cost basis. The input values differ as well: utilities are set on an ARRA-wide basis rather than Yeo’s Rotterdam-specific assumptions (§3.2.4), electricity price and grid carbon intensity are country-specific, feedstocks are drawn from explicit procurement scenarios, and carbon pricing enters at all. Those are set out in Section 3.5. Together they set a wider geographic scope and a different input basis.

Table 3.5: Capital and duty changes relative to the inherited Yeo (2025) cost data. The first three are capital changes; the fourth is an electricity-duty correction carried in OPEX. The compressor-correlation capital impact is reported with the headline comparison in Chapter 4.

Change	Route(s)	Impact
Compressor correlation restored	CON, ECO	machine cost $\sim \times 10$; see Ch. 4
N ₂ O decomposition reactors costed	CON	+\$25 M fixed capital
Catalytic hydrogenation reactor + AA recovery added	BIO	see §3.2.3
Compressor electricity tagged (duty)	ECO	+3,213 kW

3.4. Environmental assessment

GHG emissions are assessed on a cradle-to-gate basis following the GHG Protocol scope framework [120], reported per functional unit using IPCC AR6 100-year GWPs [7]. This is a GHG intensity assessment, not a full multi-impact LCA. Scope 1 covers on-site emissions (dominated by N₂O for CON, near-zero for the electrochemical routes); scope-2 covers grid electricity (dominant for ECO and BIO) and purchased process steam; scope-3 covers upstream feedstock and, for the BIO route, a biogenic carbon credit following the cradle-to-gate convention of the GHG Protocol [120].

The levelised GHG intensity is the output-weighted mean of the annual cradle-to-gate emissions. Only the grid term varies from year to year, so only it is averaged:

$$\bar{e}_{\text{GHG}} = e_{\text{scope1}} + e_{\text{steam}} + e_{\text{scope3}} + \frac{\sum_{t=1}^{N_{\text{prod}}} \psi_t \text{CI}_c(t) w_{\text{elec}}(t)}{\sum_{t=1}^{N_{\text{prod}}} \psi_t} \quad (3.7)$$

where $N_{\text{prod}} = 20$ is the number of producing years, the 22-year project life of Eq. (3.2) less the two construction years in which the plant makes nothing, ψ_t the output of year t as a fraction of design rate (0.30 and 0.80 in the two ramp years and 1 thereafter, so that the start-up years, which fall on the dirtiest early grid, carry proportionately less weight; $\sum_t \psi_t = 19.1$), e_{scope1} the on-site emission intensity (static; dominated by N₂O for CON), $\text{CI}_c(t)$ the country-specific grid carbon intensity in year t (declining over time as the electricity mix decarbonises), $w_{\text{elec}}(t)$ the total electricity consumption per kg product in year t (rising over time due to cell voltage degradation), e_{steam} the process-steam emission intensity (static; carbon from purchased process steam counted in scope-2 following Yeo’s convention), and e_{scope3} the upstream feedstock emission intensity (static). The grid-electricity scope-2 term $\text{CI}_c(t) \cdot w_{\text{elec}}(t)$ captures the interaction between grid decarbonisation and degradation-driven electricity demand. Grid carbon intensity uses each country’s 2024 lifecycle, consumption-based baseline [94] in the base case, with a supplementary temporal sensitivity rebasing the EU Reference Scenario 2020 trajectory onto this anchor; the full treatment is set out in Section 3.5.1. Scope 2 is reported location-based, from the carbon intensity of the grid supplying the plant. A market-based treatment using guarantees of origin would lower the emissions reported, but not the carbon intensity of the electricity actually delivered, and is outside this work’s boundary [121, 120].

Scope 1: N₂O abatement

For the conventional route the on-site scope-1 term is the residual process N₂O after abatement,

$$e_{\text{scope1}}^{\text{CON}} = (1 - \eta) g_{\text{N}_2\text{O}} \text{GWP}_{100}, \quad (3.8)$$

with abatement efficiency η (swept 95–100 %; 99 % central for a modern redundant thermal-destruction unit), unabated generation $g_{\text{N}_2\text{O}} = 0.30 \text{ kg N}_2\text{O/kg AA}$ (IPCC stoichiometric), and $\text{GWP}_{100} = 273$ [7, 35]; the same abated value drives the ETS liability in Eq. (3.3).

The 99 % central case is the redundant twin-train configuration costed in §3.3.3: realised annual abatement is the product of destruction efficiency and availability, so near-complete removal needs the second train. The achieved-rate evidence, named plants and economic rationale are reviewed in §2.2.1. The sweep brackets this physically: a single non-redundant train capped near 90 % as the downside, complete destruction as the upper bound.

Material efficiency

Alongside the two headline KPIs, LCOA and GHG intensity, material efficiency η_{mat} is reported as a diagnostic: the actual product yield relative to the stoichiometric maximum from the carbon-backbone feed,

$$\eta_{\text{mat}} = \frac{\dot{m}_{\text{product}}/\dot{m}_{\text{feed}}}{\nu M_{\text{product}}/M_{\text{feed}}}, \quad (3.9)$$

with $\dot{m}_{\text{product}}/\dot{m}_{\text{feed}}$ the actual mass ratio from the Aspen balance, $\nu = 1$ the product-to-feed mole ratio for all routes, and M the molar masses of key feed and product (benzene for CON and ECO, the muconic-acid content of the broth for BIO). Because the flowsheets are inherited, η_{mat} stays in line with Yeo for CON, ECO and BIO_{t3HDA}; BIO_{AA} differs, its catalytic hydrogenation step and recovery train not being present in Yeo. It is not a selection criterion, since its consequences already enter LCOA through feedstock cost and GHG through scope-3 emissions; it helps explain why the two headline KPIs differ across routes.

3.5. Scenario evaluation

This section operationalises the scenario framework of the literature review, whose scenario-framework and feedstock-decision tables (LR Tables 2.6 and 2.5) are the basis for both the coherent-future scenario analysis and the one-at-a-time sensitivity scans. It takes the dimensions, tier labels, values, and sources catalogued there and specifies how they feed the evaluation procedures. The framework is implemented in Python through a common interface returning LCOA and GHG intensity with component-level breakdowns; input interactions (e.g. cell performance amplifying the electricity-price impact through specific energy consumption) emerge naturally from joint evaluation rather than from a separate interaction model. The remainder of the section addresses the lifetime treatment of inputs (Section 3.5.1), the HNO₃ procurement scenarios (Section 3.5.2), the cell-CAPEX learning lever (Section 3.5.3), the coherent-future narratives (Section 3.5.4), the load-following sensitivity (Section 3.5.5), and the six evaluation modes that combine these (Section 3.5.6).

3.5.1. Treatment of inputs over plant lifetime

Most scenario inputs are held constant over the 20-year plant lifetime; only an input with credible time-resolved data is given an explicit trajectory. The tier-based dimensions stay flat because their Low / Central / High tiers already span the plausible future range (LR §2.4.1), so a shift in any direction falls within those bounds without a forward trajectory. Two components are the exception, carrying explicit time dependence and specified below: a supplementary temporal sensitivity on grid carbon intensity, and a supplementary rising CO₂-price trajectory.

Carbon price is held flat at €75, €115, and €145/tCO₂. Free allocation is likewise held flat, at the full $\alpha = 1$ of the current carbon-leakage regime, and bounded by its complete withdrawal at $\alpha = 0$ (Section 3.3.2, LR §2.4.1). A supplementary sensitivity adopts the CIP (2026) Net Zero rising trajectory, 140 to 210 €/tCO₂ over 2030–2050 [76], whose rationale is given in LR §2.4.1.

Grid carbon intensity is country-specific, with two treatments. The default static treatment holds each country at its 2024 lifecycle, consumption-based intensity over the plant lifetime [94], matching the GI-tier values in Table A.2. As a supplementary sensitivity, each country's intensity is projected forward by rebasing the EU Reference Scenario 2020 (REF2020) PRIMES trajectory [122] onto the 2024 lifecycle anchor:

$$CI_c(t) = CI_c^{2024} \times \frac{CI_c^{\text{REF2020}}(t)}{CI_c^{\text{REF2020}}(2024)} \quad (3.10)$$

where the 2024 denominator and all intermediate plant years are linearly interpolated between REF2020's five-year nodes. The rebasing preserves the country-specific decarbonisation pace projected by PRIMES while anchoring the absolute level to verified 2024 measurements, addressing the offset between REF2020's pre-Fit-for-55 baseline and observed 2024 intensities. The resulting per-country trajectories are shown in Figure A.2 (§A.3). REF2020 predates recent policy and market developments, so the trajectory tests how route GHG responds to grid decarbonisation pace within a plausible EU range rather than serving as a country-level prediction. Post-2022 national projections broadly support its shape, with Belgium the partial exception; that cross-check, and the static tier that brackets it, are in §A.2.

Other inputs are held flat for input-specific reasons: electricity and feedstock prices because no reliable persistent forecasts exist over the 20-year horizon; cell performance because it represents the technology-readiness state at

commissioning; H_2 , HNO_3 , and benzene as discrete procurement scenarios (§2.4.6) capturing supplier-choice exposure whose consequences do not evolve over time within a given scenario.

3.5.2. HNO_3 procurement scenario construction

In §A.5.3 the EI-A and EI-B procurement scenarios are defined through their component emission factors (regional N_2O process emission factor and upstream NH_3 cradle-to-plant-gate emission intensity); this subsection combines those components into the per-scenario total cradle-to-gate scope-3 emission intensity. All other per-feedstock parameter values are taken directly from §2.4.6 and Table 2.5 without methodology-side construction.

HNO_3 cradle-to-gate scope-3 emission intensity is varied as a two-scenario discrete procurement choice (EI dimension), capturing the supplier-choice drivers identified in §A.5.3. Per-scenario total values (Table 3.6) combine the regional N_2O process emission factor and the upstream ammonia cradle-to-plant-gate emission intensity:

$$\text{EI}_{\text{HNO}_3} = e_{\text{N}_2\text{O}} \cdot \frac{\text{GWP}_{100}}{1000} + e_{\text{NH}_3} \cdot r_{\text{NH}_3/\text{HNO}_3} \quad (3.11)$$

where $e_{\text{N}_2\text{O}}$ is the regional N_2O process emission factor (kg N_2O /t HNO_3 ; the factor of 1000 converts to per-kg HNO_3 basis), e_{NH_3} is the upstream ammonia cradle-to-plant-gate emission intensity (kg CO_2e /kg NH_3), $\text{GWP}_{100} = 273$ [7], and $r_{\text{NH}_3/\text{HNO}_3} = 0.284$ kg NH_3 /kg pure HNO_3 is the stoichiometric ratio at 95 % NH_3 conversion efficiency [35]. Per-region component values for $e_{\text{N}_2\text{O}}$ and e_{NH_3} are listed in Table A.6. Nitric acid is priced at the central of \$0.30/kg of contained HNO_3 (§A.5.3), applied directly to the model's stoichiometric pure-acid consumption: index and model share the 100 %- HNO_3 basis, so no grade correction is needed. The asymmetric \$0.22–0.50/kg sensitivity band is taken from the same section.

Table 3.6: HNO_3 procurement scenarios (EI dimension): **total cradle-to-gate scope-3** emission intensity per scenario, derived by combining the literature **components** in Table A.6 via the formula above. Aspen feed basis: 60 % HNO_3 solution; values per kg of 100 % pure HNO_3 shown for international-convention comparability.

Tier	Scenario	Scope-3 CI (kg CO_2e /kg pure)	Scope-3 CI (kg CO_2e /kg 60 % sol.)
EI-A	Modern European abated	0.93	0.56
EI-B	Non-EU coal-based unabated	3.99	2.39
Spread		4.3×	4.3×

The cradle-to-gate boundary captures the two dominant emission components (upstream NH_3 and on-site N_2O process). The autothermal catalytic NH_3 oxidation and waste-heat steam export make modern HNO_3 plants net energy-positive [123]; plant electricity for compression and cooling is small relative to these contributions and is excluded for compactness and due to lack of regional plant-specific data.

3.5.3. Capital-cost learning for the electrochemical cells

The conventional route is built from mature equipment whose cost is not expected to fall with further deployment. The electrochemical cells (the ECO electro-oxidation cell and the BIO electrochemical-hydrogenation reactor) are early-stage technologies (TRL 3–5) whose manufactured cost can fall as they scale; this potential is represented as a scenario lever applied only to those cells, not to the mature plant equipment or the commodity feedstocks. The cost of a manufactured energy technology follows an experience curve (Wright's law), $C = C_0 (1 - \text{LR})^n$ for n cumulative-capacity doublings at learning rate LR. Published learning rates for water electrolyzers, the closest manufactured analogue, cluster at 16–21 % per doubling (IRENA 2020 [109]; Schmidt et al. 18 % [124]; Reksten et al. 20 % for PEM [125]). The –30 % cell-CAPEX reduction used for the matured case corresponds to about two doublings at this rate (16–20 % gives a 29–36 % reduction; 30 % is a round and conservative value).

The lever is a maturity endpoint rather than a dated cost trajectory, because the cumulative-deployment data a year-indexed curve would need do not exist for these bespoke cells. It spans a plausible cost range in three settings: the matured value (–30 %, the learning-curve result above), the central estimate, and a high-cost case (+50 %, a cell whose cost does not come down, the loose end of the Class-4 estimate's accuracy envelope carried by the least-defined equipment group (LR §2.3.2)); only the cells move, with the mature equipment and feedstock unchanged. The endpoints frame a range rather than forecast it, and the lever exists to show how far the electrified routes' LCOA and viability depend on the cell cost falling. The three settings anchor the cell-CAPEX dimension of the coherent-future narratives (§3.5.4). Two caveats apply: the learning rate is measured for water electrolyzers and its transfer here is reasonable but unproven, and the reduction is contingent on stable supply of the cells' critical materials (vanadium in the ECO anode, bismuth in the BIO cathode).

3.5.4. Coherent-future narratives

Coherent-future narratives test how route-ranking holds up under bundled conditions and reveal interaction effects invisible to single-dimension sensitivity. Three named narratives bundle the policy, market, and infrastructure dimensions into internally coherent futures, evaluated as Mode 6 of the framework (§3.5.6). Each narrative pins feedstock-cost tier, H₂ procurement, HNO₃ procurement, benzene procurement, carbon-price tier, free-allocation share and grid-CI treatment to values that reflect a single underlying decarbonisation pace, and the electrochemical cells mature with the storyline on both of the dimensions that describe them: capital cost moves from pilot under lags to matured under accelerates (§3.5.3), and performance moves with it from *EC-Conservative* to *EC-Optimistic*. The two are bundled rather than separated because a cell that reaches its target efficiency and a cell that reaches its target cost are the same development, and pinning one while moving the other would describe a future in which the technology matured on one axis but not the other. Moving the efficiency also moves the cathodic hydrogen the electrified routes make for themselves, so that effect is carried inside the bundle rather than isolated from it. The narrative bundles are summarised in Table 3.7.

Table 3.7: Coherent-future narratives. Each column is a single decarbonisation pace; each row is a scenario dimension, given as its named tier with the value that tier takes. The columns are directionally consistent input bundles, not forecasts of which future will be realised.

Dimension	Decarbonisation lags	Reference	Decarbonisation accelerates
Feedstock cost	<i>FP-High</i> upper tier of every feedstock band	<i>FP-Central</i>	<i>FP-Low</i> lower tier of every feedstock band
H₂ procurement	<i>Grey</i> \$3.20/kg, 10 kgCO ₂ e/kg	<i>Blue</i> \$3.50/kg, 3 kgCO ₂ e/kg	<i>Green</i> \$4.40/kg (2030 STEPS), 1 kgCO ₂ e/kg
HNO₃ procurement	<i>Non-EU unabated</i> <i>EI-B</i> , 3.99 kgCO ₂ e/kg pure	<i>EU-abated</i> <i>EI-A</i> , 0.93 kgCO ₂ e/kg pure	<i>EU-abated</i> <i>EI-A</i> , 0.93 kgCO ₂ e/kg pure
Benzene procurement	<i>BEN-Fossil</i> \$1.65/kg, 1.87 kgCO ₂ e/kg	<i>BEN-Fossil</i> \$0.88/kg, 1.87 kgCO ₂ e/kg	<i>BEN-Bio</i> \$0.77/kg, 1.08 kgCO ₂ e/kg
Carbon price	<i>CP-Low</i> €75/tCO ₂ , flat	<i>CP-Central</i> €115/tCO ₂ , flat	<i>CIP (2026) Net Zero rising trajectory</i> [76] €140 to €210/tCO ₂ over 2030–2050
Free allocation α	<i>retained in full</i> $\alpha = 1$	<i>retained in full</i> $\alpha = 1$	<i>withdrawn</i> $\alpha = 0$
Grid carbon intensity	<i>static 2024 grid</i> 261 gCO ₂ e/kWh, held flat	<i>adjusted grid trajectory</i> 261 gCO ₂ e/kWh 2024 anchor, declining	<i>decarbonised supply</i> 31 gCO ₂ e/kWh 2024 anchor, declining
Cell performance	<i>EC-Conservative</i> FE 0.45/0.30 (ECO/ECH)	<i>EC-Central</i> FE 0.80/0.40	<i>EC-Optimistic</i> FE 0.90/0.60
Cell CAPEX	<i>pilot: +50 % CAPEX</i>	<i>central</i>	<i>matured: −30 % CAPEX</i>

The tier assignments in each column are indicative: each dimension sits at the value that best fits the storyline's direction rather than a forecast level, and the reasoning behind the main parameters is set out below.

One policy hinge carries two of the rows, which is what makes them coherent rather than merely co-varying. CBAM and free allocation are two ways of addressing the same carbon-leakage risk, so an adipic-acid plant loses the second only if it is brought within the first (LR §2.4.1). Nitric acid is already in CBAM scope, so what varies there is enforcement; adipic acid is outside it, so what varies here is extension. Under *lags* neither happens and the cheaper non-EU unabated nitric acid is the rational buy; under *Reference* enforcement arrives and EU-abated acid becomes rational, the allocation being retained at $\alpha = 1$ in both, the position that stands today held across the producing life; under *accelerates* extension follows and the allocation goes.

The two reach the routes differently. The allocation is benchmark-based, so it shifts all three by the same amount and moves the level rather than the ranking, whereas the nitric acid reaches the conventional route alone. Withdrawal under *accelerates* also arrives with a border charge on imported adipic acid that this model does not price, so $\alpha = 0$ is a cost to the European plant rather than a loss of standing against the producers it competes with.

Feedstock supplier choice tracks policy and pace. HNO₃ follows the CBAM hinge above (§A.5.3), with one deliberate conservatism: under *accelerates* the ammonia behind EU-abated acid would itself partly decarbonise, but with no credible matched trajectory its intensity is held at the 2014 value, under-crediting that narrative. Benzene follows the bio-feedstock assumption: fossil naphtha-cracker benzene stays rational under *lags* and *Reference* (BEN-Fossil; LR §2.4.6), while under *accelerates* bio-attributed BTX scales up with the parallel H₂ and electricity decarbonisation, making BEN-Bio rational (FP-Low, anchored to BioBTX's cost-parity claim [126]).

Carbon price and H₂ scale with ambition: CP-Low and Grey under *lags*, CP-Central and Blue (CCS at scale [127]) under *Reference*, and a CIP (2026) Net Zero rising trajectory [76] with Green H₂ under *accelerates*. Hydrogen is

carried on the 2024 production cost throughout except under *accelerates*, which takes the lower 2030 IEA STEPS projection [127]: a price decline is itself optimistic, so it is confined to the optimistic narrative (§A.5.2). Grid CI and feedstock cost co-vary with pace: *lags* holds the static 2024 grid and FP-High, while *accelerates* takes a grid anchored at France’s 31 gCO₂e/kWh, together with reduced fossil demand at FP-Low. France’s intensity is used as the decarbonised bound rather than as a fourth sited case: it is what an accelerated ARRRRA decarbonisation converges toward. The electricity price is held at the relieved extra-high-voltage tariff of €62/MWh in all three futures, so that grid carbon intensity varies on its own rather than being coupled to a country’s price. The cells’ capital cost matures on the same axis, pilot-stage under *lags*, central under *Reference* and matured under *accelerates*, as wider deployment funds cheaper cells (§3.5.3).

The bundles address Gap (viii). Their inputs are anchored to published scenarios (IEA STEPS, CIP Net Zero, REF2020), but they are directionally consistent shifts rather than predictions, and the assumed co-variation is itself an analytical choice; alternative correlations are conceivable, for example accelerated decarbonisation compressing carbon prices if it outpaces ETS cap reductions.

3.5.5. Load-following under dynamic electricity prices

The central case runs every route as continuous baseload at 8 000 h/yr. Separate from the tier and trajectory dimensions above, a targeted one-dimensional sensitivity tests whether operating each route at a load factor below unity, following the hourly electricity price, improves the levelised cost of adipic acid or the scope-2 GHG intensity. Two strategies bound how much this helps, differing in whether annual output is kept. The *produce-less* baseline keeps the plant at its design size and runs it fewer hours: it makes less product, so the whole fixed cost spreads over a smaller output, and it needs no storage. The *oversize* bound instead builds the cell larger and buffers an intermediate so the plant holds its output while only the cell follows the price; it saves more, but during a long cheap spell the oversized cell out-produces the downstream train, so the intermediate accumulates and holding output flat can demand storage of days to weeks. That buffer is left uncoded, so this is an optimistic upper bound. Produce-less is taken as the baseline because it needs no such storage; how the oversize case differs from it, and how its buffer is sized, are in §E.5. Reporting both bounds together shows the order of magnitude of what load-following could deliver on the two indicators for an electrified plant under grid connection, rather than a single figure. The realistic operating point (a modestly oversized cell with a small, coded buffer that rides through short price spikes) falls somewhere between them, at a position this analysis does not resolve, so the two ends bracket the achievable benefit without a separate buffer-sizing case.

The sensitivity addresses grid-connected dynamic pricing rather than on-site renewable coupling; the case for grid-connected baseload as the representative supply at industrial scale is made in LR §2.4.3. The supporting detail is collected in Appendix E: the five-step optimisation that sets the threshold, the per-country decomposition of the electricity price, the oversize bound with its buffer sizing, and the cost of cycling the cell.

Data

Electricity is priced from the observed 2024 hourly day-ahead prices for the Netherlands, Germany, and Belgium (ENTSO-E via Ember [93]; annual averages €77.3, €78.5, €70.3/MWh), which set both the commodity level and its hourly shape. Non-commodity grid charges are from E-Bridge [100] (Table A.1) and split into a flat per-MWh volumetric part and a per-kW capacity part. The volumetric part is the same cost per kilogram at any load factor, while the capacity charge, like the plant’s other fixed costs, is unchanged in total and so rises per kilogram as the plant runs fewer hours. Spread over a full 8 000-hour year that charge is worth €17.96/MWh in the Netherlands (unrelieved TenneT contracted-capacity and peak charges), €1.99 in Germany (heavily relieved under §19 StromNEV) and €1.39 in Belgium, and being fixed in total it is the main source of the country differences below. Each hour also carries its own grid carbon intensity, from the 2024 hourly Electricity Maps series [94] (lifecycle, consumption-based), so the scope-2 GHG of the flexible load follows the hours actually run. The conventional route is thermal, with no shiftable electrochemical load, so it is excluded; ECO and BIO_{AA} are analysed. The flexible load is assumed to curtail and restore without a ramp or part-load penalty.

Method

The plant follows a single price threshold p_{low} : in each hour it runs at full load if the price is at or below p_{low} and shuts otherwise (run-or-shut). The threshold and the load factor LF (the fraction of the year’s hours priced below it) are two views of the same decision, since fixing either fixes the other through the sorted price distribution. That decision is set by the optimisation below, so the load factor is a result of the analysis, not a value chosen in

advance. Run-or-shut counts only the cell's membrane wear from cycling and leaves the plant's own shut-and-restart costs out of scope, so it treats whole-plant cycling as costless. It is therefore the optimistic dispatch: it stops the whole plant for any hour above p_{low} , even an isolated one, charging nothing for the shut and the restart. A real operator would weigh that cost against the size of the spike, and the two do not always point the same way: an ordinary hour above the threshold would not be worth stopping for, an extreme one would. Where the balance falls is not resolved here, and to the extent that short spikes are ridden through rather than curtailed, fewer hours are dropped and less is saved than reported here.


The threshold has no closed form and is found numerically: for each load factor on a grid from 0.50 to 1.00 the plant runs the cheapest hours of the operating year, the resulting LCOA change is computed against flat baseload, and the load factor giving the largest reduction fixes p_{low} . That change is the sum of three terms, two penalties and one saving. The whole-plant fixed cost and the per-kW capacity charge are both unchanged in total but spread over less product, so both rise per kilogram as $(1/LF - 1)$; against them the hours actually run are cheaper than the annual average, a saving that deepens as fewer and cheaper hours are kept. The optimum balances the two. The 8 000-hour design basis leaves 784 of 2024's hours as a maintenance window, placed on the contiguous block whose mean price is closest to the annual mean, so the operating year is representative rather than optimistic. The full procedure, its five steps and the expressions for the LCOA and scope-2 GHG changes are in §E.2.

The search is repeated per country and route. The dispatch assumes perfect foresight of the day-ahead prices and so bounds the achievable saving from above. Each shut and restart of the ECO cell adds a small increment of membrane degradation that brings the stack replacement forward, so cycling is priced as a cost per start-stop, $C_{\text{cycle}} \approx \64 (€58), built from the measured degradation rate and the replacement cost (§E.6) and charged on every run-or-shut cycle; the biomass routes carry none (their cell corrodes at the cathode, not voltage rise), and the burden is small and reported with the results (§4.6). Labour, logistics, and other operating costs or revenues of load-following are outside this analysis (Chapter 6); the resulting load factors and LCOA and scope-2 GHG changes are reported in §4.6.

3.5.6. Evaluation modes

The framework supports six evaluation modes, ordered from broad overview to multi-dimensional exploration, that together map under which conditions each route becomes competitive rather than delivering a single verdict. In each mode, dimensions not actively varied are held at their Central values unless stated otherwise. Cell-performance degradation is applied in every mode as a within-lifetime mechanism (specific energy consumption, SEC, rising at the imposed rate, periodic stack replacement); the uncertainty in initial SEC, Faradaic efficiency, and degradation rate is captured by the cell-performance tier dimension swept in Mode 3.

1. **Headline comparison.** Produces top-level LCOA and GHG maps across routes and time. All five tier dimensions at Central, H_2 at Blue on the 2024 production cost [128], HNO_3 at EU-abated, fossil-derived benzene, and free allocation at the headline $\alpha = 1$ with the $\alpha = 0$ bound reported alongside it.
2. **Country.** Isolates the location effect on LCOA and GHG. Routes are compared across BE, NL, DE, with France as a renewable-grid reference (Section 1.3), under Central tier dimensions; country profiles couple each country's electricity-price tier with its grid-intensity tier.
3. **Single-dimension sensitivity.** Produces tornado charts: six dimensions swept across Low/Central/High, including feedstock and process-chemical prices, plus supplementary CAPEX sensitivity. Procurement scenarios are evaluated separately as paired-input switches, and the load-following operating-strategy sensitivity (§3.5.5) is reported here as a further one-at-a-time result.
4. **Time-dependent sensitivity.** Tests whether within-lifetime evolution changes the picture beyond Mode 3's static tier reading. Two cases (Section 3.5.1): per-country grid-CI trajectory and CIP (2026) Net Zero rising CO_2 -price trajectory [76], both evaluated against their static-tier counterparts.
5. **Pairwise interaction.** Tests whether two dimensions reinforce or offset each other; produces 3-tier $\times$ 3-tier heatmaps revealing interaction effects invisible in single-parameter sweeps.
6. **Coherent-future narratives.** Three named futures (*Decarbonisation lags*, *Reference*, *Decarbonisation accelerates*; §3.5.4), each pinning every dimension to a single decarbonisation pace and carrying the grid and carbon-price trajectories with it.



Results

This chapter reports the levelised cost of adipic acid and the cradle-to-gate GHG intensity for the three routes, with the biomass route evaluated at both its *t*3HDA and adipic-acid end points, across the evaluation space of Chapter 3.

The headline comparison (§4.1) sets each route against Yeo (2025)'s published values, so the impact of this work's corrections and extensions reads as the change from his basis; that this work's levelised-cost model first reproduces his costs on his own basis (§3.2.7) is what licenses that reading. Sections 4.3–4.6 then probe the conditions under which each route is competitive, locating each route's dominant bottleneck and where its viability turns. The research questions those results answer are stated in §1.4 and are taken in order in the concluding chapter (Chapter 5).

4.1. Headline route comparison

The research questions concern the impact of the input dimensions and of the gaps this work fills, so the comparison to the inherited model is the natural starting point: it places the results of this work against the reference and isolates what the extensions changed. Table 4.1 sets the central-case LCOA and GHG intensity against Yeo (2025), on the common basis of design set out in Chapter 3 (Table 3.1) and at the central tier of the scenario framework. The two sides are not on quite the same accounting: this work differs from Yeo in cost scope, in corrections to his cost data and in input values alike. Chapter 3 sets out those differences and §C.1 lists them input by input; their net effect is resolved in Table 4.2.

Both comparisons are close enough that the ordering should be read as conditional rather than settled. The conventional and electrochemical routes differ by 17 % on price and 12 % on carbon, and several of the input choices behind those figures shift them by a magnitude similar to the difference between the routes itself: the capital estimate on price, the nitric-acid supplier and the country grid on carbon. The N₂O abatement level reaches both, since the same abated figure sets the conventional route's scope-1 emissions and its EU ETS charge. §4.3 and §4.5 set out how far each ordering survives when they vary, so the rankings below are illustrative of the central inputs rather than fixed properties of the routes. One column needs a further caution: at the *t*3HDA endpoint the biomass route sells an intermediate rather than adipic acid, so its levelised cost is not directly price-comparable with the three adipic-acid figures, and the percentage against the conventional route should be read with that in mind (§4.2.7).

4.1.1. Levelised cost of adipic acid

The ordering Yeo reported holds; the distances between the routes do not (Table 4.1). The corrections this work makes land almost entirely on the conventional route: it rises 18 % against Yeo's figure, while the electrochemical route falls 6 % and the biomass route moves 1 %. The premiums they carried over it therefore fall, from 47 to 17 % on the electrochemical route and from 151 to 115 % on the biomass one (§2.3.1). The electrochemical route closes most of its inherited cost disadvantage; the biomass route closes little of its own. Its largest single input,

Table 4.1: Central-case results against Yeo (2025), on the heat-integrated basis with central-tier cell degradation. LCOA is on each route's product basis: adipic acid for CON, ECO and BIO_{AA}, and t3HDA for BIO_{t3HDA}. Yeo does not model BIO_{AA}, hence the dashes. Bracketed percentages are each route's difference from the conventional route within the same column. The component build-up behind every figure in this table is given in Table B.1.

Route	LCOA (\$/kg)		GHG (kgCO ₂ e/kg)	
	This work	Yeo	This work	Yeo
CON	1.87	1.58	3.39	4.79
ECO	2.19 (+17 %)	2.33 (+47 %)	3.79 (+12 %)	2.20 (−54 %)
BIO _{t3HDA}	4.03 (+115 %)	3.97 (+151 %)	2.34 (−31 %)	1.18 (−75 %)
BIO _{AA}	4.34 (+132 %)	—	2.56 (−24 %)	—

the muconic-acid broth, is not affected by the gaps addressed in this work: on its own it contributes 2.08 \$/kg of adipic acid, more than the conventional route's entire levelised cost of adipic acid. That price has no market to read it from and is modelled from a published selling price for muconic acid, so the biomass premium reported throughout this chapter is conditional on it (§4.4).

The conventional figure sits just below the 2025 EU adipic-acid market band (\$1.88–2.04/kg, the full-2025 IMARC EU range [15], Figure 4.1). The two are different quantities: a levelised cost of adipic acid is a break-even for a new build, the price at which it recovers its capital, offsites and engineering over the project life while carrying a carbon cost, whereas a market price is what existing capacity currently clears at.

These are central-case figures on the heat-integrated basis (§4.2.2) with central-tier cell degradation (§4.2.5). Of the gaps this work closes, three move the central price materially: the cost-scope and capital corrections (§4.2.1), the move to extra-high-voltage industrial electricity prices net of the relief large consumers qualify for, which §4.3 then varies by country, and the EU ETS charge (§4.2.6); §4.2 resolves all of them component by component. Heat integration, cell degradation and the cathodic-hydrogen credit each move it by less than 0.04 \$/kg on every route: they are carbon levers rather than cost levers.




Figure 4.1: Levelised cost of adipic acid decomposed by cost component, this work vs Yeo (2025), per route. Both bars are the same quantity, the price giving NPV = 0 over the plant life; Yeo reports it under the name minimum selling price (§3.3.2). The grey band is the 2025 EU adipic-acid market range.

4.1.2. GHG intensity

On carbon the inherited ranking does not survive (Table 4.1). The conventional route falls 29 % while the electrified routes rise, by 72 % on the electrochemical route and 98 % on the biomass one, so the electrochemical route crosses from 54 % below conventional to 12 % above and the biomass advantage narrows from 75 to 31 % below. Two of Yeo's assumed emission factors are what close that gap and then reverse its sign, and they sit on

opposite sides of the values measured here. He carried nitric acid at the industry-average factor of 3.9 kgCO₂e/kg against 0.93 for an abated EU supplier, which pulls the conventional route down, and electricity at an assumed 50 gCO₂e/kWh against the 261 the European grid actually delivers, which pushes the electrified routes up. Two smaller terms partly cancel on the conventional route: heat integration lowers it by a further 0.49 kgCO₂e/kg (§4.2.2), while charging it for the N₂O that survives 99 % abatement, rather than assuming complete destruction, adds 0.82 back. §4.2 resolves both bridges component by component.

Published intensities for conventional adipic acid are far higher, spanning 6.3 to 25.6 kgCO₂e/kg (LR §2.2.1). The difference is not one of boundary: the benzene term here, 1.35 kgCO₂e/kg, matches the 1.33 Corona et al. carry for cyclohexane. It is two inputs and an energy balance. Nitric acid contributes 0.80 against their 4.63, because the supplier is EU-abated rather than unabated, and residual N₂O 0.82 against 3.58, because destruction is 99 % rather than 97 % [17]; those two are about seven-tenths of the difference, and most of the remainder is process energy they charge more heavily than a heat-integrated plant on a European grid. The route modelled here is therefore lower-emitting for what it buys and how much heat it recovers rather than for its chemistry, and §4.5 shows how quickly that position erodes when any of the three is relaxed.

The sharpest departure from Yeo's ranking is between the conventional and electrochemical routes, and the scope split shows where it comes from (Figure 4.2). Against the conventional route, electrification does what is expected on the two scopes it targets: scope-1 falls from 0.89 to 0.09 kgCO₂e/kg with the N₂O, and scope-3 from 2.32 to 1.38 with most of the nitric acid, a combined saving of 1.74. Scope 2 then rises from 0.18 to 2.32 and more than reverses both. The route does not trade a concentrated N₂O emission for a smaller one elsewhere; it trades two moderate emissions for one larger one, at an exchange rate the grid sets. On the Dutch grid intensity carried in the central case, a conventional plant that destroys its N₂O to modern levels, buys abated nitric acid and is systematically heat-integrated therefore emits less per kilogram of adipic acid than the electrochemical route built to replace it.




Figure 4.2: Cradle-to-gate GHG intensity decomposed by scope, this work vs Yeo (2025), per route.

4.1.3. Material efficiency

Alongside price and carbon the routes carry a third, diagnostic measure: how much of the carbon entering as feedstock leaves as product (§3.4). The conventional and electrochemical routes convert 0.741 and 0.774 of the theoretical maximum, the biomass route 0.928 and its adipic-acid endpoint 0.872. The first three reproduce Yeo's figures, since they run on his mass flows; only the adipic-acid endpoint is new, and it falls from 0.928 to 0.872 because the polishing reactor and its recovery train yield about 6 % less adipic acid than the *t*3HDA they process could give (§4.2.7).

4.2. Corrections and extensions to the reference model

This section quantifies what each gap this work addresses is worth on the headline LCOA and GHG intensity of §4.1, separately from the corrections to Yeo's inherited cost data, which are not gaps but have to be accounted for before the extensions can be read cleanly. Table 4.2 resolves the difference between his baseline and this work component by component; several components work against one another, so a small net change on a route can hide substantial movement underneath. Every comparison to Yeo in this section stops the biomass route at *t*3HDA, since that is the product his model carries; the polishing step to adipic acid has no counterpart in his model and is treated separately in §4.2.7. One extension is left out of the bridge entirely: Yeo buys each purchased feedstock at a single generic price and upstream intensity, whereas this work treats benzene, hydrogen and nitric acid as procurement choices with their own prices and supply chains. That has no single central-case value to add to the bridge, since the central case picks one supplier per stream; what it is worth is the spread between suppliers, and that is reported in §4.4.

Table 4.2: From Yeo's MSP to this work's LCOA: difference by component (\$/kg, central case). A positive entry raises the price and the rows sum to the difference between the two lines, to within 0.01 from rounding. The starting row is Yeo's method run on this framework, so that each component difference below is a clean comparison on a common basis (§3.2.7).

Component	CON	ECO	BIO _t 3HDA
Yeo (2025) MSP, his basis	1.58	2.33	3.97
Capital ¹	+0.13	+0.07	+0.05
Fixed OPEX	+0.33	+0.39	+0.56
Feedstock	+0.03	0.00	+0.15
Utilities ¹	−0.14	−0.43	−0.56
Waste disposal	0.00	0.00	+0.02
EU ETS	−0.06	−0.17	−0.17
This work LCOA	1.87	2.19	4.03

¹ These two rows carry the cell-replacement outlays and, on the electrochemical route, the degradation energy that Yeo's static basis omits: about +0.01 \$/kg in each on the electrochemical route, and +0.01–0.02 \$/kg in Capital alone on the biomass routes, whose cathode corrodes without an energy channel.

The same difference on the carbon side resolves the other way, and it does not decompose into the same terms. Table 4.3 closes it. On the conventional route the movement is dominated by one input the routes never share: nitric acid bought from an EU-abated supplier rather than a generic one is worth −2.19 kgCO₂e/kg on its own, more than the whole net change, and it is partly given back by charging for the N₂O that survives 99 % abatement rather than assuming complete destruction. On the two electrified routes the movement is almost entirely one term in the other direction: a grid at 261 gCO₂e/kWh instead of the 50 Yeo assumes.

Table 4.3: From Yeo (2025) to this work: GHG intensity difference by component (kgCO₂e/kg, central case). A positive entry raises the intensity. The starting row is Yeo's inventory run on this framework (§3.2.7), so each component below is a clean comparison on a common basis; it differs slightly from the 4.79, 2.20 and 1.18 Yeo prints, which Table 4.1 carries as published. Rows sum to the difference between the two totals, to within 0.01 from rounding.

Component	CON	ECO	BIO _t 3HDA
Yeo (2025) inventory, this framework	4.76	2.18	1.22
Nitric acid supplier	−2.19	—	—
N ₂ O abatement level	+0.82	—	—
Benzene emission factor	+0.48	+0.46	—
Grid carbon intensity	+0.13	+1.07	+1.20
Hydrogen supply	−0.13	−0.01	—
Cell degradation	—	+0.03	—
Heat integration (steam)	−0.49	−0.09	−0.16
Cooling water	—	+0.07	+0.03
Chilled water	—	+0.04	+0.04
Process chemicals	—	+0.04	+0.02
This work GHG	3.39	3.79	2.34

4.2.1. Cost scope and capital corrections

Capital and fixed operating cost are the two rows of the bridge this section takes apart (Table 4.2). Fixed capital ends 152 % above Yeo's total capital on the conventional route, 46 % on the electrochemical and 36 % on the biomass one. Those are large movements on a sum spent once, and levelised across the plant's whole output they reach the price as \$0.13, \$0.07 and \$0.05/kg. Fixed operating cost is the larger of the two on price, at \$0.33, \$0.39

and \$0.56/kg, because it is incurred again in every producing year. The difference against Yeo is therefore less than that the plant costs more to build than that it costs more to run and to account for.

Of the movements behind the two rows, most are conventions and two are corrections. The factorial capital build-up, the fixed-cost lines it carries and the working-capital treatment that goes with it are differences of method rather than of the plant (§3.3.4). The other two are equipment Yeo mis-costed or left out: his compressor correlations sit an order of magnitude below their published values and are restored, which adds \$31 M of installed equipment to the conventional route and \$19 M to the electrochemical one, and the N₂O abatement facility, modelled but left uncoded, adds \$25 M to the conventional route alone. Both fall hardest on the conventional route, which is why it is the route that rises furthest against Yeo.

Cash flow accounts for the rest. Yeo assumes full output from year one; this work discounts a two-year construction and a two-year production ramp, worth \$0.06–0.14/kg. The source-by-source decomposition of both the capital and the fixed-operating increase, and the capital composition by build-up line, are in §C.3.

4.2.2. Heat integration

A heat-exchange network was designed for each route in Aspen Energy Analyzer under the rules of §3.2.4, and both its utility loads and its exchanger area were costed through the same engine as everything else. Against the inherited flowsheets the redesign moves LCOA by about one per cent or less on every route, but cuts the conventional route's GHG intensity by an eighth; the two are taken in turn below. Both are read at the central case, Belgium's extra-high-voltage price of 62 €/MWh with the Dutch grid intensity of 261 gCO₂e/kWh, and with the cathodic-hydrogen steam credit held out throughout so that the comparison isolates heat integration.

Recovery already in the inherited flowsheets

Yeo's electrochemical and biomass flowsheets already recover, through the exchangers they carry, most of what the network synthesised here recovers; his conventional flowsheet does not (Table 4.4). The network grids are in §D.2 and the area and installed cost of each design in §D.4, which also lists every exchanger of every route.

Everything below therefore measures the increment this work adds to Yeo's flowsheets, not the worth of heat integration itself. The network synthesised here recovers 112.6 MW on the electrochemical route against no recovery at all; 102.7 of that was already in Yeo's flowsheet and the redesign adds the remaining 9.9, which is the 91 and 9 % of Table 4.4. The biomass routes are the same shape. The conventional route is the opposite case: Yeo captures only 1.7 of the equivalent 35.5 MW, so almost all of what heat integration can do for that route was still uncaptured.

Table 4.4: Utility duties before and after integration, per route (MW). *No recovery* is the counterfactual in which every cold stream is served by utility; *Inherited* is Yeo's flowsheet run through this work's engine; *HEN* is the network synthesised here at $\Delta T_{\min} = 10^\circ\text{C}$. The two recovery columns split the heat that network recovers, *no recovery* – *HEN*, between the two designs and therefore sum to 100 %; a different approach temperature would move that denominator. There is no equivalent counterfactual on the cooling side, and chilled water has no inherited column because that cooling water serves every stream, whereas the 25–30 °C specification used here (§3.2.4) leaves the coldest streams to chilled water. The hot-side duties of §D.3 exceed the cooling shown here on the two routes that raise steam on site; that reconciliation is given there.

Route	Heating (MW)			Recovery captured		Cooling water (MW)		Chilled (MW)
	no rec.	inherited	HEN	inherited	added here	inherited	HEN	
CON	38.9	37.2	3.4	5 %	95 %	77.5	42.0	6.2
ECO	167.2	64.5	54.6	91 %	9 %	143.6	129.6	9.5
BIO _{E3HDA}	35.2	12.0	10.2	93 %	7 %	65.8	47.9	9.1
BIO _{AA}	37.4	14.2	10.3	86 %	14 %	71.8	50.6	10.5

Effect on LCOA

The network moves LCOA by about one per cent or less on every route (Table 4.5), and the reason is a near-cancellation rather than an absence of savings: on the conventional route the utility saving is the largest of the four, but the exchangers that deliver it cost almost as much to carry. On the routes Yeo had already integrated there is little left to win, and the electrochemical route is the one case where the sign reverses, because its network reaches near-complete recovery through a larger number of small matches, which cost more per unit area than the fewer large ones Yeo's design needed.

That reversal follows from the design rule rather than from the costing. The networks are synthesised for maximum recovery at a common minimum approach temperature of 10 °C so that all four sit on one basis, rather than being

cost-optimised route by route (§3.2.4). On a flowsheet Yeo had already integrated well, driving recovery to that target buys negligible utility gain for real exchanger area. Heat integration at this scale is therefore close to break-even on LCOA. Measured against a flowsheet with no integration at all, the network adds \$38.5 M of fixed capital on the conventional route and \$6.2 M on the electrochemical one, and removes \$15.9 M and \$9.7 M on the two biomass routes, whose designed networks cost less to build than the exchangers their inherited flowsheets already carried.¹ The dominant uncertainty is the exchanger type: the networks are costed floating-head, at \$27.0, \$39.8, \$12.9 and \$16.2 M of installed exchangers, and costing them as U-tube instead lowers each by about 17 % (§D.4), a swing larger than the net LCOA effect itself on three of the four routes, so the sign of that effect should not be over-read. It does not touch the GHG result, which depends on the utility duties rather than on how the exchangers are costed.

Table 4.5: Effect of the designed network on LCOA, per route (\$/kg). *Inherited* is Yeo’s flowsheet run through this work’s engine, not his published figures; *HEN* replaces both its utility loads and its exchanger network with the Aspen Energy Analyzer design. On the two biomass routes the cathodic-hydrogen steam credit is held out, so that the columns isolate the network; the conventional route evolves no cathodic hydrogen and the electrochemical route runs a deficit at the central tier, so neither is affected. *Utilities* and *exchangers* sum to the *net* to within 0.001 from rounding; *utilities* carries the same turnover-linked markup as every other cost line, and *exchangers* carries the annualised capital together with the fixed operating costs it induces, which is why it exceeds a plain annualisation of the capital difference. Costing basis and the U-tube sensitivity are in §3.2.4 and §D.4.

Route	LCOA level		Change on integration		
	Inherited	HEN	utilities	exchangers	net
CON	1.890	1.869	−0.111	+0.090	−0.021
ECO	2.178	2.189	−0.013	+0.023	+0.011
BIO _{t3HDA}	4.067	4.033	−0.004	−0.030	−0.034
BIO _{AA}	4.367	4.343	−0.009	−0.015	−0.024

Effect on GHG intensity

The network is synthesised to recover as much heat as possible, and on carbon the whole of that recovery counts: exchangers cost money but do not emit, so nothing offsets the utility saving. On the conventional route it removes 13 % of the GHG intensity the route carries with Yeo’s network in place, against under 2 % on the electrochemical and biomass routes (Table 4.6).

Table 4.6: Effect of the designed network on GHG intensity, per route (kgCO₂e/kg product), on the same basis as Table 4.5. *Steam* is the purchased steam the network displaces, at 0.202 kgCO₂e/kWh (§A.6) and net of what each route raises on site; *CW pumping* is the reduced circulation and cooling-tower-fan draw; *chilled water* is the refrigeration the 25–30 °C cooling-water specification newly requires. A negative change lowers the GHG intensity. The three sum to the *net* on the conventional and electrochemical routes; on the biomass routes a residual below 0.012 kgCO₂e/kg sits in smaller utility lines.

Route	GHG level		Change on integration			
	Inherited	HEN	steam	CW pumping	chilled water	net
CON	3.881	3.391	−0.491	−0.026	+0.027	−0.490
ECO	3.849	3.791	−0.087	−0.010	+0.039	−0.058
BIO _{t3HDA}	2.485	2.469	−0.031	−0.013	+0.039	−0.016
BIO _{AA}	2.648	2.607	−0.067	−0.016	+0.046	−0.041

Two quantities set the size. The benefit is the purchased steam the network displaces, which scales with how little recovery the inherited flowsheet already had. The penalty is the refrigeration electricity that appears when cooling water is specified at 25–30 °C rather than Yeo’s 10–15 °C, which scales with how much sub-ambient duty the route carries. The net therefore differs sharply by route (Table 4.6): the penalty is 5 % of the steam saving on the conventional route, 45 % on the electrochemical one, and 126 % on the biomass route to *t3HDA*, where heat integration would raise the GHG intensity were the smaller cooling-water pumping duty not offsetting it. On Yeo’s 10–15 °C cooling water the penalty would not arise at all, and the electrochemical route’s net benefit would be −0.097 rather than −0.058 kgCO₂e/kg.

The steam saving is part of why the conventional route’s GHG intensity comes out below the electrochemical route’s (Table 4.1): integrating a flowsheet that carried almost no recovery is worth far more in carbon than refining one that already did.

¹On the biomass routes the two sides are costed differently: Yeo’s exchangers enter as one installed-cost figure from his Aspen model (\$20.3 M for BIO), the network here shell by shell from its area. His area is never measured, so this is a cost result, not by definition a smaller network.

4.2.3. Cathodic hydrogen and the by-product credit

At the central cell both biomass routes evolve more hydrogen at the ECH cathode than they consume. Yeo credits that surplus against disposal at \$0.70/kg; this work burns it on site to raise steam instead. The two are alternatives rather than additions: hydrogen that is burned is not disposed of, so his credit is withdrawn where the steam credit is taken.

The switch costs almost nothing on price and buys a real emission reduction. Yeo's disposal credit is worth 0.016 \$/kg of product; the steam the burned hydrogen displaces is worth 0.026 \$/kg gross, of which the boiler absorbs most, leaving 0.006 \$/kg net, so changing treatment costs 0.010 \$/kg, under a quarter of a per cent of the route's price. On carbon the two are not equivalent at all: his credit carries a zero emission factor and never moves scope-2, whereas burning the hydrogen displaces purchased steam and yields only water, worth 0.127 kgCO₂e/kg on the route to *t*3HDA. On the adipic-acid basis the routes are compared on it is the smaller credit that applies, 0.045, because the polishing reactor takes most of the hydrogen (§4.2.7).

Whether there is a surplus to burn at all is set by the cell rather than the flowsheet (§3.2.5). The route to *t*3HDA has one at every tier; the route to adipic acid is self-sufficient at the central and conservative tiers but turns hydrogen-short at the optimistic one, where it buys about 60 kg/h of makeup and forfeits the steam credit altogether (§4.5.3). The electrochemical route evolves cathodic hydrogen too, and this work credits it as an avoided purchase against that route's own benzene-hydrogenation demand rather than burning it.

The per-route split of both treatments, the valuation behind these figures, the bound on what merchant sale would be worth instead, and the waste-line reconciliation are given in §B.3.

4.2.4. Electricity price and grid intensity

Electricity enters through the duty each route draws and the price and carbon intensity attached to it. This work refines the duties only slightly: the biomass route's 5.70 kWh/kg is Yeo's own Aspen figure of 5.75 less the 536 kW of hydrogen export compression, a machine the flowsheet no longer carries once the surplus is burned on site rather than compressed for sale (§4.2.3), and the electrochemical route's corrected 5.05 kWh/kg restores the two feed compressors missing from his reported 4.79 (§3.2.7), worth +0.017 \$/kg. What those duties are valued at is the difference that matters.

This work prices electricity, and its carbon, from measured ARRR data rather than from a single assumed value, with the headline at the central scenario of each: Belgium's relieved extra-high-voltage tariff of 62 €/MWh, or \$0.0682/kWh (*EP-Central*), and the Dutch lifecycle grid intensity of 261 gCO₂e/kWh (*GI-Central*). Yeo carried a single industrial price of \$0.16/kWh, without the levy relief large ARRR consumers qualify for, and a renewable supply assumption of 50 gCO₂e/kWh [9]. The effect falls almost entirely on the electrified routes and runs in opposite directions: against his reported figures the tariff lowers LCOA by 23 and 15 % on the electrochemical and biomass routes, while the grid intensity adds scope-2 emissions worth 54 and 107 % of their totals (Table 4.7). Price and grid intensity coupled country by country are in §4.3.

Table 4.7: Electricity as a cost and a carbon effect: each route's corrected duty (§3.2.7) valued on Yeo's basis and on the central case of this work. Yeo's basis is \$0.16/kWh and 50 gCO₂e/kWh [9]; the central case pairs Belgium's relieved extra-high-voltage tariff, 62 €/MWh or \$0.0682/kWh [100], with the Dutch lifecycle grid intensity of 261 gCO₂e/kWh [94]; these are the named central scenarios of the two sets, *EP-Central* and *GI-Central*, which fall in different countries. Two duties are given because they differ: the priced duty is the electricity the plant buys, chilled water included, while the duty carrying grid intensity also includes cooling-water pumping, which is charged at the cooling-water price but drawn from the grid. Each effect row is the model's own difference between the two bases rather than the duty times the price step, since the cost lines charged on turnover move with the electricity price as well.

	CON	ECO	BIO _{t3HDA}	BIO _{AA}
Priced duty (kWh/kg)	0.514	5.323	5.851	6.190
Price effect (\$/kg)	-0.051	-0.529	-0.581	-0.615
Duty carrying grid intensity (kWh/kg)	0.631	5.669	5.983	6.336
Scope-2 effect (kgCO₂e/kg)	+0.133	+1.196	+1.262	+1.337

4.2.5. Cell degradation

Neither cell holds its initial performance over a twenty-year life, and the central case carries the degradation model of §3.2.6 rather than initial-performance figures. The two cells degrade differently, and the difference shows in which KPI moves. The electro-oxidation cell degrades by voltage rise, so its specific energy climbs over the life and its membranes are replaced periodically: it costs 0.018 \$/kg and, because the extra energy is drawn from the

grid, also 0.034 kgCO₂e/kg. The electrochemical-hydrogenation cell degrades by cathode corrosion, which is a replacement cost with no energy channel, so it costs 0.015 \$/kg on both biomass end points and leaves the carbon intensity untouched. The conventional route carries no cell and is unaffected. These are central-tier rates; the tier itself, and the cost of the cell as a maturing technology, are swept in §4.5, which is where the uncertainty in them is tested.

4.2.6. EU ETS charge and abatement

This work prices each route's scope-1 emissions under the EU ETS, which no published route comparison does (§3.3). The headline results take a plant first producing in 2030, against the adipic-acid benchmark of 1.40 tCO₂e/t, full free allocation at $\alpha = 1$ (§3.3.2), and a carbon price of 115 €/tCO₂.

Scope-1 emissions are in line with Yeo's on the electrified routes, 0.090 tCO₂e/t on the electrochemical route and 0.034 on the biomass one. They diverge only on the conventional route, where he assumes the N₂O is destroyed completely and reports 0.068. The central case here takes 99 % instead, the modern fleet-median level and closer to what operating installations actually report (LR §2.2.1). The N₂O surviving that leaves 0.819 tCO₂e/t on top of 0.073 of process CO₂, lifting conventional scope-1 to 0.892.

Whether a route pays or earns depends on how its scope-1 emissions sit against the allocation, and under Article 10b(1) that allocation stands at the full benchmark for the whole producing life rather than declining through it (LR §2.4.1). All three routes sit below the benchmark, so all three are net sellers: levelised over the producing years the conventional route earns 0.064 \$/kg, the electrochemical route 0.166 and the biomass routes 0.173. The electrified routes earn most, because their scope-1 is almost nothing and very nearly the whole benchmark is surplus.

The abatement assumption is what puts the conventional route on that side of the line, and it does so narrowly. The allocation is exhausted at 98.38 % destruction, where its scope-1 reaches the benchmark, and below that the route pays rather than earns. That break-even does not move with the carbon price, because both sides of it scale with the price. Measured against the EU fleet it is clear rather than marginal: the production-weighted mean of 1.14 tCO₂e/t (LR Table 2.1) implies about 98.7 % and still sits above it, so an average European plant would also earn, by roughly half as much, taking its levelised cost of adipic acid to 1.90 and narrowing the gap to the electrochemical route from 0.32 to 0.29 \$/kg. On complete destruction the route would earn 0.168 \$/kg, level with the electrified ones, so the whole of its EU ETS position turns on the last one and a half per cent of N₂O. The carbon price, the abatement level and the free-allocation assumption are all swept in §4.5.

4.2.7. Product boundary: *t*3HDA versus adipic acid

The polishing step hydrogenates *t*3HDA to adipic acid over a Pd/AC catalyst and a second crystallisation train recovers the product. Two consequences follow, and both push the same way. The first is mass: the reactor and its recovery train turn 12,350 kg/h of *t*3HDA into 11,757 kg/h of adipic acid, a net loss of about 5 %, so every per-kilogram cost the route already carried is now spread over less product. That falls hardest on the muconic-acid broth, because the broth dominates the route's cost base to begin with.

The step also consumes a by-product the route to *t*3HDA was free to use. Hydrogen evolved at the ECH cathode is burned to raise steam when the route stops at *t*3HDA, and polishing diverts most of it to the reactor (§4.2.3). A credit is therefore withdrawn at the same moment equipment is added, and on carbon that withdrawal is the largest single term in the increment.

Both endpoints are in Table 4.1; between them the step adds \$0.32/kg and 0.22 kgCO₂e/kg. Both figures are conditional on the polishing step performing as assumed: they price the consequence of a successful *t*3HDA to adipic acid conversion rather than evidence that the step has been demonstrated at this scale (LR §2.2.3). On price the increment is spread across the build-up rather than concentrated in one line: roughly a third is the mass loss reaching the feedstock, about half is the polishing reactor and crystallisation train through the capital and fixed operating cost they induce, and the rest is utilities. On carbon it is almost entirely scope-2, and the withdrawn steam credit is the largest part of that; the recovery train's own electricity and the route's grid emissions spread over less product make up the balance, and scope-3 carries only the feedstock effect. The full component build-up for every route is given in Appendix B. The route's standing against the conventional and electrochemical benchmarks is unchanged. The second crystallisation stage earns the capital it adds: it lifts adipic-acid recovery from 91.3 to 95.7 %, leaving about 531 kg/h of dissolved acid to leave in the mother-liquor purge.

4.3. Location: price and grid intensity

This section varies electricity price and grid carbon intensity together across realistic values and measures what that variation does to each route. The three ARRA countries supply the range from observed 2024 data, 49–92 €/MWh and 150–341 gCO₂e/kWh, with France added as a bonus scenario for a near-renewable grid; each country supplies its own price and its own grid, and every other input stays central (Table 4.8, Figure 4.3).²

Location enters the model through those two terms only. Wages, construction costs, material and utility prices, taxes other than the ETS, and transport distances are held at their central values, on the assumption that any candidate site sits inside the ARRA cluster itself, an established petrochemical complex with the jetty, pipeline, storage and utility connections these plants need already in place. A site outside such a cluster would carry logistics and connection costs this comparison does not capture; the claim is only that the residual differences between the three countries are second order next to a price range of 49–92 €/MWh and a carbon range of 150–341 gCO₂e/kWh.

Table 4.8: Each country's extra-high-voltage electricity price and lifecycle grid carbon intensity, and the resulting levelised cost of adipic acid and GHG intensity per route. Every other input is central, so electricity is the only term that changes between countries; France is the low-carbon reference case. Prices are E-Bridge's 2030 projected extra-high-voltage baseload levels with full relief applied, on which every route is priced [100]; intensities are 2024 lifecycle, consumption-based values from Electricity Maps [94].

Country	Elec. (€/MWh)	Grid CI (gCO ₂ e/kWh)	LCOA (\$/kg)			GHG (kgCO ₂ e/kg)		
			CON	ECO	BIO _{AA}	CON	ECO	BIO _{AA}
Belgium	62	150	1.87	2.19	4.34	3.32	3.16	1.86
Netherlands	92	261	1.89	2.38	4.56	3.39	3.79	2.56
Germany	49	341	1.86	2.11	4.25	3.44	4.24	3.07
France (reference)	48	31	1.86	2.10	4.24	3.25	2.49	1.11

Effect on LCOA

The conventional route barely moves on either measure, 1 % on price and 4 % on carbon, because it buys little grid electricity. The electrified routes move far more, and asymmetrically: the electrochemical route spans 12 % on price against 34 % on carbon, the biomass route 7 % against 65 %. Figure 4.3 shows the same pattern in its distribution shape: both the electrochemical and the biomass points spread vertically while staying in a relatively narrower band horizontally. Against the conventional route, the electrochemical premium falls from 26 % in the Netherlands to 17 % in Belgium and 13 % in Germany, but no country moves any route far enough to change the cost ranking. Neither does removing the relief. The EP-Statutory tier prices every route at the statutory band-IG tariffs instead, which carry no extra-high-voltage reduction and stand at 84, 110 and 130 €/MWh in Belgium, the Netherlands and Germany; on it the electrochemical route reaches 2.32, 2.49 and 2.62 \$/kg while the conventional route moves only to 1.88–1.91, and at Yeo's 145 €/MWh it reaches 1.91. The relief is therefore worth more to the electrified routes than to the conventional one, as their electricity duties imply; withdrawing it widens the premium rather than reordering the routes.

Effect on GHG intensity

On carbon the ranking does shift. In Germany the electrochemical route carries 23 % more GHG emissions than the conventional one and in the Netherlands 12 % more, while in Belgium it carries 5 % less. France, the bonus scenario, shows what a near-renewable grid is worth: there the electrochemical route falls 23 % below the conventional route. The two therefore cross, and on the central inputs they reach the same carbon intensity at 181 gCO₂e/kWh, both at 3.34 kgCO₂e/kg. Above that line electrification raises the carbon intensity of adipic acid; below it, lowers it. The same fit puts the biomass route level with the conventional one at 406 gCO₂e/kWh and 3.48 kgCO₂e/kg, above every grid considered, though only 19 % above Germany's 2024 intensity.

Grid decarbonisation over the plant life

In the headline results the grids above are held at their 2024 intensity for twenty years, while all four are projected to decline over that period. Letting each follow its projected national trajectory and averaging over the producing years of a plant first producing in 2030 gives the levelised intensities of Table 4.9. Every one of them sits below the 181 gCO₂e/kWh crossing, so on the projected trajectories the electrochemical route emits less than the conventional one in all four countries, on the central abatement assumption (§4.5.2). On the static 2024 values

²France is not an ARRA country and is carried as a reference case rather than a siting candidate (§3.1); at 31 gCO₂e/kWh it is close to the lower bound of any European grid.




Figure 4.3: Levelised cost of adipic acid against cradle-to-gate GHG intensity, one marker per route and country, with every input other than electricity price and grid carbon intensity at its central value. Colour is the route, marker the country; points to the lower left are both cheaper and cleaner.

Belgium was the only ARRRA country where it did. The 2024 basis carried through the rest of this chapter is in that sense the conservative one.

Because the trajectory makes grid intensity time-dependent, commissioning year becomes a lever in its own right: a plant first producing later averages a cleaner grid over the same twenty years. It is the only time-dependent input whose effect turns on the commissioning year in the central case. The free allocation does not, since adipic acid keeps the full benchmark rather than following a phase-out schedule (§4.5.1), so commissioning year moves the carbon and not the allocation. Moving first production from 2030 to 2035 lowers the electrochemical route's footprint by between 0.03 and 0.22 kgCO₂e/kg across the three ARRRA countries, most in Germany and least in Belgium, whose grid declines least of the three.

Table 4.9: GHG intensity on a declining grid: each country's 2024 anchor against the intensity levelised over the producing years of a plant first producing in 2030, and the resulting route emissions. Trajectories rebase the EU Reference Scenario 2020 PRIMES projections [122] onto each country's 2024 lifecycle anchor [94]; per-country values are given in §A.3.

Country	Grid CI (gCO ₂ e/kWh)		GHG (kgCO ₂ e/kg)		
	2024	levelised	CON	ECO	BIO _{AA}
Belgium	150	140	3.31	3.10	1.79
Netherlands	261	103	3.29	2.89	1.56
Germany	341	128	3.31	3.04	1.72
France (reference)	31	24	3.24	2.45	1.06

4.4. Procurement scenarios

Three of this study's feedstocks are not continuous ranges but procurement decisions a buyer makes: where the nitric acid is sourced from, and what kind of benzene and hydrogen is bought; the prices and scope-3 intensities used for every feedstock in this study are collected in Table 2.5. They are treated separately because they move a route's carbon without changing its design or its site, and so remain open to a plant already built.

Whether they can be pulled is a separate question this assessment does not answer. Each presumes a supply that can be contracted at the scale a 100 kt/a plant consumes, from a source whose emission profile is documented rather than assumed, and how binding that is differs by feedstock (§3.1). What follows is what each switch is worth if it can be made, and Table 4.10 puts each alternative supplier against the central case.

Nitric acid

Nitric acid is the single largest carbon lever in the study after the N₂O abatement level. The alternative is not simply an unabated supplier: the scenario is a non-EU producer that vents its process N₂O and makes its ammonia from

Table 4.10: Where each route lands when one purchased feedstock is swapped for an alternative supplier, everything else central. The central case buys EU-abated nitric acid on gas-based ammonia, fossil benzene and blue hydrogen; the bracketed figures are the percentage change relative to it. A dash marks a feedstock the route does not buy, and the biomass routes buy none of the three and so do not appear.

Supply	LCOA (\$/kg)		GHG (kgCO ₂ e/kg)	
	CON	ECO	CON	ECO
Central case	1.87	2.19	3.39	3.79
Non-EU HNO ₃ (unabated, coal NH ₃)	1.87 (0 %)	—	6.02 (+78 %)	—
Bio-derived benzene	2.56 (+37 %)	2.85 (+30 %)	2.82 (−17 %)	3.25 (−14 %)
Grey hydrogen	1.85 (−1.0 %)	2.19 (−0.0 %)	3.79 (+12 %)	3.81 (+0.5 %)
Green hydrogen	2.06 (+10 %)	2.20 (+0.4 %)	3.28 (−3 %)	3.79 (−0.1 %)

coal, so two changes act at once (§A.5.3, §3.5.2). Together they raise the conventional route's GHG intensity by 78 % and move it from 11 % below the electrochemical route to 59 % above it; of that 2.63 kgCO₂e/kg, roughly 1.6 is the missing abatement and 1.1 the coal-based ammonia. The price is held unchanged, on the assumption that CBAM is enforced as legislated: nitric acid falls inside its scope from January 2026, so an unabated import would carry an EU-equivalent carbon cost at the border and would not undercut EU supply. Knowing where a plant's nitric acid comes from is therefore an important condition for the conventional route's carbon standing against the electrochemical one.

Hydrogen

Hydrogen is mainly a conventional-route feedstock, and the central case buys blue at the IEA's assessed 2024 NW-European production cost of \$3.50/kg against \$6.65 for green (§A.5.2), figures that are optimistic against independent Netherlands project data. The result worth noting is on grey: hydrogen carries 5 % of the conventional route's footprint on blue but 15 % on grey, which lifts the route to 3.79 kgCO₂e/kg, level with the electrochemical route on either supply (3.79 central, 3.81 on grey). On emissions blue already sits close to green, so most of the gain is in the step from grey to blue: going on to green saves a further 0.11 kgCO₂e/kg where grey costs 0.40.

On green the carbon result is fixed at 3.28 kgCO₂e/kg and the uncertainty falls on price. The conventional route holds its advantage over the electrochemical one up to a green price of about \$9/kg and loses it above that. Both IEA figures sit below the line, the 2024 production cost of \$6.65/kg and the more optimistic 2030 projection of \$4.40/kg, so on those the route is both cheaper and 0.5 kgCO₂e/kg cleaner. What a plant would actually pay sits above a production cost, and the IEA's own European market price for renewable hydrogen, 8–10 \$/kg, straddles the line, so on a purchased-contract basis the advantage is no longer assured. The independent Dutch evidence spans it, \$6.58 to \$15.06/kg for current electrolyser projects, with the CE Delft and TNO 2030 projection of €8.30/kg (about \$9.13) landing almost exactly on it (§A.5.2). Which side of that line green hydrogen settles on is therefore the open question, and the same exposure runs through blue, bought here at \$3.50/kg: every \$1/kg on it adds about 0.06 \$/kg to the conventional route's price and brings it closer to parity with the electrochemical one.

Bio-derived benzene

In a search for less dependence on fossil feedstock, bio-derived benzene is a choice a benzene-based plant can make. It raises the levelised cost of adipic acid by 37 % on the conventional route and 30 % on the electrochemical one, and lowers the carbon intensity by 17 % and 14 %. That makes it the interesting comparison against the biomass route, because it reaches a less fossil-dependent adipic acid by changing a supplier while keeping the production process, and is therefore open to plants already built as well as to new ones. Table 4.11 sets the three routes side by side on fully non-fossil feedstock, green hydrogen included: the converted conventional route reaches 2.71 kgCO₂e/kg against the biomass route's 2.56, so the biomass route stays the cleanest, but by 0.15 rather than the 0.83 it holds over that incumbent as it runs today.

Neither of the two feedstocks that decide this comparison has a public spot price: the bio-derived benzene the converted routes buy and the muconic-acid broth the biomass route buys are both modelled from proxies, so what each would actually cost is the largest open question in it.³ Figure 4.4 sweeps both across their bands. On price the conventional route keeps its advantage: it would reach the biomass route's value only if bio-derived benzene cost about 3.80 \$/kg, some 65 % above the 2.31 \$/kg high anchor of its band, which is itself a tight-supply bracket rather than an expected price. Both levelised costs of adipic acid recover full capital, which makes that crossing the

³The bio-derived benzene band is triangulated from bio-naphtha and premium proxies rather than read from an index, and the muconic-acid broth, lacking a spot market, is carried at Mokwatlo's (2024) minimum selling price for muconic acid [51] less the cost of crystallising it out of the broth, with a ±30 % tier (§A.5).

strict case: a conventional or electrochemical plant already built, with its capital sunk, could switch to bio-derived benzene at a higher price still before a new biomass plant matched it.

Table 4.11: Each route on non-fossil feedstock: bio-derived benzene and green hydrogen for the benzene-based routes, against BIO_{AA}, so all three are per kilogram of adipic acid. *Central* is the headline case of Table 4.1; all inputs other than the two feedstocks are central in both columns. Neither non-fossil price is read from a market: both are modelled from proxies, and the non-fossil columns should be read with the price bands of Figure 4.4 rather than as point estimates. BIO_{AA} buys neither benzene nor hydrogen, so its two columns are identical.

Route	LCOA (\$/kg)		GHG (kgCO ₂ e/kg)	
	Central	Non-fossil	Central	Non-fossil
CON	1.87	2.76	3.39	2.71
ECO	2.19	2.86	3.79	3.24
BIO _{AA}	4.34	4.34	2.56	2.56




Figure 4.4: Each route's levelised cost of adipic acid swept across its feedstock-price band, everything else central. The conventional and electrochemical routes are on non-fossil feedstock and swept across the bio-derived benzene price (0.77–2.31 \$/kg), the biomass route, at its adipic-acid endpoint, across the muconic-acid-broth price (1.30–2.42 \$/kg); both bands are the literature-supported ranges of Table 2.5. Every printed value is a levelised cost of adipic acid: the bar ends at the low and the high feedstock price, the marker at the central one.

4.5. Cost and carbon sensitivities

Every result so far is quoted at one setting of inputs whose future values cannot be fixed in advance: what the EU ETS will charge once free allocation has gone, how a cell will perform at commercial scale, what a plant of this kind will cost to build in 2030. This section brackets each of them between the low and high tiers sourced in the literature review and set out in §3.3.1, varying one input at a time, and reads off what landing at either end of a bracket would do to a route's levelised cost of adipic acid and GHG intensity. The width of a bracket is the distance between two values that could each turn out to be the right one, and it carries no probability.

Two of the inputs have been treated already and are not repeated here: the price and carbon intensity of the electricity a site buys (§4.3) and the terms on which power and feedstock are contracted (§4.4). This section adds the three that have no section of their own, the ETS charge, cell performance and cost, and the capital estimate. The tornado charts then set every input against the others on one axis, one chart per route, so the brackets can be compared with each other instead of read one by one. That comparison is what identifies each route's key bottlenecks and separates the parameters that decide where it ends up from those that turn out not to matter.

4.5.1. EU ETS and abatement

The effect of the EU ETS on LCOA depends on three quantities that vary in this analysis. Two are (partly) policy-dependent: how much free allocation the plant still receives, and the carbon price. The third is the plant's own scope-1 emissions, held fixed for the electrochemical and biomass routes and varied only for the conventional one, where the impact of different N₂O abatement levels is assessed. The headline results fix all three at a single setting; what follows tests what other plausible outcomes would do.

Free allocation

Free allocation is carried as a bracket rather than as a forecast, because what adipic acid receives after 2030 is not settled (LR §2.4.1): the headline takes the position that stands today, full carbon-leakage protection at $\alpha = 1$, and the bound withdraws it entirely at $\alpha = 0$, with intermediate outcomes falling linearly between the two. The credit is the benchmark of 1.40 tCO₂e/t multiplied by the carbon price, so withdrawal is worth 0.177 \$/kg at the central 115 €/tCO₂, and 0.116, 0.223 and 0.257 \$/kg under CP-Low, CP-High and the rising CIP Net Zero trajectory. On the conventional route that withdrawal is the difference between sitting just below the 2025 EU market band of \$1.88–2.04/kg, where the headline 1.87 \$/kg leaves it, and being lifted to the top of that band. What it cannot do is change the comparison, since a benchmark-based credit is the same figure on every route and moves the level at which all three sit rather than the distances between them. A European plant's standing against a non-European one is a separate question, and the answer depends on how the allocation is lost. If adipic acid is brought inside CBAM the imported product is charged in the same act, so that standing is largely preserved; if Article 10b(1) lapses without that extension the European plant pays and the importer does not (LR §2.4.1).

The carbon price

The literature review spans a wide range of prices for a plant producing from 2030, covered here by the three named scenarios of Table 2.3: CP-Low at 75, CP-Central at 115 and CP-High at 145 €/tCO₂. Across that band all three routes get cheaper, because all three hold an allowance surplus to sell: the conventional route moves by 0.04 \$/kg and the electrified routes by about 0.10. A flat price is the central treatment; the alternative is the rising CIP Net Zero trajectory of §3.5.1, linear from 140 €/tCO₂ in 2030 to 210 in 2050, which levelises to 167 €/tCO₂ over the producing years. Implementing it deepens the conventional route's credit from 0.064 to 0.093 \$/kg and the electrified routes' credits by about 0.07 \$/kg each.

Carbon pricing narrows the gap between the conventional and electrochemical routes but does not close it. The two are 0.42 \$/kg apart at a zero carbon price, and because the conventional route carries more scope-1 emissions that gap closes linearly as the price rises; parity is reached only at about 480 €/tCO₂, some four times the central price, and then only on the central abatement assumption (Figure 4.5).

Abatement

Abatement is not a policy quantity but a property of the plant, and it has to be swept because operating plants report a spread of performance rather than a single figure: the technology can destroy essentially all of the N₂O, yet achieved rates vary with the installation (LR §2.2.1). It is therefore assessed across that spread, and in combination with the carbon price, because the two enter the ETS term as a product.

At the central carbon price of 115 €/tCO₂, moving destruction efficiency across the swept 95 to 100 % changes LCOA by 0.52 \$/kg (Table 4.12). At that same price the conventional route reaches the levelised cost of adipic acid of the electrochemical route on central inputs at an abatement of 95.9 %.

On carbon the same range matters more. The GHG intensity swings by 4.10 kgCO₂e/kg, wider than the route's entire central footprint of 3.39 and wider than the distance between any two routes in this study: at 95 % it sits at nearly twice the electrochemical route's intensity, and at complete destruction it draws level with the biomass route to adipic acid (§4.5.2). The abatement a plant actually achieves is therefore among the most important parameters in this study, and a viability condition for the conventional route rather than a detail of its costing.

Table 4.12: The conventional route against N₂O destruction efficiency, every other input central, with LCOA at the central carbon price of 115 €/tCO₂. The swept range is 95–100 %; the 90 % row lies below it and is carried to indicate the level a single non-redundant train could reach rather than one it would necessarily achieve, its 85–90 % availability placing realised annual abatement near there however well the train itself destroys (§3.4).

Destruction efficiency (%)	Scope-1 (tCO ₂ e/t)	GHG (kgCO ₂ e/kg)	LCOA (\$/kg)
90 (single train)	8.26	10.76	2.80
95	4.17	6.67	2.28
98	1.71	4.21	1.97
99 (central)	0.89	3.39	1.87
99.5	0.48	2.98	1.82
100	0.07	2.57	1.77

Every level in the table is an abated plant; what changes between them is how well the destruction performs, not whether it is installed. The facility's cost is therefore worth expressing per tonne of CO₂-equivalent it removes, and

checking against two benchmarks: what operating plants report, and the carbon price it exists to avoid (§3.3.3). At the 99 % central case the train costs 0.56 €/tCO₂e, its capital together with the maintenance and insurance carried on it, falling to 0.50 once the reaction heat is credited back, since the decomposition is exothermic and raises medium-pressure steam in the heat-integrated network. That is above the 0.10–0.40 €/tCO₂e reported for a first train (LR §2.2.1), but the configuration priced here is a redundant pair. A single train, taken indicatively at 90 % realised abatement, costs 0.31 €/tCO₂e, or 0.24 net of the same credit, which sits inside the reported range. The last nine percentage points of destruction are therefore the expensive ones, and they are still bought at about two orders of magnitude less than the ETS charge they remove.

The ETS charge scales with the N₂O that survives abatement, so each abatement level has its own break-even carbon price (Figure 4.5). At 95 % destruction the two cross at 94 €/tCO₂, already below the central 2030 price of 115 €/tCO₂; at 98 % they cross at 237 €/tCO₂, at the central 99 % at 479 €/tCO₂, and at 99.5 % at 979 €/tCO₂. At complete destruction they never cross, because the conventional route's allowance surplus then just exceeds the electrochemical route's, so its levelised cost of adipic acid falls marginally faster as the carbon price rises rather than more slowly. The biomass route is not plotted: at 4.34 \$/kg at the central carbon price it lies far above the range.




Figure 4.5: Levelised cost of adipic acid against the EU ETS carbon price for the conventional route at five N₂O destruction efficiencies, with the electrochemical route on central inputs as the dashed blue line labelled ECO. Open circles mark the two crossings that fall within the plotted range, at 95 and 98 % destruction. The vertical band is the carbon-price scenarios of Table 2.3, CP-Low 75 to CP-High 145 €/tCO₂, with the dashed rule at CP-Central 115 €/tCO₂; the horizontal band is the 2025 European market range.

4.5.2. The conventional route under complete N₂O abatement

Every route comparison so far has been measured against a conventional plant at the central 99 % destruction, which still vents 0.819 kgCO₂e/kg of unabated N₂O, the largest single carbon term the route carries (§4.5.1). The hardest test for the alternative routes is that plant with the lever taken as far as it goes: complete destruction on abated European nitric acid, which independently verified project monitoring reports at the European plants that added the redundant second unit costed here (LR §2.2.1). Its scope-1 then falls to the 0.073 kgCO₂e/kg of process CO₂ that abatement does not reach and its total to 2.57 kgCO₂e/kg, ninety per cent of it scope-3 in the benzene and the nitric acid. The electrochemical route emits 3.79 kgCO₂e/kg at the headline grid intensity of 261 gCO₂e/kWh, 47 % above that plant. The gap narrows as the grid cleans, but most of it does not close: 2.31 kgCO₂e/kg of the route's emissions is carried independently of the grid, so even on carbon-free electricity it would sit within a few per cent of the fully abated plant. Every input in this comparison is held at its central value.

The biomass route to adipic acid is level with that plant on the same grid, 2.56 against 2.57 kgCO₂e/kg, and, with every other input central, is cleaner on any grid below 263 gCO₂e/kWh. Every levelised grid in this study is below that (Table 4.9); only Germany's 2024 and 2030 anchors are above it.

Complete destruction helps the plant on price as well, because destruction efficiency sets the emissions the ETS charges on and leaves the largest allowance surplus to sell. Its levelised cost of adipic acid falls from 1.87 to

1.77 \$/kg (Table 4.12). It is already the cheapest of the three at 99 % destruction, and complete destruction widens that lead: from 0.32 to 0.42 \$/kg over the electrochemical route, and from 2.47 to 2.58 \$/kg over the biomass route to adipic acid. Complete abatement therefore strengthens the conventional route on both indicators at once.

Taken one step further, that plant buying bio-derived benzene and green hydrogen would reach 1.89 kgCO₂e/kg, below every headline intensity in this study. At the central bio-derived benzene anchor it costs 2.65 \$/kg, a price modelled from proxies rather than read from a market, and it stays below the biomass route to adipic acid across the whole band carried for that feedstock (Figure 4.4). It is the cleaner of the two on all three 2024 grids as well, but the two cross at about 144 gCO₂e/kWh and every levelised trajectory of \$4.3 falls below that, so on the projected grids the biomass route is the cleaner again.

4.5.3. The electrochemical cells

Of everything the electrified routes rest on, the electrochemical cell is the part with the least data behind it. Neither cell has been operated at anything approaching this scale, so there is no public benchmark to check either against. That is what the scenario framework is for: rather than predicting what the cells will do, it reports how each route would stand if a given level of performance and cost were reached, and how strongly that standing depends on it, which is what identifies the parameter that becomes the bottleneck.

Cell performance

What is swept here is Faradaic efficiency, three tiers of it, set on both cells at once from the levels that operating and laboratory cells report (LR Table 2.4). A lower efficiency increases the energy the cell draws per kilogram of product, and increases the hydrogen it makes, on both routes but at different rates (§3.2.5). Current density and cell voltage are pinned across the tiers, so the two cells carry identical capital at every tier: performance is swept separately from cell cost, and this comparison is unaffected by the capital-estimate band of §4.5.4.

Across the tiers the specific energy rises by about half on the electrochemical route and nearly doubles on the biomass route to adipic acid, and both KPIs rise with it, though nowhere near in proportion (Table 4.13): the electrochemical route's price moves 0.10 \$/kg and its GHG intensity 0.39 kgCO₂e/kg, and the adipic-acid end-point's 0.25 \$/kg and 0.89 kgCO₂e/kg. The extra hydrogen made at a lower Faradaic efficiency is what damps them, since the surplus is burned to raise steam the route would otherwise buy, lowering scope-2, while a higher efficiency leaves the route buying hydrogen and raises scope-3. The damping is larger on the electrochemical route because its cell passes six electrons per molecule of product against the biomass cell's two, so the same change in efficiency moves three times as much hydrogen.

On the electrochemical route that damping is complete between the top two tiers. Raising Faradaic efficiency from the demonstrated 0.80 to the target 0.90 cuts specific energy by six per cent and leaves the levelised cost where it was, at 2.19 \$/kg, the two tiers staying within 0.01 \$/kg of each other at every electricity price (§B.2), so almost the whole 0.10 \$/kg span is the conservative tier alone. On carbon the same step is still worth 0.06 kgCO₂e/kg. Improving the cell beyond the point already demonstrated is thus a carbon lever on this route and not a cost one; the biomass route does not share that, the same step moving its levelised cost from 4.34 to 4.22 \$/kg.

Cell voltage is not swept in this study, but because voltage changes the energy the cell draws without changing the hydrogen it makes (§3.2.5), the tiers can be read as voltage equivalents with the hydrogen credit held out. On that reading the electrochemical route's tiers span 11 % below to 78 % above its central-scenario voltage, worth 0.02 \$/kg and 0.07 kgCO₂e/kg less at the optimistic end and 0.15 \$/kg and 0.52 kgCO₂e/kg more at the conservative one; the asymmetry reflects how close its efficiency tiers already sit to the top of the demonstrated range. The biomass route to adipic acid's tiers sit 33 % either side of its own and respond symmetrically, 0.15 \$/kg and 0.51 kgCO₂e/kg either way.

Performance and electricity together

The cell's energy demand does not reach either KPI on its own. What it costs is set by the electricity price, and what it emits is set by the grid carbon intensity, so the three inputs enter as products rather than as sums, in the same way the ETS charge is the surviving N₂O multiplied by the carbon price. Both electrified routes are therefore swept over both pairings, cell tier against electricity price and cell tier against grid intensity, nine combinations each; Table B.2 in Appendix B sets them out in full.

Between its best and its worst corner the electrochemical route's price spans 0.45 \$/kg and its GHG intensity 1.60 kgCO₂e/kg, and neither width changes where it stands. No combination makes it cheaper than the conven-

Table 4.13: The two electrified routes by cell-performance tier. Each tier sets both cells at once, the ECO and ECH performance levels of LR Table 2.4; the specific energy is the route's cell and feed-compression electricity, and differs between them. Net H₂ is the route's hydrogen balance, positive where the surplus is burned for steam and negative where makeup is bought. The central tier is the headline case, priced on Belgium's extra-high-voltage tariff and carried on the Netherlands grid intensity (§4.2.4).

Tier	ECO				BIO _{AA}			
	SEC (kWh/kg)	Net H ₂ (kg/h)	LCOA (\$/kg)	GHG (kgCO ₂ e/kg)	SEC (kWh/kg)	Net H ₂ (kg/h)	LCOA (\$/kg)	GHG (kgCO ₂ e/kg)
Optimistic	4.76	−106	2.19	3.73	4.05	−60	4.22	2.11
Central	5.05	−32	2.19	3.79	6.01	+92	4.34	2.56
Conservative	7.04	+480	2.29	4.12	7.97	+244	4.47	3.00

tional route, the closest being 0.25 \$/kg above it, and only Belgium's grid, at 150 gCO₂e/kWh the cleanest in the cluster, makes it cleaner. Its carbon standing therefore turns on the grid rather than on the cell: its emissions climb with grid intensity about nine times as fast as the conventional route's, so the two are level at 181 gCO₂e/kWh on the central cell (§4.3) and the cell tier moves that crossing only between 159 and 189 gCO₂e/kWh. Both the Dutch and the German grid fall below it on their levelised intensity over the producing years of a plant commissioned in 2030 (§4.3).

The biomass route to adipic acid spans wider still between the same two corners, 0.61 \$/kg and 2.04 kgCO₂e/kg, but its standing does not change with it either. Its cost distance from the conventional route never falls below 2.30 \$/kg, so no combination brings it near, and on carbon it is the cleaner route in every combination but one, the conservative cell on the German grid.

Cell degradation

The central-tier charge is given in §4.2.5. Degradation reaches the electrochemical route through both its specific energy and its capital, and the biomass routes through capital alone, but the uncertainty is capital on both: the 10 % end-of-life criterion is held fixed across the tiers, so the energy charge is the same on every stack life and only the replacement count moves. The same shows in the route's specific-energy trajectory, where the tiers change how often the membrane is replaced and not how far the specific energy climbs between replacements (§B.1). The spread is limited on both routes (Table 4.14), and on emissions there is no spread across the durability tiers at all: what degradation adds stays at 0.034 kgCO₂e/kg on the electrochemical route and at zero on the biomass ones, on every stack life. It is an electricity charge, however, so on the electrochemical route it does move with the grid and with the cell-performance tier, which is why it is not constant across the panels of Table B.2.

Table 4.14: What cell degradation adds, by durability tier (\$/kg), with replacements over the twenty producing years in brackets. The electrochemical route's membrane charge splits into the extra electricity the rising voltage draws and the replacement capital; the biomass cathode corrodes at constant voltage and has no energy channel.

ECO membrane				BIO _{AA} cathode	
Stack life	Energy	Capital	Total	Cathode life	Charge
40 000 h (4)	0.009	0.004	0.014	7 yr (2)	0.007
20 000 h (8)	0.009	0.009	0.018	4 yr (4)	0.015
5 000 h (32)	0.009	0.038	0.048	2 yr (9)	0.035

Cell cost

Cell cost is a separate lever from cell performance, and it rests on as little data. What the sweep asks is what a more expensive or a cheaper cell would be worth. The low end, 30 % less, is a learning-curve endpoint, between one and two doublings of installed capacity at the 16–21 % per doubling water electrolyzers have shown; the high end, 50 % more, is not a trajectory but the width of the estimate itself, the loose end of the Class-4 envelope carried by the least-defined equipment group (§3.5.3). Across that range the electrochemical route moves from 2.19 to 2.07 and 2.38 \$/kg and the biomass route to adipic acid from 4.34 to 4.14 and 4.67 \$/kg. Neither end of that range is enough to bridge the gap to the conventional route: even a fully matured cell leaves the electrochemical route 0.21 \$/kg above it, against the 0.32 of the central case and of the optimistic performance tier (Table 4.13).

4.5.4. Capital and discount rate

Capital investment is a dominant cost driver across all pathways, particularly for the capital-intensive electrified and biomass routes. To capture this sensitivity the AACE Class-4 equipment-maturity weighting defined in Methodology §3.3.3 is applied across each route according to its specific equipment composition. The band

widens with route novelty and equipment unprovenness, from $-16/ +33\%$ on the conventional route to $-26/ +45\%$ on the biomass route to adipic acid, worth -0.04 to $+0.08$ \$/kg on the first and -0.20 to $+0.35$ \$/kg on the last. Part of the fixed operating cost is charged on the capital base, maintenance at 5% of inside-battery-limits cost and taxes and insurance at a further 1.5% , so it moves with capital and adds about two-thirds as much again. Every route's band and both effects on price are tabulated in §B.4.

These bands are not independent of one another. The routes share equipment classes and a single cost basis, so an error in the correlations or in the installation factors moves them in the same direction rather than in opposition; what is genuinely independent is the electrochemical cell, half of that route's inside-battery-limits capital and the item carrying the widest group band, which the conventional route does not have at all. Taken one route at a time, neither band closes the 0.32 \$/kg between the conventional and electrochemical routes.

Sensitivity to the cost of capital behaves differently. Raising the real discount rate from the 8% baseline to 15% raises the levelised cost of adipic acid by 0.23 \$/kg on CON, 0.39 \$/kg on ECO and 0.63 \$/kg on BIO_{AA}, taking them to 2.10 , 2.58 and 4.98 \$/kg. The ranking survives, but both distances to the conventional route widen, from 0.32 to 0.48 \$/kg on the electrochemical route and from 2.47 to 2.88 \$/kg on the biomass route to adipic acid, because financing cost falls hardest on the most capital-intensive route.

4.5.5. Tornado analysis

With every driver now introduced, the tornado charts set them against one another on a single axis, and the ranking that results varies by route (Figures 4.6–4.8). Two levers are deliberately left off. The EU ETS free allocation is benchmark-based rather than emissions-based, so it shifts every route by the same 0.177 \$/kg, and the discount rate is a property of the investor rather than of the route, reported separately in §4.5.4.

Conventional route

The tornado chart (Figure 4.6) shows that the levelised cost of adipic acid of the conventional route is dominated by raw feedstock costs, with the market price of benzene, alongside potential future premiums for bio-based benzene, acting as the single largest financial lever. Because the heat-exchange network is highly integrated, external steam demand is virtually eliminated and steam price fluctuations are negligible. Carbon pricing via the EU ETS similarly shows limited impact under the baseline high-abatement scenario, though its influence grows substantially at lower abatement levels, where unmitigated N_2O emissions incur direct carbon cost penalties. Grid carbon intensity moves the route's emissions but not its price, because the EU ETS is levied on scope-1 alone.




Figure 4.6: Conventional route: one-at-a-time sensitivity, grouped by driver class and, within a group, by the larger of the row's relative swings in cost and in carbon. Bars span the tier range of $\$3.3.1$; the line is the central case. The two capital rows are different bands and must not be added.

On greenhouse gas intensity, N₂O abatement efficiency is the most critical lever. This study evaluates a baseline range of 95 to 100% destruction, and efficiencies below 95% cause emissions to escalate substantially toward current industry averages. The second largest effect stems from the supply origin of HNO₃, where sourcing from unabated, coal-based non-EU producers imposes a severe carbon penalty relative to European suppliers. Upstream supply choices for benzene and H₂, and general utility demands, have a minor effect on the route's total footprint.

Electrochemical route

The electrochemical route (Figure 4.7) keeps the benzene dependence of the conventional pathway but moves the rest of its exposure from chemical feedstocks to utilities and capital, since hydrogen is co-produced at the cathode. The immaturity of the cell and the electrodialysis unit is what puts capital high on the chart, and the same capital intensity makes the route far more exposed than the conventional one to the cost of capital, reported outside the chart in §4.5.4. Cell performance moves the levelised cost comparatively little within the bounds analysed, though a commercial cell could fall outside them. Because the route generates almost no scope-1 emissions it keeps nearly the whole benchmark allocation, so the EU ETS enters as a credit of 0.17 \$/kg rather than a charge and a higher carbon price makes the route cheaper, by about 0.10 \$/kg across the tier band.




Figure 4.7: Electrochemical route: one-at-a-time sensitivity, grouped and ordered as in Figure 4.6. *Other process chemicals* is the KOH electrolyte make-up (§4.5.5); no sensitivity on its scope-3 intensity is carried in this analysis, so the row has no carbon bar. The two capital rows are different bands and must not be added.

On greenhouse gas intensity, the route is governed by scope-2 utility demand rather than by direct process reactions. Grid carbon intensity and process steam sourcing are the primary levers, and a low-carbon grid such as the French mix evaluated in the location analysis would lower emissions beyond the range shown in the chart. Upstream benzene origin exerts a substantial influence, whereas hydrogen supply choices have a negligible impact. Cell performance has a moderate effect, because the carbon penalty of the higher electricity consumption a lower Faradaic efficiency brings is partly offset by the accompanying increase in cathodic hydrogen (§4.5.3).

Biomass route

The biomass route (Figure 4.8) is dominated by its main feedstock as well, here the muconic-acid broth rather than benzene: its price band alone is worth 1.35 \$/kg at the adipic-acid endpoint, the largest single bar anywhere in this study, and it is a modelled price rather than a market one (§4.4). Capital sits high for the same maturity reason as on the electrochemical route, and the larger capital base makes this the route most exposed of the three to the cost of capital (§4.5.4). Cell performance moves it only moderately, because a change in efficiency is partly offset by the cathodic hydrogen it shifts.




Figure 4.8: Biomass route: one-at-a-time sensitivity on the adipic-acid basis, grouped and ordered as in Figure 4.6. *Other process chemicals* is the H_2SO_4 that crystallises the product and the $\text{Ca}(\text{OH})_2$ that neutralises the wastewater (§A.5.5); no sensitivity on their scope-3 intensities is carried in this analysis, so the row has no carbon bar. The two capital rows are different bands and must not be added.

On greenhouse gas intensity, cradle-to-gate emissions are relatively low, which makes electricity demand the main driver of what remains. Grid carbon intensity is the primary lever, and lower-carbon power would reduce emissions beyond the range shown in the chart. Cell performance has a notable relative impact, since efficiency variations alter the electricity demand per unit of product against an already low baseline.

Two readings cut across all three. The carbon price acts in the same direction on all three routes, because all three hold an allowance surplus to sell: a higher price lowers the conventional route's levelised cost of adipic acid by 0.04 \$/kg across the tier band and the electrified routes' by about 0.10, so it narrows the gap without reordering anything. And the grid carbon intensity has no effect on any route's cost, because the ETS is levied on scope-1 only and the electrified routes have almost none, so electricity is a carbon lever for them and a price lever only through its tariff.

4.6. Load-following under dynamic electricity prices

The central case runs every route as continuous baseload. This section tests whether concentrating operation in the cheapest hours of the 2024 day-ahead series lowers the LCOA or the scope-2 GHG intensity (§3.5.5). The conventional route is thermal and has no shiftable electrochemical load, so the two electrified routes are analysed, and load factor is the share of the plant's 8 000 operating hours that it runs, so $LF = 0.875$ is 7 000 h. Of the two strategies set out there, *produce-less* stops the whole plant in the expensive hours and makes less product, while *oversize* holds output by building the cell $1/LF$ larger and needs an uncosted intermediate buffer; the results are read from *produce-less*, the strategy a plant can be built to without one.

Electricity is priced here from the realised 2024 hourly series in each country rather than from the central tier the headline results are built on, so what follows are changes measured against this section's own baseload, the potential improvement load-following can deliver, rather than revised values for either headline measure.

Effect on the electricity price

What a lower load factor offers is cheaper electricity, and how much cheaper is decided by national grid regulation rather than by the wholesale market (Figure 4.9). The commodity falls everywhere as only the cheaper hours are kept, but only half the grid tariff follows: the capacity half is billed on contracted connection and peak draw, which a plant running at full load whenever it runs does not reduce, so spread over less electricity its cost per megawatt-hour rises as $1/LF$. How the tariff divides differs sharply by country (§E.1): the Dutch is entirely capacity-based, the Belgian mostly per megawatt-hour, and the German small on both only because of its §19 StromNEV relief, which is conditional on running enough full-load hours.

The three countries therefore behave differently as the load factor falls. The Netherlands falls fastest at first, but its capacity charge limits what further cheap hours are worth, so the effective price decreases only from €97/MWh at $LF = 1.00$ to €82 at $LF = 0.50$. Germany has the lowest baseload cost of the three at €48/MWh, but its relief shrinks at 7 500 full-load hours and is withdrawn entirely at 7 000, $LF = 0.875$, the dashed line in its panel, where the price jumps €26/MWh to above its own baseload level. Belgium starts above Germany and falls throughout, from €52 to €23, crossing below it at that threshold: below $LF = 0.875$ it is the cheapest of the three by a wide margin, €40/MWh below Germany and €59 below the Netherlands at half load. There is therefore real potential in Belgium for a plant able to operate at a low load factor. The *oversize* bound buys the same hours and differs only in how much flexible load is shifted into them, so on this measure the two strategies stay within €1/MWh of each other at every load factor.




Figure 4.9: Effective electricity price against load factor, decomposed; the values are tabulated in Table E.2. Bars stack the commodity paid over the hours actually run, the network and capacity charge, and the remaining levies and taxes, with the ICC/CO₂ compensation drawn below zero as the credit it is; the marked line is the net. The dashed line in the German panel is 7 000 full-load hours, $LF = 0.875$, below which the §19 StromNEV relief is lost. Drawn for the *produce-less* case.

Effect on LCOA

Almost none of that potential reaches the levelised cost of adipic acid (Figure 4.10). On the *produce-less* baseline the optimum is $LF = 1.00$ in four of the six cases: the biomass route to adipic acid in all three countries, and the electrochemical route in Germany. The two exceptions, the electrochemical route in the Netherlands and in Belgium, do fall below baseload at $LF^* = 0.99$, but are worth 0.004 and 0.001 \$/kg, under 0.2 % of that route's levelised cost of adipic acid; past them every curve turns upward and keeps rising to the end of the sweep.

The reason is the cost structure of these two plants rather than the electricity market. Capital charge and fixed operating cost come to 1.04 \$/kg on the electrochemical route and 1.67 on the biomass route to adipic acid, and are incurred however many hours the plant runs. Dropping to $LF = 0.95$ means 5 % less product to spread that cost over, which adds 0.055 \$/kg on the electrochemical route, while the cheaper hours bought in exchange return at most 0.032, in Germany: the trade gives up close to twice what it wins back, and more than twice as much again on the biomass route to adipic acid, whose fixed cost is higher against a similar electricity bill. A process whose cost sat in its utilities rather than in its capital would find real value in a low load factor, and would find most of it in Belgium; neither route studied here is such a process.

The *oversize* bound, which holds annual output by building the cell $1/LF$ larger and stopping only the cell, does not change that. Its optima are the rings on the dashed curves of Figure 4.10, at $LF^* = 0.94$ – 0.98 , and the best of them is worth 0.007 \$/kg, under 0.7 % of that route's own levelised cost of adipic acid, bought with three to seven days of intermediate storage that is not costed here (§E.5). The extra cell degradation this dispatch causes is not a leading constraint either: at those optima the electrochemical cell cycles 127 to 219 times a year, which costs under 0.00015 \$/kg (§E.6), although that partly reflects how close those optima sit to baseload.

Effect on GHG intensity

Drawing power from the grid in the cheap hours means the hours given up are the gas- and coal-dominated ones, so the same dispatch that raises the price lowers the carbon intensity of the power drawn. At half load the electrochemical route emits 0.40 kgCO₂e/kg less in the Netherlands and 0.45 less in Germany, around a tenth of




Figure 4.10: Change in levelised cost of adipic acid against load factor, measured from $LF = 1.00$ on the same dispatch. The upper panels enlarge the shaded band of the lower ones, $LF = 1.00$ to 0.90 , on a scale twenty times finer. Solid lines are the *produce-less* baseline, dashed lines the *oversize* bound, and rings mark each cost-optimal load factor; where a line carries no ring its optimum is $LF = 1.00$, at the origin. The two steps in the German curves are its $\$19$ StromNEV relief tiers, the relief narrowing below 7 500 full-load hours ($LF = 0.9375$) and being withdrawn below 7 000 ($LF = 0.875$).

its intensity, and the biomass route to adipic acid 0.45 and 0.50, nearer a sixth of its own. Belgium gains least, 0.18 and 0.20 kgCO₂e/kg on the two routes, because its hourly carbon intensity varies far less across the year, so there are fewer dirty hours to avoid. That reduction is nonetheless bought at 0.96 and 1.13 \$ per kgCO₂e avoided on the electrochemical route, far above the carbon price used anywhere in this study, and at the cost-optimal load factors of the *oversize* bound it falls to 0.015–0.043 kgCO₂e/kg, under 1.1 % of a route’s intensity. The full sweep is plotted in §E.4.

Taken together, every effect measured in this section is small against the levers ranked in §4.5 above. The best load factor is worth 0.007 \$/kg, and that only with the buffer left uncosted, against 0.27 \$/kg for the electricity-price tier on its own. How the plant is dispatched is not a lever of consequence for these two routes, and baseload remains the central case.

4.7. Coherent-future narratives

The sensitivities in this work mostly move one input at a time, or at most a pair of them together, which is enough to identify each route’s key bottlenecks and what each component is worth, but not what a real future does, in which the inputs are correlated rather than independent. A bundle of them can land somewhere no single-input sweep reaches. The coherent-future narratives move all of them at once. Three bundles, *Decarbonisation lags*, *Reference* and *Decarbonisation accelerates*, each set eight scenario dimensions to the values they would plausibly take if decarbonisation ran at one pace throughout; the dimensions, the tier each takes in each future and the reasoning behind those assignments are given with the bundles in the methodology (§3.5.4, Table 3.7).

They are not forecasts, and the correlations between the dimensions are assumed rather than derived, so what follows compares the three futures rather than saying which of them arrives. Against the central case, *Reference* holds every headline input except the grid carbon intensity, which follows the declining Dutch trajectory rather than the static 2024 value; *lags* holds that static value too, and free allocation is withdrawn only under *accelerates*.

Effect on LCOA

Reference reproduces the headline price on every route, because every dimension it moves that reaches price is already at its central value.

Decarbonisation lags holds the largest deviation, raising every route: by 0.79 \$/kg on the conventional route, 0.91 on the electrochemical one and 1.31 on the biomass route to adipic acid. Feedstock is what does it, because the lagging world takes the upper tier of every band: fossil benzene at 1.65 \$/kg against the 0.88 the reference world pays, and the muconic-acid broth, already 2.08 of the biomass route's 4.34 \$/kg, at the top of its own. The routes carrying the most feedstock exposure move furthest, which is why the ranking survives a future that raises all three. The three futures are decomposed by cost component in Figure 4.11.

Decarbonisation accelerates moves them in opposite directions. The conventional route rises 0.26 \$/kg, almost exactly the free allocation it loses at that future's levelised carbon price. The electrochemical route absorbs the identical withdrawal and rises only 0.03, because the matured cell offsets it on two counts at once: capital 30 % below today's on the learning curve, and a Faradaic efficiency of 0.90 against the central 0.80, which lowers the electricity the cell draws for every kilogram it makes. The biomass route falls 0.82, almost all of it the cheaper muconic-acid broth. The conventional route is therefore the only one dearer in both directions away from the reference world, its feedstock bill rising in one and its carbon bill in the other.




Figure 4.11: Levelised cost of adipic acid under the three coherent futures, decomposed by cost component on the build-up of Figure 4.1. The EU ETS term is net and is drawn below zero where the route sells more allowances than it surrenders.

The cost ranking never inverts, but it comes close to doing so once. The conventional route is cheapest in all three futures, by 0.44 \$/kg where decarbonisation lags, 0.32 in the reference world and 0.09 where it accelerates. Read against each route's own price the first two are the same margin, 17 % in both, and only the accelerating future changes the picture, at 4 %. What closes the gap is narrower than the bundle that produces it. The matured cell reaches the electrochemical route alone, and the rising carbon price reaches the conventional route hardest because it carries an order of magnitude more scope-1, worth 0.046 \$/kg of the 0.23 the gap closes by. The withdrawal of free allocation, the largest single policy move in that bundle, is benchmark-based: it raises both routes by the same amount and compresses nothing.

Effect on GHG intensity

Reference emits less than the headline on every route, because the grid declines over the producing years instead of being held at its 2024 value and levelises to 103 gCO₂e/kWh against the headline's 261. The routes do not fall equally: 0.10 kgCO₂e/kg on the conventional route against about 0.9 and 1.0 on the electrified ones.

Decarbonisation lags raises every route, and the conventional one most: 3.13 kgCO₂e/kg above its reference level, against 1.22 on the electrochemical route and 1.43 on the biomass route to adipic acid. The dominant term on the

conventional route is its nitric-acid supply, worth 2.63 kgCO₂e/kg on its own (§4.4); on the electrified routes the difference is the static grid against the declining one. The scope split behind each future is shown in Figure 4.12.

Decarbonisation accelerates lowers every route, but not through the same input. The conventional route falls 0.73 kgCO₂e/kg below its reference level, most of it the switch from fossil to bio-derived benzene rather than the grid, which reaches it barely at all; the electrochemical route falls 1.00, roughly half benzene and half grid, since it buys the same feedstock and draws far more power; and the biomass route to adipic acid falls 0.50, effectively all of it the grid at a levelised 24 gCO₂e/kWh, because it buys no benzene. The gap between the electrochemical and biomass routes closes with it, from 1.33 kgCO₂e/kg in the reference world to 0.83 here, against 1.23 in the headline.




Figure 4.12: Cradle-to-gate GHG intensity under the three coherent futures, decomposed by scope on the build-up of Figure 4.2.

The headline ordering does not survive any of the three. In the headline the electrochemical route sits 12 % above the conventional one; in every future it sits below, by 36 % where decarbonisation lags, 12 % in the reference world and 26 % where it accelerates, with the biomass route cleanest throughout. The reversal comes about differently at either end. Where decarbonisation accelerates it is the clean grid, which the crossing at 181 gCO₂e/kWh (§4.3) puts the electrified routes on the near side of; where it lags the grid never moves, staying at the headline's static 261 gCO₂e/kWh, and scope-3 does the work instead, carrying 97 % of the conventional route's rise as it more than doubles from 2.32 to 5.35 kgCO₂e/kg, 2.63 of that unabated nitric acid and 0.40 grey hydrogen, both procurement choices the same absence of policy makes the rational buy (§4.4).

Conclusions

This work set out to compare the conventional, electrochemical and biomass routes to adipic acid on levelised cost and cradle-to-gate GHG intensity, on one basis, and to test each of them across the range of market, policy and performance conditions a European plant could plausibly face rather than at one fixed set of inputs. §5.1 answers the eight sub-questions of §1.4, in the route grouping they are posed in, and introduces no new results; §5.2 resolves them into the bottlenecks and viability conditions the central question asks for, and §5.3 what would most reduce the sector's emissions.

No route leads on both measures, and none of the three standings is set by the process alone. On the central inputs the conventional route is the cheapest, the electrochemical route the most carbon-intensive, and the biomass route the lowest on cradle-to-gate GHG intensity at roughly twice the price of either; but each of those standings turns as well on what the plant buys and on how it is run, which is what the findings below establish. Against the reference model this work builds on, Yeo (2025) [9], the corrections and the added scope narrow the distance the comparison inherited, and narrow it from the conventional side (§4.2), leaving the conventional and electrochemical routes 0.32 \$/kg and 0.40 kgCO₂e/kg apart: that pair is the comparison worth testing under varying conditions.

5.1. Answers to the research questions

Abatement decides the conventional route's standing on price as well as GHG intensity. Its destruction efficiency sets how much scope-1 emission the EU ETS can charge, so the two act as a product. At the central 99 % destruction the route emits 3.39 kgCO₂e/kg and reaches the electrochemical route's central price only at about 480 €/tCO₂. At 95 % it emits 6.67 kgCO₂e/kg, nearly twice as much and above the electrochemical route's own intensity, and that same crossing falls to 94 €/tCO₂, below the central 2030 price of 115. Free allocation cannot reorder anything, being benchmark-based and worth the same to every route, but it is not small: withdrawing it in full raises every route's levelised cost by 0.177 \$/kg.

A fully abated plant on European nitric acid captures nearly all of what a change of route offers on carbon. At complete destruction, which the redundant twin destruction train costed here is built to approach and which the European plants reporting it installed a second unit to reach (LR §2.2.1), it emits 2.57 kgCO₂e/kg against the 6.3 to 25.6 kgCO₂e/kg published life-cycle estimates report for conventional adipic acid, and the last percentage point of destruction is worth 0.82 kgCO₂e/kg against this study's own central 3.39. At an LCOA of 1.77 \$/kg it would then be the cheapest of the three and cleaner than the electrochemical route on every grid modelled here, France's near-nuclear one included. With bio-derived benzene and green hydrogen as well it would reach 1.89 kgCO₂e/kg and undercut the biomass route to adipic acid on all three ARRRRA grids as they stood in 2024, and on price at every anchor in the bio-derived benzene band, the two meeting on carbon at about 144 gCO₂e/kWh (§4.5.2). Neither non-fossil feedstock is yet made at this scale.

Nitric-acid supply is the biggest procurement lever on carbon. An unabated non-EU supplier, venting its own N₂O and making its ammonia from coal, raises the conventional route's intensity by 78 % at no change in price, nitric acid being inside CBAM scope (§4.4), taking it from 11 % below the electrochemical route to 59 % above.

The conventional route's carbon standing is therefore constrained by what a buyer can verify about that supply and how free it is to choose the supplier at all. Hydrogen is the smaller procurement decision and reaches mostly the conventional route as well, the electrified ones evolving their own at the cathode: across grey, blue and green that route's price varies by 0.21 \$/kg and its intensity by 0.51 kgCO₂e/kg.

Bio-derived benzene could be a cheaper path to fossil independence than the biomass route. Switching benzene alone moves both benzene-based routes toward non-fossil feedstock at a price, about 0.7 \$/kg for about 0.55 kgCO₂e/kg on each. Complete fossil independence, benzene and hydrogen both, puts the conventional route at a carbon intensity comparable to that of the biomass route to adipic acid on the static 2024 Dutch grid, and at 2.76 \$/kg against 4.34; that ordering holds across the price ranges the literature supports for both feedstocks. Neither feedstock price is read from a market, and neither the muconic-acid broth nor bio-derived benzene, modelled here on the catalytic fast-pyrolysis route, is yet made at this scale (LR §2.4.6).

Electrification relocates the emission rather than removing it. On the ARRA grids that would actually supply it, rather than the near-renewable power assumed in the baseline of Yeo (2025), the electrochemical route gives back most of its carbon advantage: what it removes in process N₂O and nitric acid the grid returns in scope-2, and on the Dutch and German 2024 grids it ends above the conventional route rather than well below it. The condition is a single number, a grid below 181 gCO₂e/kWh, which of the three only Belgium's meets in 2024, and which all three meet once their decarbonisation trajectories are levelised over a 2030 plant's producing years. Moving to France's near-nuclear grid takes 1.30 kgCO₂e/kg off the route, 34 % of its intensity, at no change to the process.

EHV electricity tariffs and their country-specific exemptions have a significant impact on the electrified routes' competitiveness. Buying at the relieved tariff a plant of this scale qualifies for, rather than at the generic industrial figure used in the baseline of Yeo (2025), lowers the electrochemical route's electricity bill by 0.53 \$/kg and the biomass route to adipic acid's by 0.62, against 0.05 on the conventional route (§4.2.4). Differences across the ARRA grids, mostly in tariffs and the policies attached to them, are worth 0.27 \$/kg on the electrochemical route, which puts its premium over the conventional one at 26 % in the Netherlands and 13 % in Germany.

Baseload operation remains the economic choice for the electrified routes. Concentrating their production into the cheaper, less fossil-dependent hours has real potential on both indicators: between baseload and half load the effective electricity price falls by €15 to €30/MWh in the Netherlands and Belgium, and the carbon intensity of the power falls with it, by 0.18 to 0.50 kgCO₂e/kg across the two routes and the three countries. None of it reaches their levelised cost of adipic acid: capital and fixed operating cost are incurred however many hours the plant runs, so per kilogram they rise faster than the electricity price falls. The carbon saved at half load costs 0.96 to 1.13 \$ per kgCO₂e avoided. Under the operating and storage assumptions carried here, load-following is not an economic lever for processes this capital-intensive.

Cell degradation has little impact on the modelled KPIs of either electrified route. It reaches the two differently, the electro-oxidation cell by voltage rise, carrying both a replacement charge and a rising energy demand, the hydrogenation cell by corrosion at constant voltage, carrying a replacement charge alone. Within the degradation bounds adopted here, accounting for degradation raises the levelised cost by well under 1 % at the central stack life and by about 2 % at the shortest. On carbon it adds 0.034 kgCO₂e/kg on the electrochemical route and nothing on the biomass ones, the same at every stack life, because the end-of-life criterion is held fixed across the tiers.

Neither biomass endpoint competes without a substantial price premium. Hydrogenating *trans*-3-hexenedioic acid to adipic acid costs 0.32 \$/kg and 0.22 kgCO₂e/kg, which leaves its ordering unchanged and widens its premium. That premium is mainly the feedstock: the muconic-acid broth alone contributes 2.08 \$/kg, more than the conventional route's entire levelised cost. At the *t*3HDA endpoint the position is uncertain rather than settled, a premium having been argued for on the retained double bond but never demonstrated commercially (LR §2.4.4). On carbon the route has the lowest cradle-to-gate GHG intensity of the three, 31 % below the conventional route at the *t*3HDA endpoint and 24 % at the adipic-acid endpoint.

Complete pinch analysis moves GHG intensity, not LCOA, and moves it most on the conventional route. How much it moves depends on how much recovery the inherited flowsheet already carried: complete integration removes 0.49 kgCO₂e/kg from the conventional route, 13 % of its intensity, against under 2 % on the electrified ones. That saving exceeds the 0.40 kgCO₂e/kg separating the conventional and electrochemical routes in the headline case and is part of why the conventional route comes out below, so the carbon ordering of that pair should not be read independently of the integration each flowsheet is assumed to carry. On levelised cost it is close to break-even everywhere, within about 1 %, the utility saving and the exchanger capital that delivers it

nearly cancelling.

Joint variation can reorder the routes, and no ARRRR site is best on both indicators. Sweeping destruction efficiency and carbon price together reorders the routes at 95 % abatement and leaves them untouched at 99, and sweeping cell performance against electricity price and grid intensity moves the electrochemical route by 0.45 \$/kg and 1.60 kgCO₂e/kg across the nine combinations of each pairing. Coupling inputs by country asks what a real site offers rather than what a parameter does: Germany gives that route its lowest electricity price and its worst carbon position. The narrative bundles ask what a future in which several inputs move at once would do, and in the accelerating one the electrochemical route comes within 0.09 \$/kg of the conventional route on price. What a route is worth is therefore a function of the combination it sits in, which no single-input sweep recovers.

5.2. Key bottlenecks and viability conditions

Each route turns as much on policy and on circumstances outside the plant as on the process itself. A key bottleneck for the conventional route is its own abatement efficiency, and where that efficiency is low, the carbon price becomes one as well: what the plant pays on the N₂O it still vents then moves its standing sharply, where at near-complete destruction it barely does. Another is its nitric acid, meaning both the freedom to choose the supplier and the ability to establish what that supply emits. A dominant lever on price is benzene, but its price is not a bottleneck for competitiveness across the routes, the most competitive alternative, the electrochemical route, depending on the same feedstock.

For the electrochemical route it is electricity: low-carbon supply for its carbon standing, and for its price the grid tariffs and the policy attached to them, the CO₂-compensation relief granted in Germany and Belgium accounting for most of what separates the ARRRR countries, together with the feasibility constraint of securing a grid connection at all. Its electrochemical cell is a constraint on both feasibility and competitiveness, its performance and its capital cost alike unproven at this scale.

For the biomass route it is the cost of its main feedstock, high and not yet set by a market. It inherits the same constraints as the electrified routes and still carries a premium over both, so at the adipic-acid endpoint it turns on whether a buyer will pay that premium for a molecule identical to the incumbent's, or on a muconic-acid broth reaching a much lower cost than any yet demonstrated, and available at the appropriate scale. At the t3HDA endpoint it sells a different molecule, and turns instead on whether the reactive double bond that molecule retains is worth a premium of its own.

5.3. Implications for the wider sector

Set against the sector rather than against a single European site, the biggest lever for reducing the emissions of adipic-acid production is not a change of route at all. It is abating the direct N₂O that unabated capacity still vents, and buying nitric acid from producers that have done the same and do not make their ammonia from coal. Global N₂O from adipic-acid production in 2021 has been estimated at about 142 MtCO₂e, with up to 94 % of it attributed to Chinese capacity [12], well above its share of world output (LR §2.1), and what separates the two shares is destruction efficiency. Those two moves therefore reach the largest part of what the sector emits, and reach it far more cheaply than any change of route. Destruction costs 0.56 €/tCO₂e for the redundant pair costed here (§4.5.1) and 0.10 to 0.40 as reported by operating plants [34]. Changing route costs thousands of dollars for the same tonne: about 2,060 \$/tCO₂e avoided on the biomass route and about 2,980 at its adipic-acid endpoint.¹ The electrochemical route saves nothing at the central grid, emitting more than the incumbent there; on France's near-nuclear grid it saves at about 320 \$/tCO₂e, still hundreds of times what abatement costs. Abatement acts on the direct emissions alone, however: the dependence on fossil benzene that both the conventional and the electrochemical route carry remains, and only a change of feedstock or of route ends it.

What could create a reason to abate is a price signal on imported adipic acid. The EU ETS prices the N₂O a European plant vents and CBAM prices the nitric acid it imports, but adipic acid itself sits outside CBAM scope, so a producer selling into Europe carries no charge for what it vents. Bringing it inside would give plants exporting to Europe a reason to abate; moving the sector as a whole to near-fully abated production needs an incentive that reaches producers wherever they sell.

¹Extra cost over emissions saved, both against the conventional route as costed here: 2.16 \$/kg over 1.05 kgCO₂e/kg on the biomass route and 2.47 over 0.83 at its adipic-acid endpoint (Table 4.1); the French figure uses that row of Table 4.8.

Recommendations

What this work does not settle is stated below as the work that would settle it. Some items follow from scoping decisions, the routes compared, the plant size, the impact category, the region; others from modelling choices made within the comparison and bracketed by the sweeps reported with the results. The two structural simplifications behind the model, the electrochemical cells as imposed operating points and the inherited Aspen flowsheets, are stated with the model scope in the methodology (§3.2.7). The first items apply to the comparison as a whole; the rest follow the route grouping used in the preceding chapters, and the last looks beyond the plant.

Investigate alternative routes and feedstocks. Three routes are compared here, and within the biomass route only the electrochemical hydrogenation variant, whose hydrogenation step is costed on conditions demonstrated for muconic acid rather than for isolated *t*3HDA, so its adipic-acid endpoint is a costed extension of the route rather than a demonstrated process. Thermocatalytic hydrogenation of purified muconic acid is reported to be cheaper wherever no *t*3HDA premium exists [101]. That variant, and the alternatives to the two routes proposed here, could be assessed under the same framework, which would show how far the comparison extends beyond these three. Bio-derived benzene is likewise represented by a single production route, catalytic fast pyrolysis, and what the alternatives to it would cost and emit sets how far the bio-benzene cases hold.

Assess and optimise the scale of each route. The three are compared at one design basis and scale by different mechanisms: conventional capital sits in large single trains, where cost per kilogram falls steeply with size, electrochemical capacity is added as further cells rather than larger ones, and the biomass route is dominated by a broth whose cost is proportional to output and whose scale may in turn be limited by feedstock logistics. Each therefore has its own economic optimum, and running the framework across a range of capacities would show whether the ranking holds at the size each route would be built at.

Extend the environmental assessment beyond carbon. Environmental performance is assessed here as cradle-to-gate GHG intensity, with material efficiency alongside it as a diagnostic, and the two bio-derived feedstocks follow different biogenic conventions, which makes the bio-benzene cases conservative on carbon. Adding land and water use, eutrophication and toxicity would test whether the biomass route's lead survives the categories where a fermentation-derived feedstock would be expected to carry burdens a carbon-only comparison cannot see, and the same assessment on bio-derived benzene would show whether converting an existing plant to non-fossil feedstock is a cheaper way out of fossil carbon than building an alternative route at all.

Carry the inputs as time paths rather than tiers. Only grid carbon intensity, the carbon price and the effects of cell degradation are carried as paths over the plant lifetime; the remaining inputs are held at scenario-tier values, so their within-lifetime variation is bracketed by the sweep rather than represented. Free allocation is one of them, and it is legislated only as far as 2030, so the comparison rests on a bound where the rules beyond that date are not yet set. Giving every input its own trajectory, and repeating the comparison as those rules resolve, would show how the routes separate over a plant's life rather than at a set of fixed points.

Cost the routes on route-specific financing. Capital is estimated from published correlations [64] at the accuracy the AACE Class-4 band brackets, and discounted at a single 8 % real rate, which is a technology-neutral basis for

comparing routes rather than an assumption about any of them. A project-finance treatment, with the higher rate an emerging-technology project would attract, would show what an investor rather than a comparison sees.

Extend and update the geographical scope. Only the ARRRA countries are evaluated, together with a near-decarbonised supply represented by France's observed grid intensity, and of those countries only electricity price, grid intensity and the policy attached to them are varied, labour, land and other site-dependent costs being held common. Extending the comparison would test it against sites with different cost structures and policy environments, and the grid trajectories are worth refreshing as well, Belgium's resting on a scenario that predates the 2025 reversal of its nuclear phase-out.

Establish what abatement achieves in service. Destruction levels are taken from reported performance, which does not show how far they hold over a plant lifetime, nor what sustaining them costs. Settling both would fix what the alternatives have to beat, and that benchmark is worth re-establishing as it moves, decarbonising European steam crackers taking carbon off fossil benzene as it proceeds.

Establish what a buyer can verify about its nitric acid. The conventional route's carbon standing turns on the upstream emissions of purchased nitric acid, taken here from published factors and scenario tiers. Whether a European buyer can establish what a given supplier's plant emits, verify it, and switch supplier on that basis is a question of documentation and contracting rather than of process, and it decides how much of the spread between the supply cases is available in practice.

Obtain cell performance and cost data at pilot scale. All cell-performance data come from laboratory-scale, small-area electrodes, and the one electro-oxidation durability study available [39] pairs high initial Faradaic efficiency with rapid decline without separating the chemistry from early-stage design; cell capital cost is estimated rather than quoted. Pilot data would settle whether initial performance and degradation trade off, and a quotation for a stack at this scale would replace the estimate its capital cost rests on; together they are one of the most important extensions for the feasibility of either electrified route.

Resolve site connection and supply mode. Grid access is modelled as a binary feasibility constraint with cost captured through the country-specific extra-high-voltage tariff, leaving connection queuing, substation-reinforcement contributions, curtailment exposure, port and pipeline access and counterparty supply security unquantified, and the electrified routes charged the national grid intensity throughout. Establishing whether a given site can be connected at all, and whether it can support a power purchase agreement or a dedicated generator with storage, is the single change most likely to move their carbon standing.

Refine the load-following mode and cost the supply alternatives. The load-following sensitivity assumes the flexible load curtails and restores without a ramp-rate constraint or part-load penalty, in common with the closest prior ARRRA treatment [91], and assumes perfect foresight, the threshold being set on the realised year rather than on a forecast; cycling is treated asymmetrically, the electro-oxidation start-stop cost being an AEM-cell analogue rather than a measurement on this cell. Adding forecast error, ramp constraints and the labour, logistics and balancing-market terms would establish how much of the reported saving survives. Following the price signal is also only one way of dealing with a variable supply. Costing a dedicated generator instead, oversized so that it covers the load for more of the year, turns on how its capital is allocated between the electricity the process takes and the surplus it cannot use, and on what that surplus is assumed to earn; settling that allocation is what would make the two supply modes comparable.

Determine the bio-derived feedstock costs at commercial scale. No merchant supply exists at this scale for any of the non-fossil feedstocks, so none of their prices is read from a market: the muconic-acid broth, which decides the biomass route on its own, is modelled from a published minimum selling price [51], and bio-derived benzene is triangulated from proxies on the catalytic fast-pyrolysis route alone. A bottom-up cost study of fermentation and recovery at commercial scale, and supplier quotations or offtake terms where a producer can give them, would replace the single modelled figure the biomass route's cost standing rests on.

Compare the policy instruments against the abatement they trigger. Which instrument would deliver the abatement that §5.3 identifies as the sector's largest reduction, a price signal on imported adipic acid, purchaser-led agreements of the kind the Nitric Acid Climate Action Group runs [33], or procurement standards that specify abated supply, is a policy question this study brackets rather than answers; each would be measured against the cost of the abatement it is meant to trigger.

Bibliography

- [1] United States Environmental Protection Agency. *Background Report AP-42 Section 6.2: Adipic Acid Production*. Tech. rep. Prepared by Pacific Environmental Services, Inc. Research Triangle Park, NC: US EPA, Office of Air Quality Planning and Standards, July 1994. URL: <https://www.epa.gov/sites/default/files/2020-10/documents/b06s02.pdf>.
- [2] Stijn Van de Vyver and Yuriy Román-Leshkov. “Emerging catalytic processes for the production of adipic acid”. In: *Catalysis Science & Technology* 3.6 (2013), pp. 1465–1479. doi: 10.1039/C3CY20728E.
- [3] Mordor Intelligence. *Adipic Acid Market Size & Share Analysis — Growth Trends and Forecast (2026–2031)*. Market report. Accessed August 2026. 2026. URL: <https://www.mordorintelligence.com/industry-reports/adipic-acid-market>.
- [4] Grand View Research. *Adipic Acid Market Size, Share & Trends Analysis Report, 2030*. 2024. URL: <https://www.grandviewresearch.com/industry-analysis/adipic-acid-market> (visited on 11/01/2024).
- [5] ICIS. *Adipic Acid (AA) Market Overview*. ICIS Chemical Market Report. Accessed March 2026. 2023.
- [6] L. Schneider et al. *A Defossilised EU Petrochemical Production System: Techno-Economic Implications*. Tech. rep. Sci4Climate.NRW, 2024.
- [7] IPCC. *Climate Change 2021: The Physical Science Basis. Contribution of Working Group I to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change*. Cambridge University Press, 2021.
- [8] Rebecca Dowling et al. *Harnessing Power-to-X for a Sustainable Carbon Future: Prospects for the ARRRR Cluster*. Tech. rep. TNO/VoltaChem, 2025. URL: <https://publications.tno.nl/publication/34644893/Wr7CZxV7/dowling-2025-harnessing.pdf>.
- [9] Edward Yeo. “Adipic Acid Production: Process Modeling of the Benchmark and Two Electrochemical Approaches. A Process Systems Modeling Perspective”. Faculty of Mechanical Engineering. MSc thesis. Delft, The Netherlands: Delft University of Technology, 2025. URL: <https://resolver.tudelft.nl/uuid:ce927455-d68f-4b7f-80ef-3638ac2d0780> (visited on 09/10/2026).
- [10] International Energy Agency. *Oil Market Report 2023*. Tech. rep. IEA, 2023.
- [11] Simone Tagliapietra. *How will the Iran conflict hit European energy markets?* Bruegel First Glance. Accessed March 2026. Mar. 2026. URL: <https://www.bruegel.org/first-glance/how-will-iran-conflict-hit-european-energy-markets>.
- [12] Ali Hasanbeigi and Adam Sibal. *Stopping a Super-Pollutant: N₂O Emissions Abatement from Global Adipic Acid Production*. Tech. rep. Global Efficiency Intelligence, Apr. 2023. URL: <https://www.globalefficiencyintel.com/s/Adipic-acid-report-4242023-rev.pdf>.
- [13] Minqi Liang et al. “Four decades of full-scale nitrous oxide emission inventory in China”. In: *National Science Review* 11.3 (2024), nwad285. doi: 10.1093/nsr/nwad285.
- [14] Business Analytiq. *Adipic Acid Price Index*. Online procurement analytics platform. Accessed June 2025. 2025. URL: <https://businessanalytiq.com/procurementanalytics/index/adipic-acid-price-index/>.
- [15] IMARC Group. *Adipic Acid Prices, Trend, Chart, Demand and Supply*. Online market intelligence report. Accessed June 2025. 2025. URL: <https://www.imarcgroup.com/adipic-acid-pricing-report>.
- [16] Alchempro. *Adipic Acid 2024–25: Oversupply, China Dominance, and Bearish Pricing Trends*. Tech. rep. Accessed March 2026. Alchempro, Nov. 2025. URL: <https://www.alchempro.com/industry-article/11087/adipic-acid-2024-25-oversupply-china-dominance-and-bearish-pricing-trends>.

- [17] Andrea Corona et al. "Life cycle assessment of adipic acid production from lignin". In: *Green Chemistry* 20.16 (2018), pp. 3857–3866. doi: 10.1039/c8gc00868j.
- [18] Rio Aryapratama and Matty Janssen. "Prospective life cycle assessment of bio-based adipic acid production from forest residues". In: *Journal of Cleaner Production* 164 (2017), pp. 434–443. doi: 10.1016/j.jclepro.2017.06.222.
- [19] European Commission. *Commission Implementing Regulation (EU) 2025/2287 of 11 November 2025 imposing provisional anti-dumping duties on imports of adipic acid originating in the People's Republic of China*. Official Journal of the European Union, 12 November 2025. Accessed March 2026. 2025. URL: https://eur-lex.europa.eu/eli/reg_impl/2025/2287/oj.
- [20] Jeovanna Ríos et al. "A critical review on the progress and challenges to a more sustainable, cost competitive synthesis of adipic acid". In: *Green Chemistry* 23.9 (2021), pp. 3172–3190. doi: 10.1039/D1GC00638J.
- [21] Jan C. J. Bart and Stefano Cavallaro. "Transiting from Adipic Acid to Bioadipic Acid. 1, Petroleum-Based Processes". In: *Industrial & Engineering Chemistry Research* 54.1 (2015), pp. 1–46. doi: 10.1021/ie5020734.
- [22] Klaus Weissermel and Hans-Jürgen Arpe. *Industrial Organic Chemistry*. 4th ed. Weinheim: Wiley-VCH, 2003.
- [23] M. Larry Campbell. "Cyclohexane". In: *Ullmann's Encyclopedia of Industrial Chemistry*. Weinheim: Wiley-VCH, 2011. doi: 10.1002/14356007.a08_209.pub2.
- [24] Michael Tuttle Musser. "Cyclohexanol and Cyclohexanone". In: *Ullmann's Encyclopedia of Industrial Chemistry*. Weinheim: Wiley-VCH, 2011. doi: 10.1002/14356007.a08_217.pub2.
- [25] Ahmed Abutaleb and Mohammad Ashraf Ali. "A comprehensive and updated review of studies on the oxidation of cyclohexane to produce ketone-alcohol (KA) oil". In: *Reviews in Chemical Engineering* 38.7 (2022), pp. 769–797. doi: 10.1515/revce-2020-0059.
- [26] Michael T. Musser. "Adipic Acid". In: *Ullmann's Encyclopedia of Industrial Chemistry*. Weinheim: Wiley-VCH, 2000. doi: 10.1002/14356007.a01_269.
- [27] Sira Suren et al. "Measurement on the solubility of adipic acid in various solvents at high temperature and its thermodynamics parameters". In: *Fluid Phase Equilibria* 360 (2013), pp. 332–337. doi: 10.1016/j.fluid.2013.10.003.
- [28] Qi Wang et al. "Life cycle assessment for the direct synthesis of adipic acid in microreactors and benchmarking to the commercial process". In: *Chemical Engineering Journal* 234 (2013), pp. 300–311. doi: 10.1016/j.cej.2013.08.059.
- [29] Caliane B. B. Costa, Aline C. da Costa, and Rubens Maciel Filho. "Mathematical Modeling and Optimal Control Strategy Development for an Adipic Acid Crystallization Process". In: *Chemical Engineering and Processing: Process Intensification* 44.7 (2005), pp. 737–753. doi: 10.1016/j.cep.2004.08.004. URL: <https://www.sciencedirect.com/science/article/abs/pii/S025527010400203X>.
- [30] A. Shimizu, K. Tanaka, and M. Fujimori. "Abatement Technologies for N₂O Emissions in the Adipic Acid Industry". In: *Chemosphere – Global Change Science* 2 (2000), pp. 425–434. doi: 10.1016/S1465-9972(00)00024-6.
- [31] Daniel Scholz and Felix Nickel. *Handbook on N₂O Mitigation Experience in Germany/Europe*. Tech. rep. FutureCamp Climate GmbH, commissioned by Deutsche Gesellschaft für Internationale Zusammenarbeit (GIZ) and the NACAG initiative, Nov. 2022. URL: https://climatecooperation.cn/wp-content/uploads/2023/04/N2O_handbook_EN.pdf.
- [32] European Commission. *Update of benchmark values for the years 2021–2025: benchmark curve factsheets*. DG Climate Action. Accessed August 2026. 2021. URL: https://climate.ec.europa.eu/system/files/2021-10/policy_ets_allowances_bm_curve_factsheets_en.pdf.

- [33] Wolfram Jörß, Sylvie Ludig, and Lambert Schneider. *Mitigation potentials for emissions of nitrous oxide from chemical industry in industrialised countries world-wide*. Tech. rep. Berlin: Öko-Institut e.V., commissioned by GIZ for the Nitric Acid Climate Action Group (NACAG), Mar. 2023. URL: <https://www.oeko.de/fileadmin/oekodoc/NACAG-N2O-mitigation-potentials.pdf> (visited on 08/15/2026).
- [34] Lambert Schneider, Michael Lazarus, and Anja Kollmuss. *Industrial N₂O Projects Under the CDM: Adipic Acid, A Case of Carbon Leakage?* Working Paper WP-US-1006. Stockholm Environment Institute, 2010. URL: <https://mediamanager.sei.org/documents/Publications/Climate/sei-adipicacidleakage-9oct2010.pdf>.
- [35] Heike Mainhardt and Dina Kruger. “N₂O Emissions from Adipic Acid and Nitric Acid Production”. In: *IPCC Good Practice Guidance and Uncertainty Management in National Greenhouse Gas Inventories*. Intergovernmental Panel on Climate Change (IPCC), 2001.
- [36] Ascend Performance Materials. *Ascend reaches GHG abatement milestone*. Press release, 1 December 2023. Dec. 2023. URL: <https://www.ascendmaterials.com/news/ascend-reaches-ghg-abatement-milestone>.
- [37] F. Kapteijn, J. Rodriguez-Mirasol, and J. A. Moulijn. “Heterogeneous Catalytic Decomposition of Nitrous Oxide”. In: *Applied Catalysis B: Environmental* 9 (1996), pp. 25–64. doi: 10.1016/0926-3373(96)00016-1.
- [38] R. W. van den Brink. *Combined Catalytic Removal of N₂O and NO_x from the Tail Gases of a Nitric Acid Plant*. Tech. rep. Energy Research Centre of the Netherlands (ECN), 2001.
- [39] Xiang Liu et al. “Electrosynthesis of adipic acid with high Faradaic efficiency within a wide potential window”. In: *Nature Communications* 15 (2024), p. 7685. doi: 10.1038/s41467-024-51951-0.
- [40] Zhenhua Li et al. “Electrocatalytic synthesis of adipic acid coupled with H₂ production enhanced by a ligand modification strategy”. In: *Nature Communications* 13 (2022), p. 5009. doi: 10.1038/s41467-022-32769-0.
- [41] Ran Wang et al. “Electrifying adipic acid production: Copper-promoted oxidation and C–C cleavage of cyclohexanol”. In: *Angewandte Chemie* 134 (2022), e202214977. doi: 10.1002/ange.202214977.
- [42] Yingshuai Jia et al. “Directional Electrosynthesis of Adipic Acid and Cyclohexanone by Controlling the Active Sites on NiOOH”. In: *Journal of the American Chemical Society* 146 (2024), pp. 1282–1293. doi: 10.1021/jacs.3c05898.
- [43] Shariful Kibria Nabil et al. “Acid-base chemistry and the economic implication of electrocatalytic carboxylate production in alkaline electrolytes”. In: *Nature Catalysis* 7 (2024), pp. 330–337. doi: 10.1038/s41929-024-01107-6.
- [44] Marcelo Carmo et al. “A Comprehensive Review on PEM Water Electrolysis”. In: *International Journal of Hydrogen Energy* 38.12 (2013), pp. 4901–4934. doi: 10.1016/j.ijhydene.2013.01.151.
- [45] Alexander Buttler and Hartmut Spliethoff. “Current Status of Water Electrolysis for Energy Storage, Grid Balancing and Sector Coupling via Power-to-Gas and Power-to-Liquids: A Review”. In: *Renewable and Sustainable Energy Reviews* 82 (2018), pp. 2440–2454. doi: 10.1016/j.rser.2017.09.003.
- [46] John R. Varcoe, Plamen Atanassov, Dario R. Dekel, et al. “Anion-Exchange Membranes in Electrochemical Energy Systems”. In: *Energy & Environmental Science* 7 (2014), pp. 3135–3191. doi: 10.1039/C4EE01303D.
- [47] Dario R. Dekel. “Review of Cell Performance in Anion Exchange Membrane Fuel Cells”. In: *Journal of Power Sources* 375 (2018), pp. 158–169. doi: 10.1016/j.jpowsour.2017.07.117.
- [48] John S. Newman and Karen E. Thomas-Alyea. *Electrochemical Systems*. 3rd ed. Hoboken, NJ: Wiley-Interscience, 2004. ISBN: 9780471477563.
- [49] Allen J. Bard and Larry R. Faulkner. *Electrochemical Methods: Fundamentals and Applications*. 2nd ed. New York: Wiley, 2001. ISBN: 9780471043720.
- [50] John E. Matthiesen et al. “Electrochemical Conversion of Biologically Produced Muconic Acid: Key Considerations for Scale-Up and Corresponding Technoeconomic Analysis”. In: *ACS Sustainable Chemistry & Engineering* 4.12 (2016), pp. 7098–7109. doi: 10.1021/acssuschemeng.6b01981.

- [51] Sekgetho C. Mokwatlo et al. "Bioprocess development and scale-up for *cis,cis*-muconic acid production from glucose and xylose by *Pseudomonas putida*". In: *Green Chemistry* 26 (2024), pp. 10152–10167. doi: 10.1039/D4GC03424D.
- [52] Derek R. Vardon et al. "Adipic Acid Production from Lignin". In: *Energy & Environmental Science* 8 (2015), pp. 617–628. doi: 10.1039/C4EE03230F.
- [53] Miguel Suástegui et al. "Combining Metabolic Engineering and Electrocatalysis: Application to the Production of Polyamides from Sugar". In: *Angewandte Chemie International Edition* 55.7 (2016), pp. 2368–2373. doi: 10.1002/anie.201509653.
- [54] K. M. Draths and J. W. Frost. "Environmentally Compatible Synthesis of Adipic Acid from D-Glucose". In: *Journal of the American Chemical Society* 116.1 (1994), pp. 399–400. doi: 10.1021/ja00080a057.
- [55] Han-Na Lee et al. "Corynebacterium Cell Factory Design and Culture Process Optimization for Muconic Acid Biosynthesis". In: *Scientific Reports* 8.1 (2018), p. 18041. doi: 10.1038/s41598-018-36320-4.
- [56] J. E. Matthiesen et al. "Electrochemical Conversion of Muconic Acid to Biobased Diacid Monomers". In: *ACS Sustainable Chemistry & Engineering* 4.6 (2016), pp. 3575–3585. doi: 10.1021/acssuschemeng.6b00679.
- [57] Marco Nazareno Dell'Anna et al. "Electrochemical hydrogenation of bioprivileged *cis,cis*-muconic acid to *trans*-3-hexenedioic acid: from lab synthesis to bench-scale production and beyond". In: *Green Chemistry* 23 (2021), pp. 6456–6468. doi: 10.1039/d1gc02225c.
- [58] Sofia Capelli et al. "Bio-adipic acid production by catalysed hydrogenation of muconic acid in mild operating conditions". In: *Applied Catalysis B: Environmental* 218 (2017), pp. 220–229. doi: 10.1016/j.apcatb.2017.06.060.
- [59] Sofia Capelli et al. "Bio Adipic Acid Production from Sodium Muconate and Muconic Acid: A Comparison of two Systems". In: *ChemCatChem* 11 (2019), pp. 3075–3084. doi: 10.1002/cctc.201900343.
- [60] International Organization for Standardization. *ISO 14040:2006 – Environmental management – Life cycle assessment – Principles and framework*. Geneva, Switzerland, 2006.
- [61] International Organization for Standardization. *ISO 14044:2006 – Environmental management – Life cycle assessment – Requirements and guidelines*. Geneva, Switzerland, 2006.
- [62] Iris Vural-Gürsel et al. "Improving energy efficiency of process of direct adipic acid synthesis in flow using pinch analysis". In: *Industrial & Engineering Chemistry Research* 52.23 (2013), pp. 7827–7835. doi: 10.1021/ie4002052.
- [63] AACE International. *Cost Estimate Classification System – As Applied in Engineering, Procurement, and Construction for the Process Industries*. Tech. rep. Recommended Practice 18R-97. Morgantown, WV: AACE International, 2020.
- [64] Gavin Towler and Ray K. Sinnott. *Chemical Engineering Design: Principles, Practice and Economics of Plant and Process Design*. 3rd. Oxford, UK: Butterworth-Heinemann / Elsevier, 2022. ISBN: 978-0-12-821179-3. doi: 10.1016/C2019-0-02025-0.
- [65] National Energy Technology Laboratory. *Quality Guidelines for Energy System Studies: Cost Estimation Methodology for NETL Assessments of Power Plant Performance*. Tech. rep. DOE/NETL-2011/1455. U.S. Department of Energy, National Energy Technology Laboratory, 2021.
- [66] Boaz Izelaar et al. "Techno-economic assessment of different small-scale electrochemical NH₃ production plants". In: *Energy & Environmental Science* 17 (2024), pp. 7983–7998. doi: 10.1039/D4EE03299C.
- [67] INVISTA. *INVISTA Celebrates Inauguration of 400,000-ton/year ADN Plant at the Shanghai Chemical Industry Park*. 2022. URL: <https://www.invista.com/news/invista-celebrates-inauguration-of-400-000-ton-year-adn-plant-at-the-shanghai-chemical-industry-park> (visited on 08/15/2026).
- [68] Toray Industries, Inc. *Toray and PTT Global Chemical Agree to Explore Mass Production Technology to Build Supply Structure for Non-Edible Biomass-Derived Nylon*. Accessed February 2025. 2024.
- [69] BASF. *BASF drives forward structural adjustments at the Ludwigshafen site and closes production plants for adipic acid, CDon and CPon*. Press release P-24-269, 29 August 2024. Accessed July 2026. 2024. URL: <https://www.basf.com/global/en/media/news-releases/2024/08/p-24-269>.

- [70] World Bank. *State and Trends of Carbon Pricing 2025*. Washington, DC: World Bank, 2025. ISBN: 978-1-4648-2255-1. URL: <https://documents.worldbank.org/en/publication/documents-reports/documentdetail/099052726140533053>.
- [71] Moutaz Altaghlibi and Giovanni Gentile. *ESG Economist: Scenarios Shaping EU ETS Prices*. Tech. rep. Accessed March 2026. ABN AMRO Group Economics, Sustainability Research, Sept. 2025. URL: <https://www.abnamro.com/research/en/our-research/esg-economist-scenarios-shaping-eu-ets-prices>.
- [72] European Parliament and Council of the European Union. *Directive (EU) 2023/959 of the European Parliament and of the Council of 10 May 2023 amending Directive 2003/87/EC establishing a system for greenhouse gas emission allowance trading within the Union*. Official Journal of the European Union, L 130, pp. 134–203. Accessed March 2026. 2023. URL: <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:32023L0959>.
- [73] Ariadne Project. *The EU-ETS price through 2030 and beyond: A closer look at drivers, models and assumptions*. Tech. rep. Ariadne Research Project, 2022. URL: https://ariadneprojekt.de/media/2023/01/Ariadne-Dokumentation_ETSWorkshopBruessel_December2022.pdf.
- [74] Robert C. Pietzcker, Sebastian Osorio, and Renato Rodrigues. “Tightening EU ETS targets in line with the European Green Deal: Impacts on the decarbonization of the EU power sector”. In: *Applied Energy* 293 (2021), p. 116914. doi: 10.1016/j.apenergy.2021.116914.
- [75] International Energy Agency. *Net Zero by 2050: A Roadmap for the Global Energy Sector*. Tech. rep. Paris: IEA, 2021. URL: <https://www.iea.org/reports/net-zero-by-2050>.
- [76] Copenhagen Infrastructure Partners. *Charging Ahead: Europe’s Energy System to 2050 – Affordability, Resilience and Clean Energy*. Tech. rep. Copenhagen: Copenhagen Infrastructure Partners, 2026. URL: <https://www.cip.com/media-insights/2050/>.
- [77] European Commission. *Commission Implementing Regulation (EU) 2026/1412 of 26 June 2026 determining revised benchmark values for free allocation of emission allowances for the period from 2026 to 2030 pursuant to Article 10a(2) of Directive 2003/87/EC of the European Parliament and of the Council*. Official Journal of the European Union, L 2026/1412, 29 June 2026. Accessed August 2026. 2026. URL: <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:32026R1412>.
- [78] European Commission. *Commission Implementing Regulation (EU) 2021/447 of 12 March 2021 determining revised benchmark values for free allocation of emission allowances for the period from 2021 to 2025 pursuant to Article 10a(2) of Directive 2003/87/EC of the European Parliament and of the Council*. Official Journal of the European Union, L 87, 15 March 2021, pp. 29–50. Accessed March 2026. 2021. URL: <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:32021R0447>.
- [79] European Commission. *Commission Delegated Decision (EU) 2019/708 of 15 February 2019 supplementing Directive 2003/87/EC of the European Parliament and of the Council concerning the determination of sectors and subsectors deemed at risk of carbon leakage for the period 2021 to 2030*. Official Journal of the European Union, L 120, 8 May 2019, pp. 20–25. Accessed March 2026. 2019. URL: https://eur-lex.europa.eu/eli/dec_del/2019/708/oj.
- [80] European Union. *Regulation (EU) 2023/956 of the European Parliament and of the Council of 10 May 2023 establishing a Carbon Border Adjustment Mechanism*. Official Journal of the European Union, L 130, 16 May 2023, pp. 52–104. 2023. URL: <https://eur-lex.europa.eu/eli/reg/2023/956/oj/eng>.
- [81] European Commission. *Report from the Commission to the European Parliament and the Council pursuant to Article 30 of Regulation (EU) 2023/956 establishing a Carbon Border Adjustment Mechanism*. Tech. rep. COM(2025) 783 final. Accessed March 2026. Brussels: European Commission, Dec. 2025. URL: <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=COM:2025:783:FIN>.
- [82] European Commission. *Proposal for a Regulation of the European Parliament and of the Council amending Regulation (EU) 2023/956 as regards the extension of its scope to downstream goods and anti-circumvention measures*. COM(2025) 989 final, Brussels, 17 December 2025. Dec. 2025. URL: <https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=COM:2025:989:FIN>.

- [83] European Commission. *Proposal for a Directive of the European Parliament and of the Council amending Directive 2003/87/EC and Decision (EU) 2015/1814 as regards driving competitiveness and cost-effective decarbonisation*. Tech. rep. COM(2026) 616 final. Brussels: European Commission, July 2026. URL: <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=COM:2026:616:FIN>.
- [84] S&P Global Commodity Insights. *BASF to cease adipic acid production in Ludwigshafen, Germany*. News, 30 August 2024. Accessed July 2026, 2024. URL: <https://www.spglobal.com/commodity-insights/en/news-research/latest-news/chemicals/083024-basf-to-cease-adipic-acid-production-in-ludwigshafen-germany-spokesperson>.
- [85] TenneT TSO. *Netcapaciteitskaart*. Live Dutch HV grid capacity map. 2025. URL: <https://www.tennet.eu/nl/de-elektriciteitsmarkt/congestiemanagement/netcapaciteitskaart>.
- [86] T. Dubbeling and S. H. Tindemans. “Congestion Management in the Day-Ahead Timeframe: Lessons From the Netherlands”. In: *Proceedings of the 21st International Conference on the European Energy Market (EEM)*. IEEE, 2025.
- [87] Elia and Fluvius. *Joint Action Plan: Congestion Management for Businesses*. Joint press release, 27 March 2025, 2025. URL: https://www.elia.be/en/press/2025/03/20250327_congestion-management-for-businesses.
- [88] Amprion GmbH. *Market Report 2024*. Tech. rep. Amprion GmbH, 2024. URL: https://www.amprion.net/Dokumente/Strommarkt/Marktbericht/2024/amprion_market_report_2024.pdf.
- [89] Maximilian Cegla et al. “Flexible Process Operation for Electrified Chemical Plants”. In: *Current Opinion in Chemical Engineering* 39 (2023), p. 100898. doi: 10.1016/j.coche.2023.100898.
- [90] Svenja Bielefeld, Miloš Cvetković, and Andrea Ramírez. “Should We Exploit Flexibility of Chemical Processes for Demand Response? Differing Perspectives on Potential Benefits and Limitations”. In: *Frontiers in Energy Research* 11 (2023), p. 1190174. doi: 10.3389/fenrg.2023.1190174.
- [91] J. Luo et al. “Impact of intermittent electricity supply on a conceptual process design for microbial conversion of CO₂ into hexanoic acid”. In: *Chemical Engineering Research and Design* 205 (2024), pp. 364–375. doi: 10.1016/j.cherd.2024.04.005.
- [92] M. Kashanian, S. Tohidi, and S. M. Ryan. “Viability of a distributed manufacturing network for chemicals from biomass using local renewable-powered electrochemistry”. In: *The Engineering Economist* (2025). doi: 10.1080/0013791X.2025.2541925.
- [93] Ember. *European Wholesale Electricity Price Data*. Dataset derived from the ENTSO-E Transparency Platform, hourly day-ahead prices. 2025. URL: <https://ember-energy.org/data/european-wholesale-electricity-price-data/>.
- [94] Electricity Maps. *Hourly Carbon Intensity Data, 2024 (lifecycle, consumption-based)*. Dataset, NL/DE/BE/FR; underlying sources ned.nl, ENTSO-E, RTE. 2025. URL: <https://portal.electricitymaps.com>.
- [95] Ember. *European Electricity Review 2025*. Ember. 2025. URL: <https://ember-energy.org/latest-insights/european-electricity-review-2025/>.
- [96] August Gladik, Marc Riedel, and Rüdiger-A. Eichel. “Anion exchange membrane electrolysis at work—Investigating impact of starting parameters and start-stop operation on cold start behavior and degradation”. In: *Journal of Power Sources* (2025). doi: 10.1016/j.jpowsour.2024.235878.
- [97] Simone Mucci, Alexander Mitsos, and Dominik Bongartz. “Cost-optimal power-to-methanol: flexible operation or intermediate storage?” In: *Journal of Energy Storage* 72 (2023), p. 108614. doi: 10.1016/j.est.2023.108614.
- [98] Ammonia Energy Association. *Flexible ammonia synthesis: shifting the narrative around hydrogen storage*. 2023. URL: <https://ammoniaenergy.org/articles/flexible-ammonia-synthesis-shifting-the-narrative-around-hydrogen-storage/>.
- [99] Eurostat. *Electricity prices for non-household consumers – bi-annual data (from 2007 onwards) [nrg_pc_205]*. Eurostat online data browser, custom extraction. 2024. URL: https://ec.europa.eu/eurostat/databrowser/view/nrg_pc_205/default/table.


- [100] E-Bridge Consulting. *Electricity Cost Assessment for Large Industry in the Netherlands, Belgium, Germany and France – Full Report*. Tech. rep. E-Bridge Consulting GmbH, 2024. URL: <https://www.rijksoverheid.nl/documenten/rapporten/2024/04/03/e-bridge-report-electricity-cost-assessment-for-large-industry-in-the-netherlands-belgium-germany-and-france-full-report>.
- [101] Sampath Gunukula and Robert P. Anex. “Techno-economic analysis of multiple bio-based routes to adipic acid”. In: *Biofuels, Bioproducts and Biorefining* 11.5 (2017), pp. 897–907. doi: 10.1002/bbb.1797.
- [102] M. K. Patel et al. *Medium and Long-Term Opportunities and Risks of the Biotechnological Production of Bulk Chemicals from Renewable Resources — The BREW Project*. Tech. rep. Utrecht University / European Commission DG Research, 2006. doi: 10.24406/publica-fhg-293032.
- [103] BusinessAnalytiq. *Commodity Price Indexes*. Online procurement analytics platform. 2026. URL: <https://businessanalytiq.com/procurementanalytics/>.
- [104] David Miller, Katy Armstrong, and Peter Styring. “Assessing methods for the production of renewable benzene”. In: *Sustainable Production and Consumption* 32 (2022), pp. 184–197. doi: 10.1016/j.spc.2022.04.019.
- [105] B. Linnhoff and E. Hindmarsh. “The Pinch Design Method for Heat Exchanger Networks”. In: *Chemical Engineering Science* 38.5 (1983), pp. 745–763. doi: 10.1016/0009-2509(83)80185-7.
- [106] ASHRAE. *ANSI/ASHRAE/IES Standard 90.1-2019, Energy Standard for Buildings Except Low-Rise Residential Buildings*. Tech. rep. Peachtree Corners, GA: American Society of Heating, Refrigerating and Air-Conditioning Engineers, 2019. URL: https://www.ashrae.org/file%20library/technical%20resources/standards%20and%20guidelines/standards%20addenda/90_1_2019_x_20220204.pdf.
- [107] National Institute of Standards and Technology. *Hexanedioic acid (adipic acid), CAS 124-04-9: phase change data*. 2026. URL: <https://webbook.nist.gov/cgi/cbook.cgi?ID=C124049&Mask=4>.
- [108] McKenzie Hubert et al. *Clean Hydrogen Production Cost Scenarios with PEM Electrolyzer Technology*. DOE Hydrogen Program Record 24005. U.S. Department of Energy, Hydrogen and Fuel Cell Technologies Office, May 2024. URL: <https://www.hydrogen.energy.gov/docs/hydrogenprogramlibraries/pdfs/24005-clean-hydrogen-production-cost-pem-electrolyzer.pdf>.
- [109] International Renewable Energy Agency. *Green Hydrogen Cost Reduction: Scaling up Electrolysers to Meet the 1.5°C Climate Goal*. Tech. rep. Abu Dhabi: IRENA, 2020. URL: <https://www.irena.org/publications/2020/Dec/Green-hydrogen-cost-reduction>.
- [110] Tom Wirtanen et al. “Cathodic Corrosion of Metal Electrodes — How to Prevent It in Electroorganic Synthesis”. In: *Chemical Reviews* 121.17 (2021), pp. 10241–10270. doi: 10.1021/acs.chemrev.1c00148.
- [111] David Sopher and Andreas Gieseler. “Electrode Consisting of an Iron-Containing Core and a Lead-Containing Coating”. U.S. patent 5593557 (United States). BASF SE. 1997. URL: <https://patents.google.com/patent/US5593557A/en>.
- [112] Panji B. Tamarona, Thijs J. H. Vlugt, and Mahinder Ramdin. “OpenPyTEA: An open-source Python toolkit for techno-economic assessment of chemical process plants and energy systems with economic sensitivity and uncertainty evaluation”. In: *SoftwareX* 35 (2026), p. 102816. doi: 10.1016/j.softx.2026.102816.
- [113] KPMG. *Cost of Capital Study 2024: The New Dilemma — Balancing Interest Rates and Growth*. Tech. rep. KPMG AG Wirtschaftsprüfungsgesellschaft, 2024. URL: <https://assets.kpmg.com/content/dam/kpmgsites/ch/pdf/cost-of-capital-study-2024.pdf>. coredownload.inline.pdf.
- [114] Hydrogen Tools (HyARC). *Lower and Higher Heating Values of Fuels*. Pacific Northwest National Laboratory for the U.S. Department of Energy. 2024. URL: <https://h2tools.org/hyarc/calculator-tools/lower-and-higher-heating-values-fuels>.
- [115] Chemical Engineering. *The Chemical Engineering Plant Cost Index*. Chemical Engineering magazine, Access Intelligence LLC. 2025. URL: <https://www.chemengonline.com/pci-home/>.

- [116] Chemical Engineering. *The Chemical Engineering Plant Cost Index (January 2010)*. Chemical Engineering magazine, Access Intelligence LLC. 2010. URL: <https://www.chemengonline.com/pci-home/>.
- [117] CHEManager. *Climate-friendly Production*. Dec. 2009. URL: <https://www.chemanager-online.com/en/news/climate-friendly-production>.
- [118] Eurostat. *Producer prices in industry, capital goods, euro area (annual data, index 2021=100)*. 2024. URL: https://ec.europa.eu/eurostat/databrowser/view/sts_inpp_a.
- [119] European Central Bank. *Euro foreign exchange reference rates: US dollar, annual average*. 2024. URL: https://www.ecb.europa.eu/stats/policy_and_exchange_rates/euro_reference_exchange_rates/.
- [120] World Resources Institute and World Business Council for Sustainable Development. *Scope 2 Guidance: An Amendment to the GHG Protocol Corporate Standard*. Tech. rep. World Resources Institute, 2015. URL: <https://ghgprotocol.org/scope-2-guidance>.
- [121] European Parliament and Council. *Directive (EU) 2018/2001 of the European Parliament and of the Council on the promotion of the use of energy from renewable sources (recast)*. Official Journal of the European Union, L 328/82. 2018. URL: <https://eur-lex.europa.eu/legal-content/EN/ALL/?uri=CELEX:32018L2001>.
- [122] European Commission, Directorate-General for Energy and Directorate-General for Climate Action. *EU Reference Scenario 2020: Energy, Transport and GHG Emissions — Trends to 2050*. Tech. rep. MJ-02-21-816-EN-N. Luxembourg: Publications Office of the European Union, 2021. DOI: 10.2833/35750. URL: https://energy.ec.europa.eu/data-and-analysis/energy-modelling/eu-reference-scenario-2020_en.
- [123] A. Hoxha and B. Christensen. “The Carbon Footprint of Fertiliser Production: Regional Reference Values”. In: *Proceedings 805, International Fertiliser Society* (2019). URL: https://www.fertilizerseurope.com/wp-content/uploads/2020/01/The-carbon-footprint-of-fertilizer-production_Regional-reference-values.pdf.
- [124] O. Schmidt et al. “Future cost and performance of water electrolysis: An expert elicitation study”. In: *International Journal of Hydrogen Energy* 42.52 (2017), pp. 30470–30492. DOI: 10.1016/j.ijhydene.2017.10.045.
- [125] Anita H. Reksten et al. “Projecting the future cost of PEM and alkaline water electrolyzers: a CAPEX model including electrolyser plant size and technology development”. In: *International Journal of Hydrogen Energy* 47.90 (2022), pp. 38106–38113. DOI: 10.1016/j.ijhydene.2022.08.306. URL: <https://www.sciencedirect.com/science/article/pii/S0360319922040253>.
- [126] Bio Based Press. *BioBTX collects €80 million for its renewable aromatics technology*. Bio Based Press. June 2024. URL: <https://www.biobasedpress.eu/2024/06/biobtx-collects-e-80-million-for-its-renewable-aromatics-technology/>.
- [127] International Energy Agency. *Global Hydrogen Review 2024*. Tech. rep. Paris: IEA, 2024. URL: <https://www.iea.org/reports/global-hydrogen-review-2024>.
- [128] International Energy Agency. *Northwest European Hydrogen Monitor 2025*. Tech. rep. Paris: IEA, May 2025. URL: <https://www.iea.org/reports/northwest-european-hydrogen-monitor-2025>.
- [129] Autoriteit Consument & Markt. *Ontwikkeling netkosten tot en met 2050 en de kostenverdeling via nettarieven*. Tech. rep. ACM (Netherlands Authority for Consumers and Markets), 2024.
- [130] Agency for the Cooperation of Energy Regulators. *Gas and Electricity Key Developments in Europe – 2025 Report*. Tech. rep. ACER, 2025.
- [131] ENTSO-E. *European Resource Adequacy Assessment 2025 Edition – Executive Report*. Tech. rep. European Network of Transmission System Operators for Electricity, Dec. 2025.
- [132] Federal Planning Bureau. *Energy Outlook for Belgium towards 2050 (May 2024 edition)*. Tech. rep. Brussels: Federal Planning Bureau, May 2024. URL: <https://www.plan.be/en/data/energy-outlook-belgium-towards-2050-may-2024>.

- [133] PBL Netherlands Environmental Assessment Agency. *Climate and Energy Outlook of the Netherlands 2024 (Klimaat- en Energieverkenning 2024)*. Tech. rep. The Hague: PBL, 2024. URL: <https://www.pbl.nl/en/publications/climate-and-energy-outlook-of-the-netherlands-2024>.
- [134] Umweltbundesamt. *Projektionsbericht 2024 für Deutschland — Treibhausgas-Projektionen 2024*. Tech. rep. Dessau-Roßlau: Umweltbundesamt, 2024. URL: <https://www.umweltbundesamt.de/publikationen/treibhausgas-projektionen-2024-ergebnisse-kompakt>.
- [135] Réseau de Transport d'Électricité (RTE). *Bilan Prévisionnel 2023–2035 — Mise à jour de la première période de Futurs énergétiques 2050*. Tech. rep. La Défense: RTE, 2023. URL: <https://www.rte-france.com/actualites/bilan-previsionnel-transformation-systeme-electrique-2023-2035>.
- [136] S&P Global Commodity Insights. *Bio-Chemicals 2025: Insights into the global biochemicals market*. S&P Global Commodity Insights special report. 2025. URL: <https://www.spglobal.com/energy/en/news-research/special-reports/chemicals/bio-chemicals-2025>.
- [137] Emma A. R. Zuiderveen et al. “Evaluating the Environmental Sustainability of Alternative Ways to Produce Benzene, Toluene, and Xylene”. In: *ACS Sustainable Chemistry & Engineering* 12.13 (2024), pp. 5092–5104. doi: 10.1021/acssuschemeng.3c06996.
- [138] Marian Rosental, Thomas Fröhlich, and Axel Liebich. “Life cycle assessment of carbon capture and utilization for the production of large-volume organic chemicals”. In: *Frontiers in Climate* 2 (2020), p. 586199. doi: 10.3389/fclim.2020.586199.
- [139] Business Analytiq. *Benzene Price Index*. Online procurement analytics platform. Accessed May 2026. 2026. URL: <https://businessanalytiq.com/procurementanalytics/index/benzene-price-index/>.
- [140] U.S. Energy Information Administration. *Crude oil and petroleum product prices increased sharply in the first quarter of 2026*. Today in Energy. 2026. URL: <https://www.eia.gov/todayinenergy/detail.php?id=67424>.
- [141] PlasticsEurope. *Eco-profile: Benzene; production mix, at plant, EU-28*. PlasticsEurope Life Cycle Inventory database, ILCD dataset UUID 38f29fc1-dc9a-4835-b8a4-ed4e1ec20ff7. 2013. URL: <https://plasticseurope.lca-data.com/datasetdetail/process.xhtml?uuid=38f29fc1-dc9a-4835-b8a4-ed4e1ec20ff7>.
- [142] J. G. Vogtländer. *Idemat 2026 database*. Delft University of Technology, Sustainability Impact Metrics. 2026. URL: <https://www.ecocostsvalue.com/data-tools-books/>.
- [143] Petrochemicals Europe. *European Market Overview — Aromatics (BTX) Supply 2024*. Industry association statistical compilation. 2025. URL: <https://www.petrochemistry.eu/about-petrochemistry/petrochemicals-facts-and-figures/european-market-overview/>.
- [144] BioBTX. *BioBTX announces the start of construction on its first-of-a-kind circular aromatics plant*. July 2026. URL: <https://biobtx.com/fid/>.
- [145] Anellotech. *Bio-TCat technology is ready for commercialization*. Accessed 2026-07-17. 2021. URL: <https://www.anellotech.com/blogs/anellotechs-bio-tcat-r-technology-is-ready-for-commercialization>.
- [146] International Energy Agency. *Towards Hydrogen Definitions Based on Their Emissions Intensity*. Tech. rep. Paris: IEA, 2023. URL: <https://www.iea.org/reports/towards-hydrogen-definitions-based-on-their-emissions-intensity>.
- [147] European Hydrogen Observatory. *Hydrogen production costs*. Dataset, reference year 2024, Clean Hydrogen Partnership. 2025. URL: <https://observatory.clean-hydrogen.europa.eu/hydrogen-landscape/production-trade-and-cost/cost-hydrogen-production> (visited on 09/05/2026).
- [148] L. Eblé and M. Weeda. *Evaluation of the levelised cost of hydrogen based on proposed electrolyser projects in the Netherlands*. Tech. rep. TNO 2024 R10766. TNO, May 2024.
- [149] CE Delft and TNO. *Afnameverplichting groene waterstof [Off-take obligation for green hydrogen]*. Tech. rep. Delft: CE Delft, 2023.

- [150] Business Analytiq. *Nitric Acid Price Index*. Online procurement analytics platform. Accessed June 2026. 2026. URL: <https://businessanalytiq.com/procurementanalytics/index/nitric-acid-price-index/>.
- [151] Intratec. *Nitric Acid Prices*. Primary commodity price series. Accessed 2026. 2025. URL: <https://www.intratec.us/solutions/primary-commodity-prices/commodity/nitric-acid-prices>.
- [152] Thunder Said Energy. *Nitric acid: production costs?* Energy research consultancy data-file. 2022. URL: <https://thundersaidenergy.com/downloads/nitric-acid-production-costs/>.
- [153] IMARC Group. *Nitric Acid Prices, Trend, Chart, Demand, Market Analysis, News, Historical and Forecast Data Report 2026 Edition*. Market pricing report. Accessed June 2026. 2026. URL: <https://www.imarcgroup.com/nitric-acid-pricing-report>.
- [154] ChemAnalyst. *Nitric Acid Price Trend and Forecast*. Pricing database. Accessed June 2026. 2026. URL: <https://www.chemanalyst.com/Pricing-data/nitric-acid-1142>.
- [155] Xiang Liu, Amgad Elgowainy, and Michael Wang. "Life cycle energy use and greenhouse gas emissions of ammonia production from renewable resources and industrial by-products". In: *Green Chemistry* 22 (2020), pp. 5751–5761. doi: 10.1039/D0GC02301A. URL: <https://www.osti.gov/servlets/purl/1657142>.
- [156] IPCC. *2019 Refinement to the 2006 IPCC Guidelines for National Greenhouse Gas Inventories. Volume 3: Industrial Processes and Product Use, Chapter 3: Chemical Industry*. Tech. rep. Switzerland: Intergovernmental Panel on Climate Change, 2019. URL: <https://www.ipcc-nggip.iges.or.jp/public/2019rf/index.html>.
- [157] B. Neuffer, N. Frank, and M. Desai. *Available and Emerging Technologies for Reducing Greenhouse Gas Emissions from the Nitric Acid Production Industry*. Tech. rep. Research Triangle Park, NC: U.S. Environmental Protection Agency, Office of Air and Radiation, 2010. URL: <https://www.epa.gov/sites/default/files/2015-12/documents/nitricacid.pdf>.
- [158] I. Garcia-Herrero et al. "Environmental challenges of the chlor-alkali production: Seeking answers from a life cycle approach". In: *Science of the Total Environment* 580 (2017), pp. 147–157. doi: 10.1016/j.scitotenv.2016.10.202.
- [159] CarbonCloud. *Potassium Hydroxide (KOH), E525*. 2026. URL: <https://apps.carboncloud.com/climatehub/product-reports/id/1394351136979>.
- [160] Euro Chlor. *Eco-profile of the Chlor-Alkali Process — September 2022 Update*. Tech. rep. Euro Chlor, 2022. URL: https://www.eurochlor.org/wp-content/uploads/2022/09/Eco_profile-study-update-September-2022.pdf.
- [161] CarbonCloud. *Sulfuric Acid (H₂SO₄), E513*. 2026. URL: <https://apps.carboncloud.com/climatehub/product-reports/id/1266801278346>.
- [162] European Lime Association. *Environmental Product Declaration: Hydrated Lime, EPD-IES-0013450:001 (S-P-13450)*. International EPD System, externally reviewed by Rina Consulting. 2022. URL: <https://www.environdec.com/library/epd13450>.
- [163] A. Laveglia et al. "Hydrated lime life-cycle assessment: Current and future scenarios in four EU countries". In: *Journal of Cleaner Production* 369 (2022), p. 133224. doi: 10.1016/j.jclepro.2022.133224.
- [164] BusinessAnalytiq. *Potassium Hydroxide Price Index*. 2026. URL: <https://businessanalytiq.com/procurementanalytics/index/potassium-hydroxide-price-index/>.
- [165] BusinessAnalytiq. *Sulfuric Acid Price Index*. 2026. URL: <https://businessanalytiq.com/procurementanalytics/index/sulfuric-acid-price-index/>.
- [166] ChemAnalyst. *Calcium Hydroxide Price Trends, Chart, Index and News*. 2025. URL: <https://www.chemanalyst.com/Pricing-data/calcium-hydroxide-1204>.
- [167] Intergovernmental Panel on Climate Change. *2006 IPCC Guidelines for National Greenhouse Gas Inventories, Volume 2: Energy*. Ed. by H. S. Eggleston et al. Chapter 2, Table 2.3: natural gas default of 56 100 kgCO₂/TJ on a net calorific value basis. Hayama, Japan: Institute for Global Environmental Strategies, 2006. URL: <https://www.ipcc-nggip.iges.or.jp/public/2006gl/>.

- [168] Department for Energy Security and Net Zero. *Greenhouse Gas Reporting: Conversion Factors 2026*. Tech. rep. Fuels worksheet, natural gas on a net calorific value basis. London, UK: Department for Energy Security and Net Zero, June 2026. URL: <https://www.gov.uk/government/publications/greenhouse-gas-reporting-conversion-factors-2026>.
- [169] Filipa Franco et al. "A Systematic Performance History Analysis of a Chlor-Alkali Membrane Electrolyser under Industrial Operating Conditions". In: *Applied Sciences* 9.2 (2019), p. 284. doi: 10.3390/app9020284.
- [170] Euro Chlor. *Hydrogen: how made, uses and resources*. 2026. URL: <https://eurochlor.org/products/hydrogen/>.
- [171] Richard Turton et al. *Analysis, Synthesis, and Design of Chemical Processes*. 5th ed. Boston: Prentice Hall, 2018. ISBN: 978-0-13-417740-3.
- [172] Bundesministerium der Justiz. *Stromnetzentgeltverordnung (StromNEV), § 19 Abs. 2 — Individuelle Netzentgelte*. 2005. URL: https://www.gesetze-im-internet.de/stromnev/__19.html.



Model inputs and factors

This chapter collects the inputs the model runs on and their derivations. It is the supporting material behind the literature review: the main text there states each input's level and the choice made for it, and what follows here is the evidence and the reasoning that fix them. The sections run in the order the inputs enter the model. First the electricity tariff build-up for a large extra-high-voltage consumer in the three ARRRRA countries (§A.1); then the grid carbon intensities (§A.2) and the rebased trajectories along which the Mode 4 sensitivity (§3.5.6) projects them (§A.3); the CIP (2026) rising carbon-price trajectory used in the same mode (§A.4); the per-feedstock price and scope-3 emission-intensity anchors (§A.5); the utility greenhouse-gas emission factors (§A.6); and the durability and performance evidence behind the electrochemical cell tiers (§A.7). The consolidated scenario values are in the literature review, Tables 2.5 and 2.6.

The rest of the supporting material follows in its own chapters: the comparison against Yeo (2025) in Appendix C; the heat-exchange networks, and the stream lists they were designed from, in Appendix D; and the load-following analysis in Appendix E. The exchanger-by-exchanger listing behind the networks is printed there in full.

A.1. Electricity prices in ARRRRA

Electricity cost inputs differ by route because the conventional and electrified routes operate at different consumption scales and therefore face different tariff structures. The electrified routes' implied annual demand at 50,000–100,000 t/a AA-equivalent falls in the same scale as the E-Bridge (2024) EHV-baseload reference profile (1 TWh/yr, 125 MW, 8 000 FLH), where country-specific relief instruments materially reduce the effective cost; the conventional route is primarily steam-driven and its electrical demand falls in a smaller statutory consumption band. This section establishes the relevant price level for each, drawing on Eurostat statutory prices and the E-Bridge benchmark for large EHV-connected industry. Eurostat's bi-annual electricity price survey for non-household consumers provides the statutory reference for industrial electricity prices by consumption band [99]. For the largest band, band IG (≥ 150 GWh/a), the 2024 annual average excluding VAT was €130/MWh in Germany, €110/MWh in the Netherlands, and €84/MWh in Belgium, against an EU-27 average of €113/MWh. These values reflect a post-crisis stabilisation following the 2022 peak (€190–225/MWh across ARRRRA) and are confirmed as representative by H1 2025 Eurostat data [99]. They represent invoiced prices aggregated across all voltage levels, before country-specific EHV relief is applied.

For extra-high-voltage (EHV) connected baseload consumers the effective cost is substantially lower. E-Bridge (2024), commissioned by the Dutch Ministry of Economic Affairs (EZK) and submitted to Parliament, provides an explicit benchmark for this consumer profile [100]. After applying each country's relief instruments, effective 2024 costs are €46/MWh in Germany, €56/MWh in Belgium, and €95/MWh in the Netherlands (Table A.1).

E-Bridge notes that fixed EHV charges such as the Dutch *vastrecht* become negligible above 500 GWh/a consumption [100]. At this consumption scale, the chemicals sector is eligible in Germany and Belgium for Indirect Cost Compensation (ICC), a national State Aid measure approved under EU guidelines that partially offsets the ETS-related carbon cost embedded in electricity prices; the Netherlands terminated its equivalent scheme in 2023.

ICC is listed separately from commodity and non-commodity costs in Table A.1 [100].¹ The conventional route's substantially lower electrical demand places it in a different consumption band, addressed below.

Table A.1: All-in effective electricity costs for large EHV-connected baseload industrial users (1 TWh/a, 125 MW, 8 000 FLH reference profile) across the three ARRA countries: 2024 observed values and E-Bridge 2030 projections. Total = Commodity + Non-commodity + ICC, non-commodity net of all applicable relief before ICC; ICC is shown negative, as a saving. Source: E-Bridge [100].

Country	Commodity ¹ (€/MWh, 2024)	Non-comm. (€/MWh, 2024)	ICC relief (€/MWh)	Total 2024	Total 2030 (EP scenario)	Key policy instrument
Germany	75.9	5.3	−35.3	46	49 <i>EP-Low</i>	§19 StromNEV relief (≥7 000 FLH baseload)
Netherlands	74.7	19.9	—	95	92 <i>EP-High</i>	No ICC since 2023; highest network charges
Belgium	73.5	7.1	−25.0	56	62 <i>EP-Central</i>	ICC for chemicals; Flanders certificate relief




Figure A.1: All-in electricity cost breakdown for large EHV-connected baseload consumers across ARRA countries, 2024 and 2030 (E-Bridge projection). Bars stack non-commodity and commodity costs, with ICC/policy relief as a downward green segment; the line marks the net all-in cost. Reference profile: 1 TWh/a, 125 MW, 8 000 FLH EHV consumer; values in Table A.1. Source: E-Bridge [100].

Figure A.1 illustrates the cost composition: Germany's low effective cost is driven primarily by §19 StromNEV network charge relief and ICC (the downward segment), while the Netherlands carries the full non-commodity burden with no equivalent reduction. The 2030 projections preserve this ranking, with net all-in costs broadly stable as falling commodity prices are largely offset by rising network charges and shrinking ICC relief [100, 129]. These 2030 values anchor the electricity price scenarios defined below.

Country-level cost differences and uncertain commodity and network charge trajectories make a single deterministic electricity price an unreliable basis for assessment [130, 131]. For the electrochemical routes, three scenarios are anchored to the E-Bridge 2030 projections: EP-Low (€49/MWh, Germany), EP-Central (€62/MWh, Belgium), and EP-High (€92/MWh, Netherlands), all with full EHV relief applied [100]. A supplementary scenario, EP-Statutory, uses Eurostat band IG statutory prices without EHV-scale relief (€130/MWh Germany, €110/MWh Netherlands, €84/MWh Belgium [99]), providing an upper-bound reference broadly consistent with Yeo (2025). The gap between the statutory prices and the relieved EHV tiers is the country-specific relief, which is large enough to reorder the three countries: Germany has the most relief (§19 network reduction and ICC) and so moves from the most expensive country on statutory prices to the cheapest once relief is applied, while the Netherlands, which

¹The 2024 commodity figure combines actual day-ahead prices for January and February 2024 with February 2024 forwards for the rest of the year; the 2030 totals take a projected commodity of ≈€65/MWh for all three countries, on European wholesale market convergence [100].

ended its ICC scheme in 2023 and receives neither, becomes the most expensive. Belgium, with ICC only, sits between them.

For the conventional route, the process is primarily steam-driven [9] with a small electrical demand (~6 MW, 49.9 GWh/a at 8 000 FLH). Spreading fixed network charges such as the Dutch *vastrecht* over fewer MWh raises its per-unit cost, though its lower capacity requirement may partly offset this. As an energy-intensive chemicals consumer it also qualifies for the same relief instruments as the electrified routes (ICC and the German §19 Strom-NEV network relief), which hold its effective price per MWh close to the relieved E-Bridge figures above rather than at the unrelieved statutory level. All routes are therefore priced at that relieved level (Table A.1), with the effect of a higher electricity price on competitiveness tested through Yeo (2025)'s €145/MWh sensitivity [9].

France is the fourth country in the same E-Bridge assessment, on the same reference profile and the same relieved basis, and its €48/MWh is taken from there [100]. Why it is carried at all is set out with the grid intensities in §A.2.

A.2. Grid carbon intensity

The tier values quoted in §2.4.3 are annual means of Electricity Maps 2024 hourly data on a lifecycle, consumption-based basis [94]. The three ARRR countries occupy structurally distinct positions: Belgium 150 gCO₂e/kWh (nuclear and offshore wind, with residual gas), the Netherlands 261 gCO₂e/kWh (gas-dominated dispatch), and Germany 341 gCO₂e/kWh (residual coal after the nuclear exit), a spread of more than a factor of two, against an EU-27 average of 213 gCO₂e/kWh [95], with direct consequences for the GHG balance of electrified routes (Table A.2).

Table A.2: Grid carbon intensity of the ARRR countries and France, 2024 (gCO₂e/kWh, lifecycle, consumption-based). Values are annual means of Electricity Maps 2024 hourly data [94]; they anchor the three core scenario tiers, and France's value defines the GI-Renewable supplementary scenario. The EU-27 average is shown as a reference [95].

Country	2024 CI	Scenario tier
Belgium	150	<i>GI-Low</i>
Netherlands	261	<i>GI-Central</i>
Germany	341	<i>GI-High</i>
EU-27	213	<i>reference</i>
France	31	<i>GI-Renewable</i>

The prior electrochemical AA TEA by Yeo (2025) assumed 50 gCO₂e/kWh [9], well below any ARRR country's 2024 grid intensity, understating scope-2 emissions regardless of plant siting.

All three countries are on documented downward decarbonisation trajectories, with the ranking Belgium < Netherlands < Germany robust across the ENTSO-E European Resource Adequacy Assessment (ERAA) 2025 projection scenarios [131]. A fourth scenario tier, *GI-Renewable*, bounds the analysis from below using France's 2024 observed lifecycle intensity of 31 gCO₂e/kWh [94] as an empirical anchor for predominantly low-carbon supply, grounding the renewable-electricity case in an observed European grid rather than an assumed value. In the location analysis France also supplies its own electricity price of €48/MWh (§A.1), so that case is a decarbonised supply on both counts rather than a clean grid at an assumed cost.

To capture the potential impact of grid decarbonisation on route scope-2 GHG intensity over the plant's 20-year lifetime, grid carbon intensity is treated as time-dependent in a supplementary sensitivity case, complementing the 2024-static base case anchored on Table A.2. To the author's knowledge, the EU Reference Scenario 2020 (REF2020) [122] is the most recent published source providing country-level grid carbon intensity forecasts at Member State resolution to 2050. Its trajectories form the basis of the time-dependent sensitivity, rebased onto the 2024 lifecycle anchor; the full procedure is set out in Section 3.5.1.

REF2020 was finalised in 2020 and predates recent policy changes, notably Belgium's 2024 nuclear lifetime extension and the planned post-2035 gas-fired backfill of the nuclear phase-out. Cross-checked against post-2022 national projections [132, 133, 134, 135], three of the four trajectories broadly support the REF2020-rebased shape; FPB (2024) is the partial exception, projecting Belgian production-basis CI essentially flat at ~145 gCO₂e/kWh through 2050 (gas-fired backfill of the post-2035 nuclear phase-out) versus the rebased trajectory's decline to 129 gCO₂e/kWh. The static GI-Low tier (150 gCO₂e/kWh, held flat over the plant lifetime) brackets this divergence. The full per-country trajectories follow in §A.3.

A.3. Rebased grid carbon intensity trajectories

The supplementary time-dependent sensitivity defined in Section 3.5.1 (Eq. (3.10)) projects each country's grid carbon intensity forward by rebasing the EU Reference Scenario 2020 (REF2020) PRIMES country trajectory [122] onto its 2024 lifecycle anchor [94]; intermediate plant years are linearly interpolated between the five-year REF2020 nodes (2030 / 2035 / 2040 / 2045 / 2050). The resulting trajectories are shown in Figure A.2 and the reference-year values listed in Table A.3; both are cross-checked against post-2022 national projections (§A.2).




Figure A.2: Country-specific grid carbon intensity (gCO₂e/kWh, lifecycle, consumption-based) used in the supplementary time-dependent sensitivity. Solid lines: rebased REF2020 trajectories anchored on each country's 2024 lifecycle value [94]. Dashed lines: 2024 lifecycle static baseline held flat in the base case.

Table A.3: Per-country grid CI trajectory values (gCO₂e/kWh, lifecycle, consumption-based) at the REF2020 reference years, anchored on each country's 2024 lifecycle value [94]. France stays near the decarbonised floor across the horizon, where the REF2020 projection is flat and mildly non-monotonic: its 2050 value edges up to 23 from 22 gCO₂e/kWh rather than falling further, following the reference series rather than any artefact of the rebasing.

Country	2024	2030	2035	2040	2045	2050
Belgium	150	154	148	136	133	130
Netherlands	261	129	122	103	82	79
Germany	341	276	161	100	79	78
France	31	27	25	26	22	23

A.4. Carbon price trajectory

The Mode 4 time-dependent sensitivity on the carbon price applies the CIP (2026) Net Zero rising trajectory [76] against the static CP-Low / Central / High tiers used in the other evaluation modes. Figure A.3 shows the Net Zero trajectory (operationally used in Mode 4) together with the Competitive & Resilient and Slow Transition trajectories from the same CIP report (shown for comparison) and the three static thesis tiers (dashed). Static tier values are anchored to the literature ranges discussed in LR §2.4.1; the rising-trajectory rationale is in the same section.

The CIP industry-scenario trajectories plotted in Figure A.3 were reconstructed by graphical extraction from the published CIP (2026) figure [76]. The Net Zero trajectory was digitised at three reference years (2030, 2040, 2050), yielding anchor values of approximately €142, €185, and €210 per tCO₂, and linearly interpolated between the anchors with a smoothed transition at the 2040 inflection point; the Competitive & Resilient and Slow Transition trajectories were reconstructed analogously from their respective anchor years. The reproduced curves

are intended as a close approximation of the original CIP trajectories rather than an exact replica: minor deviations from the source figure are possible owing to the limited precision of graphical extraction, and the trajectories should be interpreted as representative of the shape and order of magnitude of the CIP scenarios rather than as exact numerical reproductions.




Figure A.3: Carbon-price trajectories used in this thesis. Solid lines: CIP (2026) industry-scenario trajectories [76]; Net Zero is the rising- CO_2 -price trajectory operationally applied in Mode 4 (§3.5.6), Competitive & Resilient and Slow Transition are shown for comparison. Dashed lines: thesis static carbon-price tiers (CP-Low / Central / High at €75 / 115 / 145 per tCO_2 , held flat across the modelled lifetime in static-tier evaluation modes). The CP-Low tier (€75/ tCO_2) coincides with the CIP Slow Transition trajectory and the two lines overlap visually in the figure.

A.5. Feedstock treatment

The five feedstock blocks below give the price and scope-3 emission-intensity anchors carried in the scenario framework, the source of each and the construction of each band. The framing, the common pricing basis and the choice between discrete procurement scenarios and continuous sensitivity ranges are in §2.4.6; the resulting values are collected in Table 2.5.

A.5.1. Benzene

Benzene is a primary purchased feedstock for the CON and ECO routes. It is evaluated as a binary procurement choice between fossil-derived (BEN-Fossil) and bio-attributed (BEN-Bio) supply, pairing price and scope-3 emission intensity per scenario (Table A.4).

Table A.4: Benzene procurement scenarios: BEN-Fossil and BEN-Bio price tiers (\$/kg; low/central/high), scope-3 CI, and supply pathway. Prices from a 2024–2025 European spot average [136, 9] (fossil) and BioBTX (2024) [126] with S&P Global (bio); scope-3 CI from Zuiderveen (2024) [137] and Rosental (2020) [138]. Band construction is given in the text.

Scenario	Price (\$/kg)	Scope-3 CI (kg CO_2 e/kg)	Pathway
BEN-Fossil	0.39 / 0.88 / 1.65	1.87	European naphtha-cracker
BEN-Bio	0.77 / 1.77 / 2.31	1.08	Biomass, catalytic fast pyrolysis (BioBTX/Anellotech family; one of several routes, see text)

Fossil benzene

The BEN-Fossil central of \$0.88/kg is a 2024–2025 European spot average, close to the ICIS-2024 contract price (\$0.89) adopted by Yeo (2025) [9]. Benzene swings sharply year-to-year (a high 2024 near \$1.05/kg against a soft, oversupplied 2025 near \$0.74/kg [136]), so a two-year average is used and the 2026 Strait-of-Hormuz spike (~\$1.2/kg [139, 140]) is excluded. The band spans the observed range, from the 2015–2019 trough (\$0.39/kg)

to the April 2022 crisis peak (\$1.65/kg), wider than the modal $\pm 30\text{--}50\%$ TEA tolerance [64] as benzene's oil linkage warrants.

The BEN-Fossil scope-3 CI of 1.87 kgCO₂e/kg is the cradle-to-gate intensity of fossil benzene from Rosental et al. (2020) [138] (Table 4), for BTX from catalytic reforming of naphtha under mass allocation. Two independent sources give the same value: the PlasticsEurope EU-28 eco-profile [141] and Idemat 2026 [142], both 1.86 kgCO₂e/kg. Rosental's Table 4 also reports 5.25 kgCO₂e/kg including end of life; the difference, 3.38, is benzene's stoichiometric combustion CO₂ ($6 \times 44.01/78.11$), which falls outside the product-gate boundary of this study and is therefore excluded.

Two caveats attach to that figure. No public emission-factor data distinguishes among fossil benzene suppliers, ISCC PLUS certifying supply provenance but not CI values, so the split here is limited to fossil versus bio. And 1.87 kgCO₂e/kg rests on one defensible LCA allocation choice and could shift under alternative methods, EU supply being 94 % oil-derived and mostly from steam-cracker BTX co-product slates [143]; the fossil-versus-bio comparison should therefore be read as a directional signal that bio is lower-CI rather than as a precise quantitative gap.

Bio-attributed benzene

BEN-Bio price is varied across \$0.77 / 1.77 / 2.31/kg (low / central / high). A minimum premium of \$1.00/kg over the H1 2025 fossil benzene reference, as reported by S&P Global [136], gives the central BEN-Bio anchor of \$1.77/kg. Because that premium is a floor rather than a mid-market quote, the central is a lower bound on what bio-attributed benzene would cost, which makes the converted-route comparison of \$4.4 optimistic on their side rather than the biomass route's; this roughly 2 $\times$ premium is driven by niche-only demand at current commercial scale (BioBTX's first plant). The low \$0.77/kg anchors to BioBTX's claim of cost parity with fossil BTX at commercial scale [126], representing a future where bio-scale-up has closed the gap. The high \$2.31/kg uses the upper-end bio-naphtha premium proxy from the same S&P Global 2025 report [136] as a tight-supply bracket. No public contract pricing exists for bio-attributed benzene (commercial confidentiality); the range is triangulated from the three cited sources above rather than read directly from a single market price index.

Bio-attributed benzene is not a single process. Miller et al. [104] evaluate nine renewable-benzene routes, from catalytic pyrolysis and gasification through lignin depolymerisation to algal routes, and stress that all remain technologically immature. Only the catalytic fast-pyrolysis family is near commercial, and neither of its two developers is currently building on biomass: BioBTX reached a final investment decision in July 2026 on a plant converting about 20 kt/a of mixed plastic waste into about 10 kt/a of aromatic oil, for start-up in 2028, with biomass a stated future extension [144]; Anellotech's Bio-TCat has run more than 5,000 hours on pine wood at pilot scale and is offered for licensing, but no commercial plant has been built [145]. BEN-Bio is therefore modelled on this pyrolysis route, at TRL 6 [104], the highest among the biomass routes assessed and the maturity relevant to near-term procurement; the price above and the intensity below are that route's values, not a generic bio-benzene figure.

BEN-Bio's CI of 1.08 kgCO₂e/kg applies the 42 % reduction Zuiderveen et al. (2024) [137] report to that fossil value, as they report the BTX mixture, not benzene alone. Their 42 % is the reduction for one route on one feedstock, catalytic fast pyrolysis of crude glycerol (the BioBTX process), against a fossil-BTX reference of about 5.2 kgCO₂e/kg on a cradle-to-grave basis. Across the wider route space the figure moves substantially and is not even one-signed: Zuiderveen et al. note that starch-based aromatics have been reported above their fossil counterpart, while wood-chip routes with CO₂ capture can be net-negative, and their own biomass case falls to -0.4 to 1.6 kgCO₂e/kg under 2050 carbon-recycling scenarios. The 1.08 value is therefore included to show the potential of this route rather than as a precise or general figure.

Bio-aromatics are also severely supply-constrained. BioBTX's first commercial plant produces about 10 kt/a of aromatic oil, with no published benzene share of that slate [144], against the ~ 70 kt/a of benzene a 100 kt/a adipic acid train consumes, so BEN-Bio is a partial-substitution option today; capacity expansion over 2028–2050 could relax this.

A.5.2. Hydrogen

Hydrogen feeds the CON and ECO benzene-hydrogenation steps and the BIO polishing reactor. For ECO, cathodic H₂ co-product covers most of the hydrogenation demand, with any balance purchased (§2.2.2). It is evaluated as a three-way procurement choice across grey, blue, and green production types, pairing price and

scope-3 emission intensity per scenario. The scenario unit is the production type rather than a market-priced supplier choice, because each pathway has chemistry-determined emission intensity and technology-determined cost [146]. The three scenarios cover the most commonly available H₂ types (Table A.5): grey (steam methane reforming, no CCS), blue (SMR + CCS), and green (renewable electrolysis, here wind-powered as the cheapest electrolytic pathway in NW Europe at \$5.7–7.6/kg in 2024 vs \$8–16/kg for solar PV [128]).

Table A.5: Hydrogen procurement scenarios. The central anchors are the IEA's assessed 2024 levelised costs of production for Northwestern Europe, not transaction prices [128]: unabated gas-based reforming averaged \$3.2/kg in 2024, gas-based production with CCUS just below \$3.5/kg, and wind-powered electrolysis \$5.7–7.6/kg, of which \$6.65 is the midpoint. The 2030 STEPS values are midpoints of the IEA's published ranges (grey 2.1–3.5, blue 2.1–3.1, green 3.6–5.2 \$/kg). Scope-3 CI from IEA Global Hydrogen Review 2024 [127].

Type	2024 NW-Europe (\$/kg H ₂ , production cost)	2030 STEPS (\$/kg H ₂)	Scope-3 CI (kgCO ₂ e/kg H ₂)	Pathway
Grey	3.20	2.80	10	SMR, no CCS
Blue	3.50	2.60	3	SMR + CCS (~70 % capture)
Green	6.65	4.40	1	Renewable electrolysis (wind)

A purchased contract sits above a production cost by the producer's margin, delivery and, for renewable hydrogen, a scarcity premium. The IEA's current European market price for renewable hydrogen, 8–10 \$/kg, brackets the European Hydrogen Observatory's cost for grid-connected electrolysis rather than the directly connected pathway costed here [128, 147]. On the Observatory's 2024 country data reforming with capture runs 3.84–4.31 \$/kg across the three ARRA countries, so the blue tier sits at the low end of that band.

The IEA STEPS column is a sensitivity rather than the anchor. It is a stated-policies projection, and its costs include a CO₂ price of \$140/t [128] against the €115/t used here, so adopting it would import a second carbon price into a feedstock cost. That inclusion is also why grey exceeds blue at the midpoint, since the charge falls on the unabated pathway.

TNO 2024 [148] reports €5.98–€13.69/kg (\$6.58–\$15.06) for current Dutch electrolyser projects at 100–200 MWe, and CE Delft and TNO 2023 [149] project €8.30/kg (\$9.13) for 2030. The \$6.65 green anchor sits at the bottom of that range and the STEPS midpoint of \$4.40 below it, so both are optimistic against what a Netherlands project currently costs. The anchor is kept because it is a production cost on the same 2024 basis as the other two tiers, but any result resting on green hydrogen being cheap should be read against the higher project figures, which \$4.4 quantifies.

A.5.3. HNO₃

HNO₃ supply is the conventional route's most significant GHG lever: its scope-3 intensity swings widely with the supplier's abatement state, dominating the route footprint for a high-intensity supplier (Yeo (2025) [9] attributes about 3 of 4.79 kgCO₂e/kg AA to nitric acid; Fig. B.1, LR §2.2.1) but much lower for an abated EU supplier. The two procurement scenarios below quantify this supplier-choice exposure, bracketing an abated EU producer against a non-EU unabated one. Price and scope-3 intensity are treated as separate dimensions because CBAM extends EU ETS-equivalent carbon costs to HNO₃ imports from 1 January 2026 [80], progressively equalising the carbon-cost component across suppliers.

The HNO₃ central is the BusinessAnalytiq European index, \$0.30/kg of contained HNO₃ [150], matching the model's stoichiometric pure-acid consumption, so no grade conversion is applied (§3.5.2). The contained basis is established rather than assumed. Fuming nitric acid, at 98–100 % HNO₃, trades in Europe at \$0.24/kg [151]; an as-traded 68 % solution would have to price near \$0.17–0.18/kg, since a third of its mass is water, so an index sitting above the near-pure price cannot be quoted on a diluted one. Yeo (2025) reads the same index the same way, at face value against a pure-basis consumption [9]. Nitric acid is made from ammonia, whose cost tracks European gas [152], giving it a one-sided exposure to the upside. The 2026 European index sits at \$0.29–0.33/kg, mid-range between an oversupplied Northeast Asia (\$0.23/kg) and a tighter North America (\$0.43/kg) [153], with IMARC projecting roughly 3–4 % annual growth through 2030. The sensitivity band is therefore asymmetric, \$0.22 to \$0.50/kg (–27 % / +67 %): the low approaches the global oversupplied floor, and the high spans the 2022 gas-crisis peak (\$0.45/kg [154], Nord Stream curtailment) with headroom for the rising outlook.

A renewable (green-ammonia) nitric acid route is not evaluated as a separate procurement scenario: the dominant decarbonisation lever for conventional HNO₃ is on-site N₂O abatement, which the EI-A scenario already captures, while emerging green-ammonia capacity is currently routed to fertiliser production rather than sold as merchant green nitric acid, leaving no green-HNO₃ price or settled CI to anchor a scenario.

HNO₃ cradle-to-gate scope-3 CI varies along two axes characterised in the LCA literature: N₂O process emissions at the HNO₃ plant (depending on abatement state) and upstream ammonia CI (depending on feedstock and plant location). Two primary-source datasets feed this thesis: Fertilizers Europe [123] (updated for the EU by Oeko-Institut/NACAG [33]) for per-region N₂O process emission factors, and Liu et al. [155] for ammonia cradle-to-plant-gate emission intensity. Both are consistent with the IPCC 2019 Refinement [156] and the EU ETS Phase 4 product benchmark of 0.230 tCO₂e/t HNO₃ [78]. Yeo [9] anchored the conventional route's HNO₃ scope-3 to a single industrial-average value [157]; this thesis instead constructs two regional cradle-to-gate scenarios for an ARRRA-located producer, using the primary-source component values listed in Table A.6: (i) *Modern European abated supplier* (EI-A), a typical EU plant under EU ETS Phase 4 with N₂O catalytic abatement and natural-gas-derived ammonia; and (ii) *Non-EU coal-based unabated supplier* (EI-B), representative of contemporary Chinese supply. The combined total cradle-to-gate scope-3 CI per scenario, derived from these components, is given in the methodology (§3.5.2, Table 3.6).

Table A.6: Primary-source **component** values (N₂O process emission factor and upstream NH₃ carbon intensity) feeding the EI-A and EI-B HNO₃ procurement scenarios, in native reporting units. N₂O values from Fertilizers Europe [123]; NH₃ cradle-to-plant-gate values from Liu et al. [155]. The combined **total cradle-to-gate** scope-3 CI per scenario, derived from these components, is given in methodology §3.5.2, Table 3.6.

Scenario	N ₂ O process EF (kg N ₂ O/t HNO ₃)	Upstream NH ₃ CI (kgCO ₂ e/kg NH ₃)
EI-A	0.7 (EU 2014)	2.6 (SMR; range 2.1–3.6)
EI-B	7.4 (China, CDM-era)	6.95 (coal; range 6.1–7.8)

A.5.4. MA broth

MA broth has no public spot market and is held at \$1.86/kg, the value Yeo (2025) [9] carries for the same stream and derives from Mokwatlo (2024) [51] as that study's minimum selling price for muconic acid less the cost of crystallising it, which leaves the clarified broth that enters the flowsheet as feed. Against Mokwatlo's base-case selling price of \$2.70/kg the implied crystallisation wedge is \$0.84/kg. Price is treated as a $\pm 30\%$ sensitivity range around the central (modal chemical-TEA practice [64]), giving \$1.30–\$2.42/kg. Holding that wedge fixed, Mokwatlo's realistic operating envelope (yield 0.4–0.6 mol/mol, productivity 0.4–1.0 g L⁻¹h⁻¹, across its three pretreatments) translates to roughly \$1.50–\$2.10/kg of broth, so the tier overshoots it on both sides. That is the intended behaviour for a twenty-year horizon, since fermentation performance over a plant life can plausibly move beyond currently demonstrated parameters, and it adds headroom for commodity-market uncertainty the study does not carry.

The broth scope-3 emission intensity is carried at a central 0.38 kgCO₂e/kg MA-content, with a sensitivity of 0.25–0.55 kgCO₂e/kg, from the reference model's feedstock table [9], which derives it from Mokwatlo (2024) on the same basis as the price [51]. The value is *net* of the biogenic carbon embodied in the product: adding back the 1.86 kgCO₂e/kg carried by the six carbons of muconic acid recovers 2.24 kgCO₂e/kg, inside the range Mokwatlo report for the hydrogenated product. It is therefore a net cradle-to-gate figure, and no further embodied-carbon credit is applied to the biomass routes (methodology §3.4). This range is adopted as a continuous sensitivity coupled to the price tier without additional headroom, because within a fixed glucose pathway the dominant CI drivers are the ones Mokwatlo varies. The coupling pairs low price with low CI (high-titre, low-impurity operating point) and high price with high CI (low-titre end of the Mokwatlo operating-parameter range). No discrete procurement-scenario tier is built: MA broth has no commercial market and feedstock-pathway alternatives are out of scope (Section 2.2.3).

A.5.5. Process chemicals

Process chemicals (KOH electrolyte make-up in ECO; H₂SO₄ crystallisation and Ca(OH)₂ wastewater neutralisation in BIO) are small by OPEX share (< 7 %) and small in route GHG except for BIO, where H₂SO₄ and Ca(OH)₂ jointly contribute approximately 22 % of t3HDA-route CO₂e per Yeo's stream tables [9]. The conventional route buys none of the three. Price is propagated as a one-at-a-time sensitivity across the bands of Table A.7; the scope-3 intensities are held at their central values and no sensitivity is carried on them in this analysis (Chapter 6). Band widths differ by chemical: KOH spans the chlor-alkali allocation choice (economic versus mass) and EU-grid vintage; H₂SO₄ spans the production-route mix from smelter by-product at the low end to virgin-sulphur burning toward the upper end; Ca(OH)₂ is bounded below by the stoichiometric calcination CO₂ floor (0.78 kgCO₂/kg CaO) and above by the European sector-average cradle-to-gate.

Table A.7: Process-chemicals price and scope-3 CI anchors. Prices are swept across the band; the intensities are held at their central value and the low and high figures show the spread in the literature rather than a sensitivity carried here. CI sources: KOH [158, 159, 160]; H₂SO₄ [161]; Ca(OH)₂ [162, 163]. Price sources: [164, 165, 166].

Chemical	Price low / central / high (USD/kg)	Scope-3 CI (kgCO ₂ e/kg)	low / central / high
KOH ECO electrolyte make-up	0.62 / 0.69 / 0.83	1.12 / 1.48 / 2.0	
H ₂ SO ₄ BIO crystallisation	0.13 / 0.19 / 0.35	0.025 / 0.14 / 0.16	
Ca(OH) ₂ BIO wastewater	0.12 / 0.14 / 0.20	0.85 / 0.94 / 1.05	

The Ca(OH)₂ band is the one that has dated: 2026 European indices quote 0.23 and 0.51 \$/kg for grades they do not specify, both above the high tier used here. Carrying the higher of them would add about 0.07 \$/kg to the biomass routes and nothing to the other two, so it changes no ordering.

A.6. Utility emission factors

The greenhouse-gas emission factors for the process utilities are collected here. The grid electricity factor is country-dependent and is given with the location analysis (§4.3); the feedstock cradle-to-gate scope-3 factors are set out in §A.5.5 (Table A.7), and the nitrous-oxide GWP₁₀₀ of 273 is given with the environmental method (§3.4). The on-site steam factor enters the results as scope-2, following Yeo's utility convention, and steam-generation credits are applied at the same factor with the opposite sign.

Table A.8: Utility greenhouse-gas emission factors, per unit of duty (kgCO₂e/kWh).

Utility	Factor	Basis and source
On-site steam (LP and MP)	0.202	Natural-gas combustion on a net calorific value basis [167], held flat off the grid-intensity trajectory
Cooling water	0	Pumping power already counted in the electrical duty

The steam factor is the default combustion factor for natural gas, 56 100 kgCO₂/TJ on a net calorific value basis, which is 0.202 kgCO₂/kWh [167]; the same figure, with the methane and nitrous oxide terms added, stands at 0.20199 kgCO₂e/kWh in the 2026 United Kingdom conversion factors [168]. It is applied here per kilowatt-hour of steam duty rather than per kilowatt-hour of gas fired, which assumes that the whole net calorific value of the fuel reaches the process as steam enthalpy. The steam term is therefore a lower bound. A gas-fired steam boiler running at 85 to 92 % on the same basis would carry 0.22 to 0.24 kgCO₂e/kWh, which would add about 0.15 kgCO₂e/kg to the electrochemical route and less than 0.03 to each of the other three, without changing the order of the four.

A.7. Electrochemical cell performance: the evidence base

The durability data for the two cells, the analogues used to bound their degradation, and the measurement behind each performance tier of Table 2.4 are set out here. The tiers themselves, and what they are used for, are in §2.4.5.

What the durability data show

Published durability data for the specific cell types in this thesis remain scarce. Liu et al. [39] operated an alkaline AEM-based membrane electrode assembly for KA oil oxidation on a NiV-LDH catalyst, the same cell architecture as the ECO route. At 200 mA cm⁻², Faradaic efficiency fell from 93 % to 81 % over 60 h while cell voltage rose by 7 %; at 300 mA cm⁻² the test was terminated after 46 h with a 22 % voltage rise, primarily due to AEM corrosion by cyclohexanone. For the BIO route ECH cell, Dell'Anna et al. [57] demonstrated approximately 70 h of continuous flow-reactor operation on a bismuth cathode and, in companion lead-cathode experiments, identified cathodic corrosion and lead leaching (hundreds of ppb in reacted broth by ICP-OES) as unavoidable durability limiters. No study on either cell type reports data beyond ~100 h, far shorter than the thousands to tens of thousands of hours quoted for mature electrochemical stacks, which range from manufacturer-stated stack lifetimes to plant service records [45, 109, 124].

Analogues for the two cells

Because directly relevant long-term data do not exist, degradation must be bounded using analogues from more mature electrochemical technologies. The two cell types require different analogues because they differ fundamentally in architecture and dominant degradation mechanism.

For the ECO cell, the dominant degradation pathways are membrane deterioration and catalyst deactivation; alkaline and AEM water electrolyzers are the closest analogue as they share nickel-class catalysts and KOH-compatible chemistry, though they use different separators (ZrO₂-based diaphragm in alkaline, anion-exchange membrane in AEM). Published stack lifetimes range from 50 000–80 000 h for PEM and 60 000–90 000 h for alkaline systems [45, 109, 124]. AEM systems, whose architecture the ECO cell shares, are far less mature: IRENA report a 2020 stack lifetime above 5 000 h against a 100 000 h target [109]. For component replacement scheduling, the chlor-alkali industry provides an industrial reference: membrane replacement every four to five years and anode recoating after eight years [169]. These values are an optimistic lower bound for the ECO cell, which additionally exposes its AEM to cyclohexanone, a membrane-corrosive agent absent in water electrolysis or chlor-alkali service [39].

For the BIO route ECH cell, the lead cathode is the primary degrading component and its corrosion rate determines the electrode replacement interval. The relevant analogue is industrial electroorganic synthesis: lead cathodes are used at scale in processes such as adiponitrile production ($>300\,000\text{ t yr}^{-1}$), where bulk lead cathodes corrode at $0.35\text{--}0.9\text{ mm yr}^{-1}$ depending on the ammonium additive regime [110]. Electrodeposited lead–copper alloy cathodes reach 0.05 mm yr^{-1} in the same work, but that is a different electrode from the bulk lead or bismuth cathode used here, so the bulk range is the relevant one. These rates reflect controlled synthetic electrolytes with quaternary ammonium additives that protect the cathode surface. The MA-to-*t*3HDA ECH cell operates under harsher conditions: the fermentation broth contains amino acids, proteins, and inorganic salts, and the muconate dianion can chelate dissolved Pb²⁺, driving further electrode dissolution [57].

In both cases the analogue values represent an optimistic lower bound, because each cell faces route-specific degradation mechanisms absent in its respective analogue. The severity of these mechanisms at industrial operating hours is unknown, motivating parametric treatment of both initial cell performance and degradation.

The evidence behind each initial-performance tier of Table 2.4 follows, first for the ECO cell and then for the ECH cell.

The ECO tiers

The three ECO tiers are the same nickel hydroxide catalyst at three levels of modification. The central tier (FE 0.80) is the vanadium-modified form of Liu et al. [39], which sustains 82 % FE at 300 mA cm^{-2} in a membrane-electrode assembly and is the only ECO catalyst demonstrated at industrial current density. The optimistic tier (FE 0.90) is the SDS-modified form of Li et al. [40], which reaches 93 % but only at low current, so it is treated as a target. The conservative tier (FE 0.45) is the unmodified catalyst, whose efficiency falls to 42 % once the current density is raised to 114 mA cm^{-2} [39]; no figure is reported for it at 300 mA cm^{-2} , so 0.45 is a rounded floor set just above that measured value for tier symmetry, rather than a value carried along the falling trend, which would place it lower still.

The ECH tiers

For the ECH cell the central tier is fixed by the operating point itself. Dell’Anna et al. [57] take 200 mA cm^{-2} at 5.7 V as their cost optimum, and processing real fermentation broth at that current density they measure a Faradaic efficiency of 40.9 %. The efficiency is strongly current-dependent, falling from 88.5 % at 50 mA cm^{-2} to 24 % at 400, while the *t*3HDA yield stays near 90 % throughout: what is lost at higher current goes to hydrogen evolution, not to the product. Pinning *j* therefore fixes the central tier at 0.40, and Yeo [9] adopts the same point, so the central case also reproduces the reference design. Because the charge lost to hydrogen evolution leaves as H₂, the cathodic hydrogen co-product and the Faradaic efficiency are inversely coupled: a higher-FE tier draws less current for the same *t*3HDA output, and a smaller share of that current is lost to hydrogen, so the evolved hydrogen falls as $(1 - \text{FE})/\text{FE}$. The two cannot be varied independently, so the hydrogen credit falls as efficiency rises, and at the optimistic tier the biomass route to adipic acid can turn from hydrogen-surplus to hydrogen-short.

The tiers either side of it describe the same cell under different conditions rather than a different cell. The optimistic 0.60 is Matthiesen et al.’s [56] lead cathode, which reached 60 % at 90 % selectivity but in a dilute model solution rather than in broth, so it stands for a cleaner feed rather than a better cell; the near-quantitative efficiencies reported under argon purging [50] do not transfer to a continuously fed reactor and are not carried as a tier. The conservative 0.30 is the efficiency the same cell delivers at 300 mA cm^{-2} , applied here at the design current to represent fouling and impurity accumulation in real broth. Faradaic efficiency is the only channel available for that downside, since the ECH degradation model treats electrode corrosion as periodic replacement and carries no specific-energy penalty. Degradation rate and replacement interval are anchored to the mature-process analogues

identified above; the values used, and the cost charged at each replacement, are set out in Section 3.2.6. Unlike the performance tiers in Table 2.4, which are efficiencies measured in the cited work, those are brackets transferred from the analogues rather than reported values, so they belong with the modelling choices rather than with the evidence reviewed here.

B

Headline results build-up

This chapter gives the component build-up behind the headline levelised cost of adipic acid and GHG intensity of Table 4.1, for each route on the central basis: the heat-integrated flowsheets of §4.2.2, cathodic hydrogen burned for steam (§4.2.3), cell degradation at the central tier, and the EU ETS charge at full free allocation, $\alpha = 1$, and 115 €/tCO₂. The two biomass rows are stated per kilogram of different products and are not a like-for-like comparison (§4.2.7).

Table B.1: Levelised cost of adipic acid and GHG intensity by component, per route (central case). The price block is the cash build-up; the emissions block is the scope split. Totals are those of Table 4.1.

Component	CON	ECO	BIO _{t3HDA}	BIO _{AA}
<i>Levelised cost of adipic acid (\$/kg)</i>				
Capital	0.273	0.478	0.691	0.777
Fixed operating cost	0.437	0.570	0.817	0.895
Feedstock	1.091	0.654	2.206	2.317
Utilities	0.048	0.565	0.419	0.462
Waste disposal	0.084	0.088	0.067	0.064
EU ETS	−0.064	−0.166	−0.173	−0.173
LCOA	1.87	2.19	4.03	4.34
<i>GHG intensity (kgCO₂e/kg)</i>				
Scope 1	0.892	0.090	0.034	0.034
Scope 2 ¹	0.181	2.324	1.601	1.786
Scope 3	2.318	1.377	0.707	0.743
GHG	3.39	3.79	2.34	2.56

¹ On the biomass routes scope-2 is net of the cathodic-hydrogen steam credit, which falls from −0.127 to −0.045 kgCO₂e/kg once the polishing reactor takes most of the hydrogen.

Components are rounded independently, so a column may differ from its total in the last place.

B.1. Cell-degradation saw-tooth

The electrochemical route's specific energy is not constant over the plant life: the cell voltage rises within each membrane campaign and resets when the membrane is replaced. The end-of-life criterion is the same 10 % voltage rise on every stack-life tier, so how far the specific energy climbs within a campaign, and its levelised midpoint, are identical by construction across the three panels of Figure B.1; only the interval between replacements is swept. What that is worth is in §4.5.3.

B.2. Cell performance against price and grid intensity

The nine combinations behind the ranges quoted in §4.5.3: each electrified route over its three cell-performance tiers crossed with the three electricity prices on cost and the three grid intensities on carbon, with the conventional route given as a reference line on each panel.




Figure B.1: Electrochemical-route specific energy consumption over the twenty producing years, one panel per stack-life tier on identical axes. Within each membrane campaign the cell voltage rises linearly to its end-of-life criterion, a 10 % rise on the cell and so 5.1 % on the route's electricity, and resets when the membrane is replaced: 5.05 kWh/kg at beginning of life to 5.30 kWh/kg at end of life. The dashed line is the campaign midpoint, the levelised specific energy consumption of 5.17 kWh/kg. The climb and its midpoint are identical in all three panels by construction; only the replacement frequency differs.

Table B.2: The two electrified routes over the cell-performance tiers and the two inputs the cell multiplies: electricity price on cost, grid carbon intensity on emissions. All other inputs are central, and the centre cell of each panel is the headline value. The conventional route is given as the last row of each panel, on the same tiers. Grid intensities are the 2024 ARRRR values of Table 4.8, in gCO₂e/kWh.

(a) ECO LCOA (\$/kg)				(b) ECO GHG (kgCO ₂ e/kg)			
Cell	EP-Low	EP-Cen.	EP-High	Cell	BE 150	NL 261	DE 341
Optimistic	2.11	2.19	2.37	Optimistic	3.13	3.73	4.16
Central	2.11	2.19	2.38	Central	3.16	3.79	4.24
Conservative	2.18	2.29	2.56	Conservative	3.25	4.12	4.74
CON	1.86	1.87	1.89	CON	3.32	3.39	3.44

(c) BIO _{AA} LCOA (\$/kg)				(d) BIO _{AA} GHG (kgCO ₂ e/kg)			
Cell	EP-Low	EP-Cen.	EP-High	Cell	BE 150	NL 261	DE 341
Optimistic	4.16	4.22	4.37	Optimistic	1.62	2.11	2.46
Central	4.25	4.34	4.56	Central	1.86	2.56	3.07
Conservative	4.35	4.47	4.76	Conservative	2.08	3.00	3.66
CON	1.86	1.87	1.89	CON	3.32	3.39	3.44

B.3. Cathodic hydrogen valuation

The steam credit of \$4.2.3 is deliberately conservative in three ways. The hydrogen is valued only as fuel, at roughly \$1.23 per kilogram of hydrogen; the boiler efficiency is taken at the middle of its cited range rather than the top; and the packaged boiler is over-costed twice over, its duty falling below the correlation's 5 000 kg/h floor so that the smallest unit is priced, and on a 15–40 bar basis that suits a much higher pressure than the 2.2 bar boiler the duty actually needs (\$3.3.2).

Merchant sale is not valued at all. There is no published netback price for by-product hydrogen at this scale, no assured offtaker, and the flowsheet carries neither the drying nor the offtake connection a sale would need. Euro Chlor report that 10 to 15 % of European chlor-alkali by-product hydrogen goes unused for want of a market or the means to move it [170], so an offtaker cannot be assumed at this scale either.

Table B.3: The two treatments of the surplus cathodic hydrogen, per biomass route. *Yeo* is the \$0.70/kg avoided-disposal credit at a zero emission factor, applied to the same surplus; the reference model does not carry the route to adipic acid, so that row is that treatment carried onto it. The two are alternatives, not additions: hydrogen that is burned is not disposed of, so the credit is withdrawn where the steam credit is taken.

Route	Yeo (2025)		This work	
	Δ MSP (\$/kg)	Δ GHG (kgCO ₂ e/kg)	Δ LCOA (\$/kg)	Δ GHG (kgCO ₂ e/kg)
BIO _{t3HDA}	−0.016	0	−0.006	−0.127
BIO _{AA}	−0.005	0	−0.000	−0.045

Burning the 273.6 kg/h of the route to *t3HDA* raises 7.8 MW of low-pressure steam, 76 % of that route's demand; the route to adipic acid has only 91.7 kg/h left once the polishing reactor has taken its feed, and raises 2.6 MW, or

25 %.

Merchant sale is not quantified, for want of a netback price, an assured offtaker and the drying and offtake connection the flowsheet does not carry. Its scale is nonetheless fixed by arithmetic: the surplus is 0.022 kg of hydrogen per kilogram of *t*3HDA, so every \$1/kg of merchant value above that \$1.23 fuel value would lower the levelised cost of adipic acid by a further \$0.022/kg. Selling rather than burning would forfeit the 0.127 kgCO₂e/kg the steam credit delivers, so the two cannot both be claimed.

The withdrawn credit is not lost value. It reappears as the steam the hydrogen displaces, which is why the waste line rises (§3.3.2) while the levelised cost barely moves.

The electrochemical route's cathodic hydrogen is credited on a different basis again, as an avoided purchase against its own benzene-hydrogenation demand, which leaves 32.3 kg/h to buy in.

B.4. Capital accuracy band by equipment group

How the band is weighted

The Class-4 accuracy band reported in §4.5.4 is a weighted mean of three group bands, and the weights are each route's own capital composition. The groups are those of §3.3.3: items sized from the converged flowsheet and costed through the Towler correlation carry $-15/ + 30$ %; items carried at the purchased price recorded in Yeo's workbooks carry $-20/ + 40$ %; and the electrochemical cells, together with the electrodialysis unit, carry $-30/ + 50$ %.

Composition and resulting bands

The composition differs enough between routes to move the band by nine percentage points on the low side. Table B.4 gives it and the band that falls out of it.

Table B.4: Installed capital by AACE accuracy group, per route (\$M, with each group's share of inside-battery-limits capital in brackets), and the weighted band. Group bands are $-15/ + 30$, $-20/ + 40$ and $-30/ + 50$ % (§3.3.3). Groups are rounded independently, so a row may differ from its total in the last place.

Route	Correlation	Direct	Cells and ED	ISBL	Band (%)
CON	71.6 (72 %)	28.3 (28 %)	—	99.8	$-16.4/ + 32.8$
ECO	73.7 (36 %)	28.9 (14 %)	99.9 (49 %)	202.5	$-23.1/ + 41.3$
BIO _{t3HDA}	14.1 (5 %)	109.2 (39 %)	160.1 (56 %)	283.4	$-25.4/ + 45.2$
BIO _{AA}	17.1 (6 %)	109.2 (36 %)	177.3 (58 %)	303.6	$-25.6/ + 45.3$

What each band is worth on the levelised cost of adipic acid is given in Table B.5. The second column moves capital alone; the third adds the part of the fixed operating cost charged on the capital base, maintenance at 5 % of inside-battery-limits cost and taxes and insurance at a further 1.5 %, which moves with it.


Table B.5: What the capital accuracy band of Table B.4 is worth on the levelised cost of adipic acid, every other input central.

Route	Band (%)	Capital alone (\$/kg)	With capital-linked fixed OPEX (\$/kg)
CON	$-16/ + 33$	$-0.04/ + 0.08$	$-0.07/ + 0.13$
ECO	$-23/ + 41$	$-0.11/ + 0.19$	$-0.18/ + 0.32$
BIO _{t3HDA}	$-25/ + 45$	$-0.17/ + 0.31$	$-0.28/ + 0.51$
BIO _{AA}	$-26/ + 45$	$-0.20/ + 0.35$	$-0.32/ + 0.57$

The weighting runs on inside-battery-limits capital. Off-site, engineering and contingency are uniform multipliers on that base, so they cancel from a share-weighted mean and cannot move the band; the N₂O abatement facility, which enters fixed capital outside the factorial build-up (§3.3.3), is held fixed while the band is applied. Carried instead at the directly-costed band, it would widen the conventional route to $-16.8/ + 33.7$ %.

Electrodialysis is weighted with the cells rather than with the correlation-costed items, because it has no Towler correlation either (§3.3.3). At roughly a fifth of the electrochemical route's inside-battery-limits capital the choice is visible rather than cosmetic: grouped with the directly-costed items that route would carry $-21.0/ + 39.2$ % and with the correlation-costed items $-20.0/ + 37.1$ %, against the $-23.1/ + 41.3$ % used here.

One qualification on the widest column. On the biomass route to adipic acid, \$17.2 M of what is grouped with the cells is the catalytic hydrogenation reactor and its recovery train rather than cell capital, so that route's share of the $-30/ + 50$ % group is wider than its electrochemical content alone would place it.



Comparison to Yeo (2025)

C.1. Input values

This section lists every input to the headline results, this work against Yeo (2025): feedstock and utility prices, emission factors, equipment cost correlations and installation factors, and financial parameters. Values that differ and drive the headline differences (the compressor and gas-turbine correlations, the grid carbon intensity, and the N₂O abatement capital) are noted as such.

The values listed under *this work* are those of the central case; those under *Yeo (2025)* are the inputs used to reproduce his results. An asterisk marks an input that differs between the two and moves the headline result.

Input	This work	Yeo (2025)	Notes
<i>Feedstock prices (\$/kg)</i>			
Benzene	0.88	0.89	
Hydrogen (blue)	3.50	3.49	
Nitric acid	0.30	0.29	
Muconic-acid broth	1.861	1.861	
KOH	0.69	0.60	
Sulphuric acid	0.19	0.13	
Lime	0.14	0.14	
<i>Utility prices</i>			
Electricity (\$/kWh)	0.0682*	0.16	the relieved EHV tariff of €62/MWh at the 1.10 USD/EUR below, against a generic industrial price
Cooling water (\$/kWh)	0.003	0.003	
Chilled water	COP 5.0*	n/a	charged as refrigeration electricity, not as a \$/kWh utility; Yeo's colder cooling water serves every stream, so his model carries no refrigeration
LP steam (\$/kWh)	0.041	0.041	
MP steam (\$/kWh)	0.048	0.048	
Wastewater disposal (\$/t)	2.2	2.2	
<i>Emission factors</i>			
Grid intensity (kgCO ₂ e/kWh)	0.261*	0.05	NL 2024 lifecycle grid, not a near-renewable supply
On-site steam (kgCO ₂ e/kWh)	0.202	0.202	
Scope-3 benzene (kgCO ₂ e/kg)	1.87	1.2	
Scope-3 hydrogen	3.0	5.2	
Scope-3 nitric acid	0.93*	3.9	EU-abated supply, not unabated
Scope-3 muconic broth	0.38	0.385	
Scope-3 KOH	1.48	0.77	
Scope-3 H ₂ SO ₄	0.14	0.14	
Scope-3 lime	0.94	0.79	
N ₂ O generation (kg/kg AA)	0.30*	0.216	IPCC stoichiometric, not the R4 flowsheet output
N ₂ O abatement	99 %*	100 %	realistic new-build vs idealised

Input	This work	Yeo (2025)	Notes
<i>Capital-cost method</i>			
Compressor correlation (a, b)	580,000, 20,000*	58,000, 2,000	Towler & Sinnott Table 7.2 values here; the erroneous constants are retained on the Yeo side for reproduction, and the gas turbine at that value on both sides
CEPCI base (2010 → 2024)	532.9 → 800	532.9 → 801	0.13 % apart; moves no reported figure
OSBL factor	0.40	n/a*	Yeo's single Lang factor carries no separate off-site line
Design & engineering factor	0.25	n/a*	likewise
Contingency	10 %	10 %	of ISBL+OSBL here, of installed cost for Yeo
Working capital	15 %	15 %	bases differ; see note below
<i>Financial parameters</i>			
Real discount rate	8 %*	15 %	real rate here; the 15 % is kept as a sensitivity
Plant life (production years)	20	20	Yeo annualised, no ramp
Operating hours (h/yr)	8000	8000	
EUR/USD	1.10	1.10	
Carbon price (€/tCO ₂)	115*	0	EU ETS applied; Yeo none
Free allocation α	1	—	bound at $\alpha = 0$
ETS benchmark (tCO ₂ e/t AA)	1.40	—	Yeo carries no ETS position
Production rate (kt/a)	100	100	

Table C.1: Every input to the headline results, this work against Yeo (2025).

Both columns are verified against the primary sources cited for each input. Yeo's capital build-up applies a single Lang factor and therefore carries no separate off-site or design-and-engineering line; his contingency and working-capital factors are as listed, and the resulting per-route amounts are given in Table C.2.

C.2. Capital-cost decomposition

This section gives the numerical backing for the capital-cost decomposition of Figure C.1: the fixed capital by equipment category for each route, this work against Yeo (2025).

The installed-equipment lines are direct costs from the same plant objects that generate the headline LCOA, re-costed here and taken from Yeo's own workbooks for his. The electro-oxidation cell is not a separate line: on both sides it sits inside the reactor block (Yeo labels it "reactors, R1–R4 incl. ECO cell"), so the reactor row carries it. Below the installed subtotal the build-up to fixed capital is shown explicitly: Yeo's single Lang contingency against this work's off-site, engineering and contingency factors, plus the N₂O abatement facility costed on the conventional route only. Working capital is recoverable and held separate from fixed capital on both sides.

Two rows are not like-for-like and should be read together rather than divided against each other. *Columns / vessels* bundles, on the Yeo side, the condenser and reboiler this work carries in the heat-exchange network, and *Heat exchangers* is that network redesigned rather than re-costed.

C.3. Sources of the cost difference

The bridge of Table 4.2 attributes the difference between Yeo's levelised cost and this work's to six components. The two this section decomposes, capital and fixed operating cost, are taken by the source of the increase rather than by the line they land on. The equipment-category view of the same capital is in §C.2; the input-by-input comparison in §C.1.

C.3.1. Cash-flow timing

One part of the difference is not cost at all but when the cash flows. Yeo's annualised-capital method assumes full output from the first year with no construction period, whereas this work discounts a two-year construction and a two-year production ramp. Removing the ramp, so that all twenty producing years run at full capacity as Yeo assumes and holding everything else at this work's basis, lowers the levelised cost to \$1.80, \$2.10, \$3.90 and \$4.20/kg for CON, ECO, BIO_{t3HDA} and BIO_{AA}: a timing effect of \$0.06–0.14/kg.

Working capital is a second timing difference. Both studies set it at 15 %, but not of the same base: here it is 15 % of the fixed capital of the plant itself, which excludes the N₂O abatement facility added outside the factorial

Table C.2: Capital cost by equipment category, this work against Yeo (2025), per route (\$M). Installed equipment above the rule; the fixed-capital build-up below it. *Columns / vessels* and *Heat exchangers* are not like-for-like across the two columns, for the reasons given above. The biomass routes carry no column or adsorber in either model, so those dashes are not missing values. Rows are rounded independently, so a column may differ from its subtotal in the last place.

Component	CON		ECO		BIO _{t3HDA}		BIO _{AA}
	Yeo	Ours	Yeo	Ours	Yeo	Ours	Ours
Heat exchangers	9.8	27.4	40.5	39.8	20.3	12.9	16.2
Reactors (incl. electrochemical cell)	17.8	13.0	75.7	73.1	155.1	160.1	163.0
Columns / vessels	11.4	6.6	26.8	12.3	—	—	—
Compressors	6.1	37.5	3.0	21.7	2.0	—	0.5
Gas / steam turbine	1.4	1.0	—	—	—	—	—
Crystallisers	17.0	12.2	15.9	11.4	72.2	51.6	64.1
Centrifuges	3.0	2.1	2.3	1.6	12.6	9.0	10.3
Electrodialysis	—	—	41.3	42.6	—	—	—
Refrigeration / cooling	—	—	—	—	47.1	48.6	48.6
Other (H ₂ boiler)	—	—	—	—	—	1.2	0.8
Installed equipment (ISBL)	66.5	99.8	205.6	202.5	309.3	283.4	303.6
Off-site, engineering, contingency	6.6	88.8	20.6	180.2	30.9	252.2	270.1
N ₂ O abatement facility	—	25.0	—	—	—	—	—
Fixed capital	73.1	213.7	226.1	382.8	340.2	535.6	573.7
Working capital (recoverable)	11.7	28.3	36.3	57.4	54.6	80.3	86.1
Total CAPEX	84.8	242.0	262.4	440.2	394.8	616.0	659.8

build-up, while Yeo takes 15 % of installed cost plus working capital, equivalent to 17.65 % of installed cost on all three of his routes. Yeo also folds it into the capital he annualises, whereas here it is drawn at start-up and released at end of life, so only its financing cost falls on the project (§3.3.3). Unlike the changes below, it lowers the burden rather than raising it.

C.3.2. Fixed capital

Fixed capital is well above Yeo on every route, and Table C.3 separates the sources. Most of the gap is not more expensive equipment but a fuller capital build-up. Yeo applied a single Lang factor to installed equipment, whereas this work uses the Towler & Sinnott factorial build-up, in which off-site, design-and-engineering and contingency costs appear as explicit lines rather than being absorbed into one multiplier (§3.3.1). Those lines are the largest entry on every route, and on the biomass route effectively the whole increase.

On most equipment categories the difference is installation alone. Reactors, crystallisers, centrifuges, electrodialysis and refrigeration carry the same purchased cost on both sides, so only the factor of Table 3.4 changes. Its direction is not uniform: fluids fall from Yeo's 4.40 to 3.20 and solids from 3.50 to 2.50, while the packaged units, electrodialysis and refrigeration, rise slightly from the 3.10 he installed them at to 3.20 here. Together these categories are what pulls the other-equipment row down \$16–22 M on every route.

Three rows are not installation comparisons and should not be read as one. The compressor row is the constant correction below, not a change of factor. The columns row is narrower in scope on this side than on Yeo's: his column cost bundles the condenser and the reboiler with the shell and the trays, whereas here those two duties sit with the heat-exchange network, which is redesigned rather than inherited. They are therefore still in the plant but cannot be followed from one row to the other, so the two columns rows are not comparable item by item. The heat-exchanger row is the same case seen from the other side, two different networks rather than the same equipment re-costed; it falls on the electrochemical and biomass routes. The conventional route is the exception, its heat-exchanger row rising \$17.7 M and nearly cancelling the reduction on its other equipment, so that route's installed cost ends within \$1 M of Yeo's. That rise is not a costing difference but a real addition of equipment: Yeo's conventional flowsheet recovers only about 5 % of its available heat, so designing its network adds exchanger area the other two routes, already well integrated in his design, do not need. Networks are designed here for all three routes (§4.2.2) rather than inherited from Yeo.

The corrections to Yeo's cost data fall unevenly (§3.2.7). Restoring the compressor constants adds about \$31 M of installed equipment to the conventional route, through its recycle and refrigeration compressors, and about \$19 M to the electrochemical route through its two large air compressors. The conventional figure is net of its gas turbine, whose cost is carried at Yeo's purchased value throughout and whose installed cost therefore falls by \$0.4 M on the installation factor alone, as the other directly-costed categories do; the biomass route is left essentially untouched, its surplus hydrogen being burned on site rather than compressed for export. The N₂O abatement facility, modelled

Table C.3: Fixed capital, Yeo (2025) to this work, by source of the difference (\$M, central case). The working-capital row is a change of treatment, not of cost: it leaves the fixed-capital base here rather than being annualised into it. Yeo levies the 15 % on installed cost plus working capital, before a separate 10 % contingency, so that row is 13.8 % of the first-row totals.

Source	CON	ECO	BIO _{t3HDA}
Yeo (2025) total CAPEX	84.8	262.4	394.8
Working capital, now recoverable	-11.7	-36.3	-54.6
Heat-exchanger network re-costed	+17.7	-0.7	-7.3
Other installed equipment (Towler vs Lang factors)	-15.3	-21.1	-16.6
Compressor correlation restored	+31.0	+18.7	-2.0
Offsite, engineering & contingency (implicit in Lang)	+82.2	+159.8	+221.3
N ₂ O abatement facility (CON only)	+25.0	—	—
This work fixed capital	213.7	382.8	535.6

by Yeo but left uncoded, is unique to the conventional route and adds \$25 M. Stacked on the fuller build-up, those two make the conventional route the one that rises most in relative terms: its fixed capital ends 152 % above Yeo's total capital, while the electrochemical route rises 46 % and the biomass route 36 %.

These figures reach the price only through an annualised charge, which compresses them considerably. Capital adds \$0.13/kg on the conventional route and \$0.07 and \$0.05/kg on the electrochemical and biomass routes (Table 4.2), against increases of 152, 46 and 36 % in the capital itself. Part of that compression is the discount rate: this work levels at 8 % real where Yeo applies a 15 % capital recovery factor, so each dollar of capital carries a smaller annual charge, and the reduction is worth most on the routes with the largest capital base. Figure C.1 shows the same comparison resolved by build-up component rather than by source of the difference.




Figure C.1: Fixed capital composition, this work against Yeo (2025), per route, by build-up line. Working capital is the top segment, shown separately because it is recoverable here rather than sunk. The electrochemical route's installed equipment already carries its electro-oxidation cell and electrodialysis stacks. Per-component values in §C.2; the full input-by-input comparison to Yeo is in §C.1.

C.3.3. Fixed operating cost

Fixed operating cost is what the plant costs to keep staffed and maintained each year, whatever it produces: labour, maintenance, insurance and overheads, with feedstock and electricity counted separately because they rise and fall with output. It is the largest positive term in the bridge, \$0.33, \$0.39 and \$0.56/kg, yet it is not a sign that this plant is more expensive to run than Yeo's. It is mostly that the method used here counts more categories of fixed cost than Yeo's did. Table C.4 separates the increase into four parts, taken in turn below, and a fifth that only the conventional route carries: maintenance and insurance on the N₂O abatement train Yeo modelled but left

uncosted (§4.2.1). The fourth part is not a difference of cost scope at all but of levelisation: fixed cost is carried in full through the commissioning years, when the plant produces a third and then four-fifths of its nameplate output.

Table C.4: Fixed operating cost, Yeo (2025) to this work, by source of the increase (\$/kg, central case). Rows sum to the fixed-cost line of Table 4.2, to within 0.01 from rounding.


Source	CON	ECO	BIO _{t3HDA}
Capital-linked lines, re-scaled (maintenance, insurance)	+0.02	0.00	0.00
Cost lines absent from Yeo's estimate	+0.19	+0.25	+0.44
Common lines on a wider basis (labour, overheads)	+0.07	+0.08	+0.06
Fixed cost carried through commissioning (ramp)	+0.03	+0.05	+0.07
N ₂ O abatement-train maintenance (CON only)	+0.02	—	—
Net increase in fixed operating cost	+0.33	+0.39	+0.56

The first part, the one that scales with capital, is close to zero, and that is the counter-intuitive result. Maintenance and insurance are charged as a percentage of installed equipment, so a larger plant might be expected to carry a proportionally larger bill. It does not, for two reasons. The capital increase (§4.2.1) sits in the off-site and engineering lines, which carry no maintenance charge; and the installed equipment itself is no larger than Yeo's on the electrified routes, because this work installs it at lower factors than his (Table 3.4). Only the conventional route moves up, by \$0.02/kg, where the corrected compressor cost is large enough to lift the equipment base. The biomass route sits at the line, its smaller equipment base against Yeo's, having no hydrogen-export compressor train at all (§4.2.3), all but cancelling the wider set of lines charged on it. The near-zero entry is therefore not that nothing changed, but a real equipment-cost increase cancelled by a lower installation factor.

The second part is most of the increase, and it is the simplest: cost lines Yeo's estimate does not include. His model carries five fixed-cost lines, labour, maintenance, insurance, supervision and overheads; the factorial method used here carries about a dozen.¹ The same plant on a longer list.

The third part, near \$0.07/kg and much the same on every route, is that for the handful of lines both estimates share this work uses a slightly fuller basis: a few more operators and a larger overhead multiplier on their wages. Yeo carries one flat labour figure for all three plants, whereas here the operator count follows Turton's correlation on the number of process steps [171]. That correlation takes the square root of the step count, so plants of very different capital cost still come out within a few operators of one another, which is why the term barely varies between routes.

¹The additions are rent of land, laboratory charges, operating supplies, interest on working capital and environmental charges, together with three costs tied to turnover rather than to the plant: patents and royalties, distribution and selling, and research and development. Those three are grouped under fixed production cost by the costing method and are kept in that line here, although they scale with turnover rather than with capital or labour.



Heat-exchange networks

This chapter supports the heat integration of §3.2.4. It gives the stream properties and curation rules the pinch analysis runs on, the network designed from them in grid representation, the exchanger-by-exchanger detail and the pinch stream lists.

D.1. Stream properties and stream-list curation

The pinch analysis of §3.2.4 runs on a curated stream list whose film coefficients and crystallisation duties are set rather than inherited.

D.1.1. Individual film coefficients

Exchanger area, and through it exchanger capital, follows from the overall heat transfer coefficient of each match, so the coefficients are set per stream rather than left at the tool default. Aspen Energy Analyzer applies a stock film coefficient of $720 \text{ kJ/h m}^2 \text{ }^\circ\text{C}$ ($200 \text{ W/m}^2 \text{ }^\circ\text{C}$) to every process stream. That is the gas line of Towler & Sinnott [64] (Fig. 19.1) and is an order of magnitude too low for the aqueous service that dominates all three flowsheets, so it would oversize almost every exchanger.

Each stream is instead assigned an individual film coefficient h by fluid category, taken as the midpoint of the range Fig. 19.1 gives for that category and rounded to a round value in $\text{W/m}^2 \text{ }^\circ\text{C}$ (Table D.1). Figure 19.1 is the right source because the tool takes one coefficient per stream irrespective of what it is matched against, and forms the overall coefficient for a match itself from the two sides,

$$\frac{1}{U} = \frac{1}{h_{\text{hot}}} + \frac{1}{h_{\text{cold}}}, \quad (\text{D.1})$$

following the same chapter (Eq. 19.1). No separate fouling resistance is added, because the values in Fig. 19.1 already carry a fouling allowance (Ch. 19, p. 825). Utility-side coefficients are left at the tool defaults, 13,489 for cooling and chilled water and 21,585 for steam, both of which fall inside the same published ranges. Those ranges are wide, so every coefficient used here is fixed only to within the spread Fig. 19.1 reports, and it is that spread rather than the midpoint convention that the exchanger areas inherit.

Categories are assigned from each stream's composition in the converged Aspen report. Four streams are overridden by hand where a phase reading alone would misclassify them: the cyclohexane coolers of the conventional and electrochemical routes are organic liquid rather than gas, since the stream sits below its boiling point; the conventional route's nitric-acid stripper recycle is aqueous rather than gas; and one conventional cooler is a developing slurry, because adipic acid crystallises out across it. As a check on the two sources being consistent, an aqueous-to-aqueous match gives $U = 3600 \text{ kJ/h m}^2 \text{ }^\circ\text{C}$ ($1000 \text{ W/m}^2 \text{ }^\circ\text{C}$) from Eq. (D.1), inside the 800 to $1500 \text{ W/m}^2 \text{ }^\circ\text{C}$ that Table 19.1 of the same chapter quotes for water-to-water service. The coefficients change area and therefore capital, but not the utility targets, which are set by duty alone.

Table D.1: Individual film coefficients by fluid category, taken as the midpoint of the range Towler & Sinnott [64] give for each category (Ch. 19, Fig. 19.1). The per-stream assignment is listed in §D.3.

Category	h (kJ/h m ² °C)
Condensing or boiling aqueous (column condensers, reboilers)	9,000
Dilute aqueous broth	7,200
Alkali solutions, dissolved-acid mother liquor	6,300
Concentrated HNO ₃ -water	5,400
Light organic liquid	4,500
Crystalliser slurry, agitated reactor jacket	2,880
High-pressure gas	1,440
Low-pressure dilute gas (vents, off-gas)	540

D.1.2. Crystallisation duty

Adipic acid crystallises out of solution in every route's recovery train, and the conventional and electrochemical routes redissolve it at 99 °C and crystallise it a second time at 17 °C. Each of those steps exchanges latent heat and so belongs in the stream list. The heat is taken as the enthalpy of fusion of adipic acid, $\Delta H_{\text{fus}} = 34.85$ kJ/mol [107], or 9.6806 kW per kmol/h. Strictly, what is released in the crystalliser is the enthalpy of solution rather than of fusion, since the acid comes out of an aqueous solution and not out of a pure melt; the fusion value is used as a first approximation to it. *t*3HDA and muconic acid map onto the same C₆H₁₀O₄ surrogate in the inherited property model, so the same value serves the biomass train.

Each crystalliser enters the stream list as its own isothermal stream at its own temperature, with the duty from the net solid formed across the block,

$$Q_{\text{cryst}} = |\Delta \dot{n}_{\text{solid}}| \Delta H_{\text{fus}}, \quad (\text{D.2})$$

with $\Delta \dot{n}_{\text{solid}}$ the solid molar flow leaving the block minus that entering it, and ΔH_{fus} the enthalpy of fusion above. The net is used rather than the absolute solid flow out of the crystalliser, which magma recirculation inflates by an order of magnitude. Crystallisation gives a hot stream and redissolution a cold one, so a cooling and a heating load enter at different temperature levels; carrying only their net would move the pinch. Sub-ambient crystallisers are served by chilled water and the 2 °C purge crystallisers as a refrigeration load.

Each crystalliser runs isothermally, its feed already brought to temperature by an exchanger upstream, so the duty the block reports is the phase change alone and the sensible conditioning stays in the separate exchanger rows. These latent duties replace what the blocks report rather than adding to it. No inherited model imposes a fixed heat of fusion on the solid, so what the blocks return is a liquor-specific heat-of-mixing term that varies with temperature and composition: near zero at the cold, dilute crystallisers of the conventional and electrochemical routes, 49 kJ/mol at the conventional route's hot dissolver in concentrated nitric acid, near zero at the electrochemical one, and 23 kJ/mol in the biomass broth. None is the fusion enthalpy and they do not agree with each other, which follows from Yeo (2025) [9] placing detailed crystallisation behaviour outside his scope rather than from an error. Because only the crystalliser rows are touched, and solid flow is unchanged across every cooler feeding them, those exchangers change temperature but not phase and the latent term is not counted twice.

D.1.3. Curation of the extracted stream list

Aspen Energy Analyzer extracts the stream list automatically, but that extraction picks up blocks that are modelling constructs rather than plant equipment, so each list is curated before use. Dropped are the bridge blocks Yeo inserted to reconcile the NRTL and electrolyte-NRTL enthalpy reference states, whose reported duties are bookkeeping corrections for physics already counted in the surrounding equipment; the idealised component separators, whose duties only reconcile the mass balance; the placeholder utility tiers; adiabatic units carrying no duty; and electrically driven duties, which are charged as electricity rather than as a thermal utility, following Yeo (2025) [9] (Appendix A). The same rules apply to every route so that utility costs and steam-related emissions stay comparable.

D.2. Heat-exchange network grid diagram

Figure D.1 shows the conventional route's network in grid representation, the form in which it was designed. Hot streams run left to right across the upper half in red and cold streams in the lower half in blue, each on its own horizontal line with temperature falling to the right. A vertical link joining two lines is one exchanger: process-process where it joins a red line to a blue one, and a utility match where it terminates on one of the utility lines at

the top or bottom. The diagram makes the structure of the recovery visible in a way the exchanger table cannot, in particular how much of the reboiler duty is met by process heat rather than steam.




Figure D.1: Heat-exchange network for the conventional route, grid representation. Hot streams run left to right in red, cold streams in blue; each vertical link is an exchanger, and a link terminating on one of the utility lines at the top or bottom is a utility match. Designed on the stream list of §D.3 under the rules of Section 3.2.4.




Figure D.2: Heat-exchange network for the electrochemical route, grid representation, on the same conventions as Figure D.1. The heavier heating duty of this route shows in the density of matches on the cold side, where the two column reboilers dominate.

D.3. Pinch stream lists

These are the inputs to the pinch analysis: every process stream that has to be heated or cooled, before any recovery is designed. What the analysis then produces from them, the exchangers and their duties, is listed in §D.4. The curated lists, 32 streams on the conventional route, 32 on the electrochemical, 15 on the biomass route to *t*3HDA and 19 on the route to adipic acid, are given below, one table per route. Each lists the stream's inlet and outlet temperature, whether it is hot or cold, its duty and its individual film coefficient h , the last assigned by fluid category as described in Section 3.2.4; the overall coefficient for a match is formed from the two streams' values. Duties are the block enthalpy change, except at the crystallisers, where the imposed heat of fusion replaces it.




Figure D.3: Heat-exchange network for the biomass route to *t*3HDA, grid representation, on the same conventions as Figure D.1. The crystallisation train sits on the chilled-water line at the top; the two steam levels at the foot serve the recovery columns.




Figure D.4: Heat-exchange network for the biomass route to adipic acid, grid representation, on the same conventions as Figure D.1. The additional streams relative to Figure D.3 are those of the polishing reactor and the adipic-acid recovery described in §3.2.3.

Stream	T_{in} (°C)	T_{out} (°C)	H/C	Duty (MW)	h	Stream	T_{in} (°C)	T_{out} (°C)	H/C	Duty (MW)	h
HX10	600.1	185	H	4.367	1440	HX42	19	87.1	C	17.171	5400
HX11	185	35	H	2.017	1440	HX44	18.9	87.2	C	1.693	5400
HX21	127.6	40	H	3.492	4500	HX46	94	99	C	0.066	2880
HX22	127.6	20	H	11.507	540	HX45_hot	99	50.2	H	0.891	2880
HX23	74.4	35	H	5.279	4500	HX45_cold	22.7	94	C	0.891	2880
HX24	224	35	H	1.550	540	R3	100.5	99.5	H	3.801	2880
HX31	154.4	30	H	0.654	4500	R4	70.5	69.5	H	11.212	2880
HX32	46.3	25	H	0.724	5400	CRYS1 ^f	17.5	16.5	H	0.800	2880
HX41	23.7	15	H	2.736	2880	CRYS2 ^f	17.5	16.5	H	0.022	2880
HX43	44.9	15	H	0.218	2880	CRYS3 ^f	17.5	16.5	H	0.811	2880
HX47	19	15	H	1.366	2880	CRYS4 ^f	98.5	99.5	C	0.822	2880
HX48	68.8	40	H	5.685	9000	COL1_cond	73.9	72.9	H	0.331	9000
HX49	87.1	45	H	18.048	5400	COL1_reb	153.9	154.9	C	4.419	9000
HX52	216.7	35	H	2.083	540	COL2_cond	69.3	68.3	H	6.722	9000
HX53	57.3	20	H	0.139	540	COL2_reb	93.3	94.3	C	13.798	9000
HX55	796	280	H	1.701	540	CNDUTYS	18.8	29.2	C	0.077	540

Table D.2: Conventional (HNO_3 oxidation): pinch stream list, 32 streams. H/C is hot or cold; h is the individual film coefficient in $\text{kJ/h m}^2 \text{ } ^\circ\text{C}$. ^f duty is the imposed heat of fusion.

Stream	T_{in} (°C)	T_{out} (°C)	H/C	Duty (MW)	h	Stream	T_{in} (°C)	T_{out} (°C)	H/C	Duty (MW)	h
HX11	718	185	H	5.806	1440	HX44_cold	18.6	77	C	9.759	7200
HX12	185	35	H	2.020	1440	HX45	76.1	75.1	H	47.099	9000
HX21	127	40	H	3.626	4500	HX46	57.8	15	H	0.442	2880
HX22	127	20	H	11.378	540	HX47_hot	96.8	95.5	H	0.478	7200
HX23	71.8	35	H	5.568	4500	HX47_cold	25	91.8	C	0.478	7200
HX24	224.3	35	H	1.550	540	HX48	95	105	C	1.624	2880
HX31	95.8	25	H	0.876	4500	HX49	25.5	13	H	2.152	2880
HX32	44.1	25	H	5.548	7200	R3	100.5	99.5	H	3.499	2880
HX33	113.1	45	H	0.280	1440	CRYS1 ^f	17.5	16.5	H	0.829	2880
HX35	100.4	25	H	23.028	7200	CRYS2 ^f	17.5	16.5	H	0.025	2880
HX41_hot	99.5	96.8	H	45.114	9000	CRYS3 ^f	17.5	16.5	H	0.835	2880
HX41_cold	93.8	94.1	C	45.114	7200	CRYS4 ^f	98.5	99.5	C	0.845	2880
HX42_hot	94.1	93.5	H	46.461	9000	COL1_cond	72.5	71.5	H	0.306	9000
HX42_cold	86.5	88.2	C	46.461	6300	COL1_reb	153.6	154.6	C	4.640	9000
HX43	24.4	15	H	3.195	2880	COL2_cond	100	99	H	33.196	9000
HX44_hot	88.2	86	H	9.759	9000	COL2_reb	99.5	100.5	C	58.265	9000

Table D.3: Electrochemical oxidation: pinch stream list, 32 streams. H/C is hot or cold; h is the individual film coefficient in $\text{kJ/h m}^2 \text{ } ^\circ\text{C}$. ^f duty is the imposed heat of fusion.

Stream	T_{in} (°C)	T_{out} (°C)	H/C	Duty (MW)	h	Stream	T_{in} (°C)	T_{out} (°C)	H/C	Duty (MW)	h
HX11_hot	95	30	H	19.057	7200	HX22_hot	97	29.6	H	3.283	6300
HX11_cold	25	81.6	C	19.057	7200	HX22_cold	22	92	C	3.283	2880
HX12	81.6	95	C	11.035	7200	HX23	95.9	100	C	1.007	2880
HX13	80	25	H	52.589	7200	HX24	18.8	15	H	1.209	2880
HX14	80	25	H	1.439	1440	CRYS2 ^f	17.5	16.5	H	0.803	2880
HX15	204.2	25	H	0.243	1440	CRYS3 ^f	96.5	97.5	C	0.826	2880
HX16	171.8	25	H	0.172	1440	BCRYS3 ^f	96.5	97.5	C	0.011	2880
HX21	25	15	H	3.310	7200						

Table D.4: Biomass to *l*3HDA: pinch stream list, 15 streams. H/C is hot or cold; h is the individual film coefficient in $\text{kJ/h m}^2 \text{ } ^\circ\text{C}$. ^f duty is the imposed heat of fusion.

Stream	T_{in} (°C)	T_{out} (°C)	H/C	Duty (MW)	h	Stream	T_{in} (°C)	T_{out} (°C)	H/C	Duty (MW)	h
HX11_hot	95	30	H	19.057	7200	HX23	95.9	100	C	1.007	2880
HX11_cold	25	81.6	C	19.057	7200	HX24	18.8	15	H	1.209	2880
HX12	81.6	95	C	11.035	7200	BCRYS3 ^f	96.5	97.5	C	0.011	2880
HX13	79.1	25	H	52.830	7200	COOL-AA	70	17	H	2.636	6300
HX14	79.1	25	H	0.450	1440	AA-CRYS ^f	17.5	16.5	H	0.743	2880
HX15	204.2	25	H	0.081	1440	HX-POLIS	24.2	70	C	2.168	2880
HX16	171.8	25	H	0.058	1440	R-POLISH	70.5	69.5	H	3.741	2880
HX21	25	15	H	3.310	7200	CRYS3 ^f	96.5	97.5	C	0.826	2880
HX22_hot	97	29.6	H	3.283	6300	CRYS2 ^f	17.5	16.5	H	0.803	2880
HX22_cold	22	92	C	3.283	2880						

Table D.5: Biomass to adipic acid: pinch stream list, 19 streams. H/C is hot or cold; h is the individual film coefficient in $\text{kJ/h m}^2 \text{ } ^\circ\text{C}$. ^f duty is the imposed heat of fusion.

Four groups of blocks are excluded from every list, on the reasoning set out in Section 3.2.4: the two property-method bridge blocks Yeo inserted to reconcile the NRTL and electrolyte-NRTL enthalpy reference states, whose reported duties, about -2.3 and $+1.6$ MW, are reference-state bookkeeping for physics already counted in the surrounding equipment; the idealised component separators, whose duties reconcile the mass balance rather than exchange heat; the two placeholder utility tiers carried in the inherited Aspen Energy Analyzer project; and adiabatic units, mixers, splitters, valves and flashes, that carry no duty.

One reconciliation against the results chapter follows from how generated steam is booked. Duty recovered as generated medium-pressure steam is credited on the hot side and appears in neither cooling column of Table 4.4, so the hot-side duties listed here exceed the cooling shown there on the two routes that raise it, by 2.45 MW on the conventional route and 1.32 on the electrochemical one. The biomass routes raise none and close exactly.

D.4. Exchanger-by-exchanger detail

These are the outputs of the pinch analysis, designed on the stream lists of §D.3. Table D.6 gives each designed network at the level the costing uses it: the exchanger and shell counts, the total area, and the duty exchanged. The listing behind those totals follows below, route by route: every exchanger with its hot and cold stream, load, area, shell count and LMTD, 179 in all, whose totals reconcile against this table.

Table D.6: Designed heat-exchange networks, per route. Areas are the total across shells; duty is the total exchanged. Installed cost is on the floating-head basis of §3.2.4. The exchanger-by-exchanger listing follows below.

Route	Exch.	Shells	Area (m ²)	Duty (MW)	Installed (\$M)
CON	54	89	13,919	89.6	27.0
ECO	53	99	24,299	307.6	39.8
BIO _{t3HDA}	31	49	5,432	92.3	12.9
BIO _{AA}	41	66	7,869	98.5	16.2

Exchangers are costed per shell on the Towler & Sinnott basis (Table 7.2): floating head, unless the match involves chilled water, in which case plate-and-frame; carbon steel; CEPCI January-2010 base 532.9 escalated to 800; installation factor 3.20. The dominant uncertainty is that type assumption. Costing every shell-and-tube unit as U-tube instead lowers the network capital by about 17 %, to \$22.3, \$32.4, \$10.7 and \$13.6 M against the floating-head values in Table D.6; the chilled-water matches stay plate-and-frame throughout. That is a swing larger than the net effect on levelised cost of adipic acid itself on three of the four routes, which is why the sign of that effect is not over-read in §4.2.2. It does not touch the GHG result, which depends on the utility duties rather than on how the exchangers are costed.

Loads are per exchanger and area is the total across shells. Process-process matches carry a stream name on both sides; a utility name on either side marks a utility match. Within each route the exchangers are sorted by area.

D.4.1. CON: conventional, HNO₃ oxidation

54 exchangers, 89 shells, 13919 m² total area, 89.6 MW total exchanged duty.

Exch.	Hot	Cold	Load (MW)	Area (m ²)	Sh.	LMTD (°C)	Exch.	Hot	Cold	Load (MW)	Area (m ²)	Sh.	LMTD (°C)
E-144	HX22	Cooling Water	5.832	1951.6	4	31.2	E-169	HX11	Cooling Water	0.701	62.7	2	49.3
E-132	HX22	HX42	1.293	1876.5	4	7.1	E-139	HX11	CRYS4	0.822	62.5	1	63.8
E-123	HX22	COL2_reb	3.240	1713.1	4	19.2	E-129	LP Steam	COL2_reb	3.284	60.7	1	259.2
E-146	HX22	HX42	0.774	1474.3	4	5.2	E-158	HX32	Cooling Water	0.588	59.9	1	46.8
E-137	HX49	HX42	6.723	1337.6	3	51.1	E-159	HX53	Cooling Water	0.102	52.8	2	7.9
E-126	R3	COL2_reb	3.801	1028.0	3	35.4	E-170	HX21	Cooling Water	1.368	51.5	1	117.9
E-136	R4	Cooling Water	11.212	396.3	1	262.3	E-165	HX32	HX44	0.136	39.9	1	13.3
E-161	HX41	Chilled water	2.736	379.5	1	66.0	E-133	HX55	COL1_reb	1.701	38.0	1	170.7
E-168	HX49	Cooling Water	9.869	311.4	1	274.6	E-151	HX43	HX42	0.160	36.4	2	15.0
E-142	COL2_cond	HX42	6.722	280.1	1	201.5	E-141	HX21	HX42	0.815	28.7	1	92.8
E-130	HX23	Cooling Water	5.279	264.3	1	164.8	E-140	HX31	COL2_reb	0.300	18.4	2	38.2
E-145	HX52	COL2_reb	0.826	240.3	2	24.4	E-166	HX11	HX45_cold	0.137	18.3	2	17.5
E-163	HX47	Chilled water	1.366	231.0	1	46.8	E-147	HX31	HX42	0.351	17.0	2	46.0
E-157	HX45_hot	HX45_cold	0.754	193.6	5	22.1	E-128	HX10	COL1_reb	1.893	15.8	1	271.4
E-156	HX24	COL2_reb	0.591	170.9	2	22.1	E-134	HX21	HX44	0.102	13.9	1	15.2
E-167	HX52	Cooling Water	0.714	160.3	2	27.9	E-175	HX53	Chilled water	0.037	13.3	1	5.6
E-173	HX22	Chilled water	0.367	154.7	1	16.5	E-172	HX45_hot	Cooling Water	0.137	7.3	1	24.4
E-154	HX48	Cooling Water	5.685	147.9	1	262.9	E-148	COL1_cond	HX42	0.331	6.9	1	59.6
E-138	HX49	HX44	1.456	137.4	3	58.9	E-122	HX52	COL2_reb	0.110	6.6	1	20.0
E-164	CRYS3	Chilled water	0.811	135.1	1	39.8	E-174	HX43	Chilled water	0.058	6.6	1	10.6
E-162	CRYS1	Chilled water	0.800	133.2	1	39.6	E-153	LP Steam	CNDUTYS	0.077	5.3	1	14.7
E-171	HX24	Cooling Water	0.511	114.6	2	24.9	E-127	HX24	COL2_reb	0.079	4.5	1	15.5
E-124	HX21	COL2_reb	1.208	108.4	2	61.0	E-160	CRYS2	Chilled water	0.022	3.6	1	4.4
E-131	HX52	COL1_reb	0.432	101.1	1	25.4	E-152	LP Steam	HX46	0.066	3.4	1	13.5
E-143	HX10	MP Steam Generation	2.449	82.4	2	146.8	E-150	HX10	COL1_reb	0.024	2.2	1	5.2
E-135	HX24	COL1_reb	0.369	79.8	1	25.0	E-125	HX31	HX42	0.003	0.1	1	0.8
E-155	HX11	COL2_reb	0.358	79.6	2	21.9	E-149	HX24	COL1_reb	0.000	0.0	1	0.0
Total			89.6	13919	89								

Table D.7: Conventional (HNO₃ oxidation): final HEN exchangers, sorted by area.

D.4.2. ECO: electrochemical oxidation

53 exchangers, 99 shells, 24299 m² total area, 307.6 MW total exchanged duty.

Exch.	Hot	Cold	Load (MW)	Area (m ²)	Sh.	LMTD (°C)	Exch.	Hot	Cold	Load (MW)	Area (m ²)	Sh.	LMTD (°C)
E-162	COL2_cond	HX41_cold	33.196	5315.9	11	5.6	E-128	HX12	HX48	0.913	86.2	1	41.6
E-151	HX22	Cooling Water	7.588	2469.0	6	22.0	E-179	HX11	COL1_reb	4.488	84.0	2	155.3
E-127	HX41_hot	HX41_cold	11.919	2081.8	5	5.2	E-167	HX12	Cooling Water	0.860	75.7	2	31.7
E-139	HX41_hot	HX42_cold	21.890	2045.9	5	10.4	E-171	HX21	Cooling Water	2.516	72.7	1	37.2
E-132	HX22	COL2_reb	2.506	1532.3	4	11.9	E-131	HX12	COL2_reb	0.195	71.8	1	8.2
E-176	HX42_hot	HX42_cold	9.894	1433.5	3	6.7	E-156	HX42_hot	HX42_cold	0.350	53.3	1	6.4
E-133	LP Steam	COL2_reb	54.582	1260.7	3	24.5	E-154	HX31	Cooling Water	0.827	43.5	2	21.0

Exch.	Hot	Cold	Load (MW)	Area (m ²)	Sh.	LMTD (°C)	Exch.	Hot	Cold	Load (MW)	Area (m ²)	Sh.	LMTD (°C)
E-175	HX41_hot	HX42_cold	10.828	1022.9	3	10.3	E-150	HX46	Cooling Water	0.297	41.5	2	11.7
E-153	HX35	Cooling Water	21.580	886.4	2	19.3	E-145	HX24	COL1_reb	0.152	38.0	1	28.3
E-136	HX22	HX48	0.326	751.5	2	3.7	E-159	HX21	HX48	0.097	32.1	1	6.2
E-158	HX45	Cooling Water	47.099	651.7	2	48.2	E-166	HX35	Chilled water	0.850	31.3	1	20.8
E-144	R3	HX42_cold	3.499	521.0	2	12.2	E-142	HX12	HX48	0.025	24.4	1	4.1
E-152	HX32	Cooling Water	4.389	491.4	2	7.7	E-130	HX24	CRYS4	0.255	18.6	1	108.7
E-157	HX43	Chilled water	3.195	445.1	1	11.5	E-174	HX33	Cooling Water	0.280	17.0	1	45.7
E-164	HX22	Chilled water	0.957	356.9	1	18.7	E-155	HX41_hot	HX47_cold	0.478	16.4	2	26.2
E-146	HX49	Chilled water	2.152	334.0	1	10.6	E-163	HX46	Chilled water	0.145	14.2	1	15.6
E-140	HX42_hot	Cooling Water	27.529	285.1	1	64.4	E-138	HX11	MP Steam Generation	1.318	7.3	1	480.5
E-177	HX32	HX44_cold	1.072	259.9	2	5.1	E-147	HX35	Chilled water	0.372	6.1	1	47.9
E-172	HX23	Cooling Water	5.568	254.3	1	23.6	E-135	HX12	HX48	0.006	5.2	1	4.3
E-141	HX42_hot	HX44_cold	8.688	211.5	1	37.1	E-178	HX47_hot	Cooling Water	0.478	5.2	1	70.8
E-161	HX24	CRYS4	0.591	193.7	2	24.3	E-168	COL1_cond	Cooling Water	0.306	4.4	1	46.6
E-149	CRYS3	Chilled water	0.835	143.9	1	8.9	E-160	CRYS2	Chilled water	0.025	4.3	1	8.9
E-148	CRYS1	Chilled water	0.829	142.9	1	8.9	E-165	HX32	Chilled water	0.087	3.1	1	21.4
E-169	HX24	Cooling Water	0.553	118.1	2	32.7	E-129	HX21	HX48	0.031	2.6	1	24.5
E-134	HX21	COL2_reb	0.982	111.6	2	10.9	E-173	HX31	Chilled water	0.050	2.5	1	21.4
E-143	HX35	HX48	0.226	110.2	1	3.6	E-137	HX12	Cooling Water	0.021	0.8	1	71.1
E-170	HX44_hot	Cooling Water	9.759	105.4	1	61.7							
Total			307.6	24299	99								

Table D.8: Electrochemical oxidation: final HEN exchangers, sorted by area.

D.4.3. BIO_{t3HDA}: biomass to the intermediate

The BIO network was transcribed from the Aspen Energy Analyzer Heat Exchangers tab without stream names, so hot/cold streams and loads are not reproduced here; area and shell count are the costed quantities and are complete. 31 exchangers, 49 shells, 5432 m² total.

Exch.	Area (m ²)	Sh.	Exch.	Area (m ²)	Sh.	Exch.	Area (m ²)	Sh.
E-144	1769.0	4	E-125	113.6	1	E-142	13.4	3
E-130	604.0	3	E-145	111.0	2	E-126	8.0	1
E-128	442.6	2	E-149	106.2	1	E-121	1.8	1
E-150	379.4	2	E-131	75.8	1	E-120	1.1	1
E-127	377.4	2	E-132	56.9	1	E-135	0.6	1
E-129	263.5	3	E-134	54.7	1	E-146	0.5	1
E-143	234.5	2	E-136	43.2	1	E-140	0.4	1
E-141	226.3	1	E-122	22.4	3	E-137	0.1	1
E-139	187.0	1	E-123	20.4	1	E-147	0.0	1
E-138	145.7	1	E-133	14.1	2			
E-124	145.0	2	E-148	13.9	1			
Total			5432	49				


Table D.9: Biomass to t3HDA: final HEN exchangers, sorted by area.

D.4.4. BIO_{AA}: biomass to adipic acid

41 exchangers, 66 shells, 7869 m² total area, 98.5 MW total exchanged duty.

Exch.	Hot	Cold	Load (MW)	Area (m ²)	Sh.	LMTD (°C)	Exch.	Hot	Cold	Load (MW)	Area (m ²)	Sh.	LMTD (°C)
E-153	HX13	Cooling Water	39.274	2306.4	6	14.0	E-135	HX11_hot	HX22_cold	0.271	44.5	1	11.5
E-139	HX11_hot	HX11_cold	4.679	748.6	4	7.8	E-149	LP Steam	CRYS3	0.826	43.2	1	27.1
E-140	HX13	HX11_cold	7.482	726.5	2	11.9	E-137	HX22_hot	HX12	0.117	29.3	1	4.3
E-136	HX11_hot	HX11_cold	1.954	670.1	2	3.7	E-164	COOL-AA	Chilled water	0.543	28.8	1	16.0
E-133	HX11_hot	HX12	2.479	435.4	1	6.5	E-126	LP Steam	HX12	1.022	22.3	1	30.6
E-160	HX11_hot	Cooling Water	7.291	420.0	1	13.6	E-145	HX22_hot	HX11_cold	0.748	18.2	1	44.9
E-141	HX11_hot	HX22_cold	2.385	318.8	3	15.9	E-131	LP Steam	HX22_cold	0.475	17.9	1	37.7
E-155	HX21	Chilled water	3.310	238.8	1	11.1	E-134	HX22_hot	HX11_cold	0.037	6.5	1	6.1
E-158	HX24	Chilled water	1.209	226.2	1	8.3	E-151	HX15	Cooling Water	0.045	5.2	3	24.7
E-142	HX13	HX-POLIS	2.168	192.4	2	22.3	E-150	HX16	Cooling Water	0.039	4.5	3	24.5
E-129	HX22_hot	HX12	0.438	181.8	3	3.2	E-165	HX14	Chilled water	0.033	4.4	1	20.9
E-146	R-POLISH	HX11_cold	3.741	175.8	1	37.3	E-132	HX15	LP Steam Generation	0.035	3.7	2	25.3
E-159	CRYS2	Chilled water	0.803	145.6	1	8.4	E-143	HX16	LP Steam Generation	0.017	3.2	2	14.4
E-152	COOL-AA	Cooling Water	2.037	143.6	2	12.7	E-157	COOL-AA	Chilled water	0.056	1.3	1	37.0
E-166	HX13	Chilled water	3.766	138.7	1	20.9	E-148	LP Steam	BCRYS3	0.011	0.6	1	27.1
E-154	AA-CRYS	Chilled water	0.743	134.8	1	8.4	E-127	LP Steam	HX22_cold	0.012	0.5	1	33.1
E-130	LP Steam	HX12	6.979	132.6	1	35.1	E-163	HX15	Chilled water	0.001	0.2	1	20.4
E-161	HX22_hot	Cooling Water	1.944	95.7	2	17.3	E-162	HX16	Chilled water	0.001	0.2	1	20.4
E-144	HX14	HX11_cold	0.417	93.7	3	15.5	E-156	HX15	Chilled water	0.001	0.0	1	56.0
E-147	LP Steam	HX23	1.007	54.7	1	26.1	E-128	HX16	Chilled water	0.000	0.0	1	161.3
E-138	HX13	HX22_cold	0.141	54.6	1	4.5							
Total			98.5	7869	66								

Table D.10: Biomass to adipic acid: final HEN exchangers, sorted by area.



Load-following under dynamic electricity prices

The load-following sensitivity of §3.5.5 asks whether running an electrified route at a load factor below unity, following the hourly electricity price, lowers its levelised cost of adipic acid or its scope-2 GHG intensity. The methodology states what is being tested and the two strategies that bound it, the produce-less baseline and the oversize case; the results are reported in §4.6. This chapter carries what sits underneath both: the procedure that sets the dispatch threshold, the inputs it runs on, and the tables the results chapter draws its figures from. Nothing here is a separate method from the one the methodology describes.

There are six sections, in the order the analysis runs. §E.1 gives the input, the split of each country's electricity price into the components the optimisation needs. §E.2 is the method that runs on it, the five-step numerical procedure that finds the dispatch threshold, which has no closed form and so is set out in full here rather than in the methodology. §E.3 and §E.4 carry the two tables the results chapter draws on, the effective electricity price and the scope-2 GHG intensity against load factor. §E.5 and §E.6 are the oversize case: the alternative strategy that holds annual output fixed by buffering an intermediate, its buffer sizing, and the cell cycling it causes.

E.1. Electricity network-cost decomposition

The effective 2024 industrial electricity price in each ARRRA country splits into the day-ahead commodity, per-MWh (volumetric) charges, and per-kW (contracted-capacity and peak) charges. Only the per-kW charges scale with $1/LF$ in the load-following optimisation (§3.5.5): a fixed annual capacity cost is spread over less output when the plant runs fewer hours, whereas the per-MWh charges are the same cost per kilogram at any load factor and cancel from $\Delta LCOA$. Table E.1 gives the components and the capacity subtotal p_{cap} used in the analysis; Germany's charge carries the graduated §19 Abs. 2 StromNEV relief [172], detailed in the table footnote.

Table E.1: Per-country electricity network-cost components (2024, €/MWh at baseload), from the E-Bridge buildup [100]. Only the per-kW capacity charge scales with load factor; the volumetric levies are flat per MWh and cancel from $\Delta LCOA$. The underlying line items are given in the E-Bridge buildup.

Component (€/MWh)	NL	DE	BE
Contracted / fixed capacity	7.58	1.99 ¹	0.62
Peak and monthly charges	10.37	–	0.77
Connection fee (Vastrecht)	0.01	–	–
Capacity subtotal p_{cap} ($\times 1/LF$)	17.96	1.99	1.39
Volumetric taxes & levies ²	1.88	3.28	5.70
ICC (CO ₂ compensation)	0.00	–35.34	–25.00
Non-commodity total	19.84	–30.07	–17.91

¹ Germany's network charge is §19-relieved to 10 % at $\geq 8\,000$ full-load hours, 15 % at $\geq 7\,500$ and 20 % at $\geq 7\,000$, and reverts to the full 19.87 (fixed) + 11.20 (variable) €/MWh below 7 000 FLH (§3.5.5). ² Flat per-MWh levies: NL energy tax; DE network-variable, §19, offshore, CHP, concession; BE management, control, market, excise, regional, green-certificate.

E.2. Threshold optimisation

The dispatch threshold p_{low} of \$3.5.5 is found by the procedure below. Steps 1 and 2 build the operating year and the price statistics, step 3 gives the LCOA change at a given load factor, step 4 selects the optimum, and step 5 applies it and gives the scope-2 GHG change.

The threshold is set to minimise the plant's cost per kilogram and has no closed form, so it is found numerically:

1. Set the operating year. The plant is designed for $H = 8\,000$ operating hours, so of the $N = 8\,784$ hours of 2024 it is down for 784 regardless of price. Those 784 hours are taken as one contiguous maintenance window, placed on the contiguous block whose mean price is closest to the annual mean; the remaining H hours are the operating year, and their prices are sorted ascending. Load factor is defined against H , so $LF = 1.00$ is 8 000 running hours and $LF = 0.875$ is 7 000, which is the basis the German §19 full-load-hour thresholds are written in.
2. For each load factor LF on a grid from 0.50 to 1.00, the plant runs the cheapest $k = \text{round}(LF H)$ hours of the operating year. In the sorted list these are $p_{(1)} \dots p_{(k)}$, so the cut-off threshold is the highest price among them, $P(LF) = p_{(k)}$, and the plant runs every hour at or below it. The commodity factor $\mu(LF)$ is the average price of the hours run divided by the annual average price over all N hours, which is what a baseload plant indifferent to timing would pay:

$$P(LF) = p_{(k)}, \quad \mu(LF) = \frac{\frac{1}{k} \sum_{i=1}^k p_{(i)}}{\frac{1}{N} \sum_{i=1}^N p_{(i)}} = \frac{\text{mean price of the hours run}}{\text{annual mean price}}. \quad (\text{E.1})$$

Because the window is placed at the annual mean rather than on peak prices, $\mu(1.00)$ is close to one.

3. Compute the LCOA change of running at that LF rather than flat baseload,

$$\Delta\text{LCOA}(LF) = \underbrace{\frac{C_{\text{fix}} \text{CRF} + C_{\text{OPEX}}^{\text{fix}}}{Q} \left(\frac{1}{LF} - 1 \right)}_{\text{whole-plant fixed cost, less output}} + \underbrace{w_{\text{elec}} p_{\text{com}} (\mu(LF) - 1)}_{\text{commodity saving } (<0)} + \underbrace{w_{\text{elec}} p_{\text{cap}} \left(\frac{1}{LF} - 1 \right)}_{\text{capacity charge}}, \quad (\text{E.2})$$

where C_{fix} is the fixed capital and CRF the annual charge the levelised-cost model places on it, 0.125 per unit of fixed capital: capital is drawn over the construction years and recovered against output that ramps, so the charge sits above the 0.102 a plain twenty-year annuity at 8% real would give. $C_{\text{OPEX}}^{\text{fix}}$ is the annual fixed operating cost (labour, maintenance, insurance, supervision, overheads), also as in Eq. (3.2); Q is the design-basis annual production, w_{elec} the route's total specific electricity demand (kWh per kg of product, independent of running hours), and $p_{\text{com}}, p_{\text{cap}}$ the commodity and capacity price components (decomposed per country in §E.1).

The three terms trade off as LF falls. The first is a penalty: the whole-plant fixed cost is unchanged in total but spread over less product, so it rises per kilogram through the growing $(1/LF - 1)$. The second is the saving: the hours run are cheaper than the annual average ($\mu < 1$), so it is negative and deepens as fewer, cheaper hours are kept. The third has the same form as the first, since the per-kW capacity charge is likewise fixed and spreads over less output as $(1/LF - 1)$; the per-MWh volumetric charges stay the same per kilogram at any load factor, so they do not affect the LCOA change. The optimal LF balances the two fixed-cost penalties against the commodity saving.

4. Take the load factor with the lowest ΔLCOA , the largest cost reduction over baseload (or baseload itself, if load-following does not lower the cost); this is the optimum LF^* , and its cut-off price is the threshold,

$$p_{\text{low}} = P(LF^*). \quad (\text{E.3})$$

5. Apply p_{low} to the operating year: classify each hour run or shut, and count the load factor, the throughput-weighted average price and carbon intensity, and the start-stops (off→on transitions). The scope-2 GHG change follows the same dispatch,

$$\Delta\text{GHG}(LF) = w_{\text{elec}} \overline{CI}_{\text{annual}} \left(\frac{\overline{CI}(LF)}{\overline{CI}_{\text{annual}}} - 1 \right), \quad (\text{E.4})$$

with $\overline{CI}(LF)$ the average intensity of the dispatched hours and $\overline{CI}_{\text{annual}}$ the country's annual mean.

E.3. Effective electricity price against load factor

Table E.2 traces the effective all-in price the flexible load pays as it is concentrated into cheaper hours. The commodity is the mean over the hours actually run, the per-kW capacity charge of Table E.1 is recovered over those hours and so rises as $1/LF$, and the volumetric levies are flat. Load factor is the share of the plant's 8 000 operating hours, so the German §19 tiers fall at $LF = 1.000$, 0.9375 and 0.875, which are 8 000, 7 500 and 7 000 full-load hours exactly.

Table E.2: Effective all-in electricity price (€/MWh) against load factor, built from the 2024 hourly day-ahead series [93] on the E-Bridge regulated cost level [100]. The turnaround window is placed on the contiguous block whose mean price is closest to the annual mean, the baseline rule of §4.6. The German step between $LF = 0.875$ and $LF = 0.85$ is the loss of §19 StromNEV relief at 7 000 full-load hours, which is why that column alone is not monotonic.

Load factor	running h/yr	Netherlands	Germany	Belgium
1.000	8 000	97.1	48.4	52.4
0.950	7 600	92.1	43.4	47.5
0.900	7 200	89.9	41.8	44.4
0.875	7 000	88.9	40.4	42.9
0.850	6 800	88.1	66.7	41.5
0.750	6 000	85.3	64.1	36.2
0.500	4 000	82.1	62.6	22.8

E.4. Scope-2 GHG intensity against load factor

The carbon side of the load-following sweep, reported in §4.6. Over the 8 784 hours of 2024 the lifecycle grid intensity averages 261 gCO₂e/kWh in the Netherlands, 341 in Germany and 150 in Belgium, with standard deviations of 112, 120 and 58 [94]; that spread is what a plant can act on, and it is why the Belgian saving is the smallest of the three. The reduction tracks the dispatch: the hours a plant drops first are the most expensive, and in these three markets they are among the most carbon-intensive as well, so the same curtailment that raises the levelised cost of adipic acid lowers the scope-2 intensity. The produce-less and oversize bounds (§E.5) nearly coincide here, because both concentrate the same flexible load into the same cleaner hours; what separates them on cost is capital, not carbon.




Figure E.1: Change in scope-2 GHG intensity against load factor, measured from $LF = 1.00$. Panels and line styles as in Figure 4.10. The two bounds nearly coincide, both concentrating the same flexible load into the same cleaner hours.

E.5. Load-following upper bound: cell oversizing and the intermediate buffer

The produce-less baseline (§3.5.5) load-follows by running the whole plant fewer hours and making less product. An alternative holds the annual output fixed by oversizing the flexible cell and buffering an intermediate; it gives

a larger saving, but only with a large, uncoded buffer, so it is reported as the optimistic upper bound. This section sets out that strategy, its cost model, and its buffer; the ECO cell start–stop cost, which applies in the baseline too, has its own section (§E.6).

E.5.1. Oversizing to hold output

To make the same annual product while running the cell only a fraction LF of the hours, the cell is built $1/LF$ larger and run at proportionally higher power when on, so annual production and annual electricity are *unchanged*; only the timing of the draw shifts into the cheaper hours. A lower load factor here does not mean less product; it means a larger cell running fewer, cheaper hours. Buffering is needed on both sides of the cell. Feed accumulates while the cell is off and is drawn down when it runs; product accumulates while it runs and is drawn down by the continuous downstream train while it is off. The feed is KA oil on the electrochemical route and muconic-acid broth on the biomass ones; the product is the cell effluent the recovery train crystallises. Only the electrochemical unit (the cell, plus electrodialysis for ECO) is oversized and buffered, so only its share of the route's electricity, the flexible fraction ϕ , follows the price. That share is large: electrolysis dominates the electricity bill while cooling is met as a thermal cooling-water duty, so from the Aspen balance $\phi = 0.949$ (ECO), 0.975 (BIO_{t3HDA}) and 0.977 (BIO_{AA}), the remainder being the continuous compressors. With only this fraction shifted, the LCOA change becomes

$$\Delta LCOA(LF) = \frac{C_{\text{cell}}(\text{CRF} + m)}{Q} \left(\frac{1}{LF} - 1 \right) + \phi w_{\text{elec}} p_{\text{com}} (\mu(LF) - 1) + \phi w_{\text{elec}} p_{\text{cap}} \left(\frac{1}{LF} - 1 \right), \quad (\text{E.5})$$

where C_{cell} is the installed capital of the oversized electrochemical block (the cell, and the electrodialysis unit on the electrochemical route), m the maintenance rate of the fixed-operating-cost model (§C.3.3), and ϕ the share of the route's electricity that block draws; CRF, Q , w_{elec} , p_{com} , $\mu(LF)$ and p_{cap} are as in Eq. (E.2).

The three terms have the same roles as in the produce-less baseline, two fixed-cost penalties growing with $1/LF$ and the commodity saving, but the expression differs in two ways. First, only the oversized cell block scales, so its annualised capital and maintenance replace the whole-plant fixed cost; the plant's labour, insurance and overheads are unchanged under oversizing, so they do not affect the LCOA change. Second, the flexible fraction ϕ weights the commodity and capacity terms, since only the cell, not the continuous train, follows the price. The optimum is found with the same five-step search as the baseline (§3.5.5), with this $\Delta LCOA$ in place of the whole-plant one. Because only the cell's capital scales, the saving is larger and the cost-optimal load factor lower than in the produce-less baseline; the two bracket the true optimum, which a joint design-and-scheduling optimisation would place between them [97].

E.5.2. The buffer, and why this is an upper bound

The buffer is sized from the running inventory: the intermediate fills at $(1/LF - 1)$ of the average throughput while the cell runs and drains at the average rate while it is off, so the working volume is the peak-to-trough of that level over the operating year. For the 2024 prices this is three to seven days of production (Table E.3). That is the product-side inventory; the feed side mirrors it, filling while the cell is off and draining while it runs, so the plant carries two working volumes of this order and neither is coded below. It stays that small only because the optimum sits close to baseload, where the cell is oversized by a few per cent and so overfills the buffer slowly; cheap hours cluster into long runs, up to 1 900 h for ECO and 2 900 h for BIO_{AA} without stopping, and a deeper optimum would fill the buffer through one of them. Even at three to seven days the capital is *not* included in the $\Delta LCOA$ above, so the oversize saving is an upper bound. The bound gains \$0.002–0.007/kg, which even a modest buffer would exceed. Flexible ammonia plants, whose synthesis loop is rigid in the way this recovery train is, carry storage worth a large share of plant capital, of order 40 % [98]. That is an analogy to a rigid continuous train fed by a flexible front end, not a costing of these vessels, and the argument does not rest on the figure: any buffer costing even a small fraction of plant capital is worth more than the saving.

E.6. ECO cell start–stop cost

Load-following shuts the cell in the expensive hours and restarts it when the price falls back (§3.5.5), so the ECO cell cycles on and off through the year. On the produce-less baseline the optimum is baseload in four of the six cases and $LF^* = 0.99$ in the other two, so cycling there is negligible (§4.6); the counts below are therefore taken

Table E.3: Load-following optima for the two electrified routes in the three ARRA countries, under the two operating bounds of §3.5.5. Changes are measured against $LF = 1.00$ on the same dispatch, so they are the load-following effect alone, and the buffer is not costed in $\Delta LCOA$. The produce-less optimum is baseload except on the electrochemical route in the Netherlands and Belgium, where it falls to $LF^* = 0.99$ and is worth 0.004 and 0.001 \$/kg.

Route	Country	p-less		oversize (upper bound)		
		LF^*	LF^*	$\Delta LCOA$	ΔGHG	buffer
ECO	Netherlands	0.99	0.97	-0.007	-0.022	4 d
ECO	Germany	1.00	0.94	-0.005	-0.043	6 d
ECO	Belgium	0.99	0.95	-0.006	-0.021	7 d
BIO _{AA}	Netherlands	1.00	0.98	-0.006	-0.018	3 d
BIO _{AA}	Germany	1.00	0.96	-0.002	-0.034	5 d
BIO _{AA}	Belgium	1.00	0.97	-0.004	-0.015	4 d

at the oversize optima, where cycling does occur. Each cycle accelerates cell-voltage degradation and brings the membrane replacement forward; the cost of one shut–restart cycle, C_{cycle} , is built from the measured degradation rate and this work’s membrane-replacement cost in four steps:

1. *Stack capital that ages.* Only the bare membrane stack degrades, priced at $c_{\text{stack}} = 209$ \$/kW [109, 108] over the cell power $P_{\text{cell}} = 32,732$ kW, so $C_{\text{stack}} = \$6.84$ M (balance of plant and electro dialysis excluded).
2. *Membrane replacement cost.* One swap is $f_{\text{rep}} = 0.35$ of that stack capital [109, 108]: $R = \$2.39$ M.
3. *Cell life per cycle.* A cycle adds $\delta_V = 5$ μV (the mid of 2–10 μV over 1 500 anion-exchange-membrane cold starts [96]) against $\Delta V_{\text{EOL}} = 0.10$ V at the central tier’s 1.9 V/cell [108], so $f = \delta_V / \Delta V_{\text{EOL}} = 2.6 \times 10^{-5}$.
4. *Cost per cycle.* $C_{\text{cycle}} = f R \approx \63 (€57).

This per-cycle cost is charged on every run-or-shut cycle: 127–219 cycles a year for ECO, an aggregate below \$0.00015/kg (§4.6). The biomass routes degrade by cathode corrosion rather than voltage rise, with no start–stop data, so BIO_{t3HDA} and BIO_{AA} carry no cycle cost. The stack cost, the replacement fraction and the degradation rate δ_V are all electrolyser analogues, taken for want of a measured ECO figure rather than from data on this cell.

Table E.4: ECO start–stop at each country’s cost-optimal load factor under the oversize bound. Each of the 127–219 yearly cycles costs $C_{\text{cycle}} \approx \63 (€57); the aggregate burden stays below \$0.00015/kg. The produce-less baseline optimises to baseload and cycles not at all.

Country	LF^*	p_{low} (€/MWh)	shut h/yr	cycles/yr	burden \$/kg
Netherlands	0.97	141	240	127	0.00008
Germany	0.94	130	479	219	0.00014
Belgium	0.95	122	400	175	0.00011

Under the oversize bound (§E.5) the continuous train stays on at its 0.70 turndown floor [91], so the cell can hold at that floor in an expensive hour rather than shut, avoiding some cycles, a three-state dispatch. Spread over an hour of floor electricity the hold is a premium $\Delta p = C_{\text{cycle}} / (\varphi_{\text{floor}} P_{\text{flex}}) = €58 / (0.70 \times 60.96 \text{ MW}) \approx €1.4/\text{MWh}$, so $p_{\text{high}} = p_{\text{low}} + \Delta p$: the cell runs full below p_{low} , holds at 70 % between, and shuts above p_{high} . Because Δp is about a euro per MWh the band is thin, and applied to the 2024 prices it moves the ECO saving by below \$0.001/kg, so the three-state refinement changes no conclusion.



Aspen Plus flowsheets

The converged Aspen Plus flowsheets behind the process models of Section 3.2 are reproduced here at full page width. They are included so that the block structure underlying the mass and energy balances can be read directly rather than inferred from the simplified process flow diagrams in the literature review. The property methods, the reactor model types and the specifications imposed on each block are given in Sections 3.2.1 and 3.2.2; this chapter adds the topology, not the specifications.

Two conventions carry across every sheet. Reactors, columns, crystallisers, centrifuges, compressors and exchangers keep the block names used throughout the thesis, so a block named in the methodology or in the heat-integration appendix can be located on the corresponding sheet. The conventional and electrochemical flowsheets are inherited from Yeo (2025) [9] without structural modification; the catalytic hydrogenation reactor and the dedicated adipic-acid recovery train that follow it are this thesis's extension (Section 3.2.3). Each sheet is the converged run the reported results are computed on. All four routes are shown: the conventional and electrochemical routes, the biomass route stopping at *l*3HDA, and the biomass route carried through to adipic acid.

The conventional sheet is reproduced without its per-stream temperature, pressure and flow annotations, so that it carries the same information as the other three; those values are reported in the stream tables rather than on the diagram.

F.1. Conventional route

Figure F.1 shows the conventional route. The reaction section runs on six RStoic blocks, R1 through R6, covering benzene hydrogenation, KA-oil air oxidation and nitric-acid oxidation in the sequence set out in Section 3.2.2. Downstream of them the sheet divides into the two RadFrac columns COL1 and COL2, the pair of RadFrac absorbers ADSB1 and ADSB2, and the product recovery train of four crystallisers CRY1 to CRY4 with three centrifuges CEN1 to CEN3. The five compressors CMPR1 to CMPR5 and the gas turbine TUR1 are worth locating, since restoring their cost correlations is the single largest addition to this route's installed equipment against the reference model (Appendix C). The HX blocks are the exchangers that the pinch analysis of Appendix D works on.

The conventional route is the largest of the four sheets and is shown in three parts, so that the block names stay legible at page width. The division follows the three sections the route is described by above: the reaction front end, the nitric-acid oxidation with its abatement train, and the product recovery train. Stream numbers run continuously across the parts, so a stream leaving the right-hand edge of one figure re-enters at the left-hand edge of the next under the same number.




Figure F.1: Aspen Plus flowsheet of the conventional route, part 1 of 3: benzene hydrogenation in R1 and cyclohexane air oxidation in R2 and R3, through KA-oil recovery at DECANT2 and COL1. Continues in Figure F.2.




Figure F.2: Aspen Plus flowsheet of the conventional route, part 2 of 3: nitric-acid oxidation in R4 and the N_2O abatement train, the absorbers ADSB1 and ADSB2, the decomposition reactors R5 and R6 and the gas turbine TUR1. Continued from Figure F.1, continues in Figure F.3.




Figure F.3: Aspen Plus flowsheet of the conventional route, part 3 of 3: the product recovery train of crystallisers CRY1 to CRY4 with the centrifuges CEN1 to CEN3. Continued from Figure F.2.

F.2. Electrochemical route

Figure F.4 shows the electrochemical route. Benzene and hydrogen enter at BEN-IN and H₂-IN and pass through the same hydrogenation and air-oxidation sequence as the conventional route, R1 to R3, with the decanter DE-CANT2 and the column COL1 recovering the KA-oil intermediate. The route then departs from the conventional one at R4: nitric-acid oxidation is replaced by the electrochemical cell, and the potassium-hydroxide electrolyte entering at KOHIN is recovered in the electrodialysis stack EDS1 with its two membrane units UEDS1A and UEDS1B. Downstream of the second column COL2 the product recovery train is the four crystallisers CRY1 to CRY4 with the centrifuges CEN1 to CEN3. The absence of any nitric-acid feed and of the ADSB absorber pair is what removes this route's N₂O abatement duty entirely.




Figure F.4: Aspen Plus flowsheet of the electrochemical route.

F.3. Biomass route to *t*3HDA

Figure F.5 shows the biomass route stopping at *t*3HDA. Muconic-acid broth enters at BROTH-IN and is concentrated through HX11–HX12 and the flash F1 before the electrochemical hydrogenation cell R1. The charge-balance blocks CB1 to CB3 and the correction blocks CORRECT1 to CORRECT3 carry the cell's ionic bookkeeping, CORRECT1 taking the ammonia stream NH3IN that restores pH across the property-method change described in Appendix G; CNDUTYS, CLR-ELEC and CLR-WTR collect its duties. Acidification at H2SO4IN and liming at LIMEIN bracket the recovery train of crystallisers CRY1 to CRY3, the batch crystalliser BCRY3 and the centrifuges CEN1 and CEN2. This sheet is the upstream half of the route in Figure F.6; everything downstream of the final centrifuge is the adipic-acid extension.






Figure F.5: Aspen Plus flowsheet of the biomass route stopping at r3HDA.

F.4. Biomass route to adipic acid

Figure F.6 shows the biomass route carried through to adipic acid. The upstream half is the *t*3HDA route of Figure F.5. The extension begins at H2-SPLIT, which divides the hydrogen header between the reactor feed and the sale stream H2-SALE, with the reactor feed compressed at CMPR4. The crystallised *t*3HDA is redissolved at M-REDIS and brought to reaction temperature at HX-POLIS, and the two streams meet in the hydrogenation reactor R-POLISH. Everything after it, the flashes F1 and F2, the crystallisers AA-CRYS and AA-CRYS2 and the centrifuge AA-CEN, is the dedicated adipic-acid recovery train, which has no counterpart in the reference model and is the reason BIO_{AA} carries a crystalliser line of its own.





Aspen Plus stream tables

The converged mass and energy balance behind each route is given here, so that the flowsheets of Appendix F can be read as a balance rather than as a topology. Every stream identifier in these tables is a line on the corresponding sheet, and the *From* and *To* columns name the blocks it connects.

Each route carries two tables. The first lists every stream with its source and destination block, temperature, pressure and total flow. The second gives the composition of the streams that cross the battery limits, the feeds and the products, which is what closes the balance: the adipic acid leaving, the benzene and hydrogen entering, and the waste streams the results chapter prices. Compositions of internal streams are not printed.

Three of the streams the biomass tables list as crossing the battery limits are not physical, and are reproduced here only because the flowsheet is cut at that point. The property method changes between streams 2-0 and 2-1, where *t3HDA* modelled on Joback-estimated properties is exchanged for *t3HDA* carrying adipic acid's, so the two halves appear as a product and a feed rather than as one internal stream. *NH3IN* restores the neutral pH that the exchange disturbs. Following Yeo (2025) [9], whose flowsheet this is, the ammonia it carries is treated as already present in the fermentation broth rather than as a purchase, since in a real plant it would have been dosed to the fermenters; it therefore carries no cost and no emission, and neither does the dummy cooler that resets the temperature across the same cut.

G.1. Conventional route

Stream	From	To	T (°C)	P (bar)	Mass flow (kg/h)	Molar flow (kmol/h)	Stream	From	To	T (°C)	P (bar)	Mass flow (kg/h)	Molar flow (kmol/h)
1-1	M11	R1	30.0	50.0	1 539	672.8	4-25	HX44	—	87.2	0.30	463.8	13.0
1-2	R1	HX10	600.1	50.0	10 336	447.6	4-26	M45	HX49	87.1	0.30	40 933	1 302
1-3	HX10	HX11	185.0	50.0	10 336	447.6	4-27	HX49	M46	n/a	n/a	0.00	0.00
1-4	HX11	S11	35.0	50.0	856.3	335.0	4-28	HX49	EOSCOR2	45.0	0.30	40 933	1 302
1-5	S11	M11	35.0	50.0	847.8	331.6	4-29	M46	CMPR5	43.0	0.30	205.4	7.39
1-6	S11	—	35.0	50.0	8.56	3.35	4-3	COL2	HX48	68.8	0.30	8 856	469.0
2-0	HX11	R2	35.0	50.0	9 480	112.7	4-30	CMPR5	—	199.5	1.00	205.4	7.39
2-1	R2	SP1	127.6	15.0	54 420	736.1	4-31	EOSCOR2	M32	45.0	0.30	40 933	1 302
2-10	HX22	SP2	20.0	15.0	25 239	865.3	4-32	CEN2	M43	18.9	1.00	329.1	2.25
2-11	HX22	M21	20.0	15.0	50 401	927.4	4-33	M43	HX45	22.7	1.00	18 667	432.6
2-12	SP2	M22	145.2	5.00	560.5	6.66	4-34	HX45	HX46	94.0	1.00	18 667	432.6
2-13	SP2	—	20.0	15.0	24 678	858.6	4-35	HX46	CRYS4	99.0	1.00	18 667	432.6
2-14	M22	HX23	74.4	1.00	41 220	590.3	4-36	CRYS4	HX45	99.0	1.00	18 667	432.6
2-15	HX23	M21	35.0	1.00	41 220	590.3	4-37	HX45	M44	50.2	1.00	18 667	432.6
2-16	M21	R2	19.0	1.00	91 621	1 518	4-38	M44	HX47	19.0	1.00	531 530	12 309
2-17	CMPR1	HX24	224.0	4.00	28 959	1 000	4-39	HX47	CRY53	15.0	1.00	531 530	12 309
2-18	HX24	CMPR2	35.0	4.00	28 959	1 000	4-4	HX48	M46	40.0	0.30	205.4	7.39
2-19	CMPR2	R2	229.0	15.0	28 959	1 000	4-40	CRY53	S43	17.0	1.00	531 560	12 310
2-2	SP1	HX21	127.6	15.0	51 736	715.8	4-41	S43	M44	17.3	1.00	512 870	11 877
2-3	SP1	—	127.6	15.0	2 683	20.3	4-42	S43	CEN3	17.3	1.00	18 696	432.9
2-4	HX21	R3	40.0	15.0	51 736	715.8	4-43	CEN3	CNDUTYS	18.8	1.00	6 449	349.2
2-5	R3	DECANT2	100.0	15.0	51 736	786.9	4-44	CEN3	—	18.8	1.00	12 246	83.8
2-6	DECANT2	COL1	107.5	15.0	49 565	674.0	4-45	CNDUTYS	—	29.2	1.00	6 449	349.2
2-7	DECANT2	—	107.5	15.0	2 172	112.9	4-5	HX48	—	40.0	0.30	8 650	461.6
2-8	COL1	M22	73.4	1.00	40 660	583.6	4-6	COL2	M41	93.8	0.30	53 768	1 399
2-9	R2	HX22	127.6	15.0	75 640	1 793	4-7	M41	HX41	23.7	25.0	617 870	16 081
3-0	COL1	HX31	154.4	1.00	8 905	90.4	4-8	HX41	CRY51	15.0	25.0	617 870	16 081
3-1	HX31	R4	30.0	1.00	8 905	90.4	4-9	CRY51	S41	17.0	1.00	617 900	16 082
3-2	R4	ADSB1	70.0	3.00	67 336	2 014	5-0	M31	CMPR3	65.6	1.00	14 681	457.8
3-3	CMPR4	ADSB1	176.5	3.00	7 526	259.9	5-1	CMPR3	HX51	575.8	30.0	14 681	457.8
3-4	R4	M31	70.0	3.00	2 443	57.4	5-10	HX54	R5	570.0	30.0	10 585	333.2
3-5	ADSB1	M31	65.1	1.00	12 238	400.4	5-11	R5	HX54	800.1	30.0	10 585	350.4
3-6	M32	HX32	46.3	1.00	45 320	1 428	5-12	HX54	HX55	796.0	30.0	10 585	350.4
3-8	HX32	—	25.0	1.00	45 320	1 428	5-13	HX55	R6	280.0	30.0	10 585	350.4
4-0	ADSB1	EOSCOR1	52.0	1.00	62 624	1 868	5-14	R6	TUR1	620.6	30.0	10 585	363.3
4-1	EOSCOR1	V1	52.0	1.00	62 624	1 868	5-15	TUR1	—	195.3	1.00	10 585	363.3
4-10	S41	M41	17.0	1.00	564 110	14 682	5-2	HX51	HX52	216.7	30.0	14 681	457.8
4-11	S41	CEN1	17.0	1.00	53 800	1 400	5-3	HX52	ADSB2	35.0	30.0	14 681	457.8
4-12	CEN1	V2	19.0	1.00	41 726	1 318	5-4	ADSB2	M32	60.3	30.0	4 387	125.5
4-13	V2	HX42	19.0	0.30	41 726	1 318	5-5	ADSB2	HX53	57.3	30.0	10 625	335.4
4-14	CEN1	M43	19.0	1.00	12 074	82.6	5-6	HX53	F1	20.0	30.0	10 625	335.4
4-15	HX42	M45	87.1	0.30	37 078	1 186	5-7	F1	ADSB2	20.0	30.0	39.9	2.12
4-16	HX42	M42	87.1	0.30	4 648	131.8	5-8	F1	HX51	20.0	30.0	10 585	333.2
4-17	M42	HX43	44.9	1.00	13 279	376.5	5-9	HX51	HX54	565.8	30.0	10 585	333.2
4-18	HX43	CRY52	15.0	1.00	13 279	376.5	AIR-IN	—	CMPR1	25.0	1.00	28 959	1 000
4-19	CRY52	S42	17.0	1.00	13 279	376.5	AIR-IN-2	—	CMPR4	25.0	1.00	7 526	259.9
4-2	V1	COL2	51.7	0.30	62 624	1 868	BEN-IN	—	R1	25.0	50.0	8 797	112.6
4-20	S42	M42	17.0	1.00	8 632	244.7	H2-IN	—	M11	25.0	50.0	691.1	341.2
4-21	S42	CEN2	17.0	1.00	4 648	131.8	H2O-IN	—	ADSB2	25.0	30.0	291.3	16.2
4-22	CEN2	V3	18.9	1.00	4 319	129.5	H2O-IN2	—	M43	25.0	1.00	6 264	347.7

Stream	From	To	T (°C)	P (bar)	Mass flow (kg/h)	Molar flow (kmol/h)	Stream	From	To	T (°C)	P (bar)	Mass flow (kg/h)	Molar flow (kmol/h)
4-23	V3	HX44	18.9	0.30	4 319	129.5	HNO3-IN	—	R4	25.0	1.00	60 874	1 931
4-24	HX44	M45	87.2	0.30	3 855	116.6							

Table G.1: Conventional route: Aspen Plus stream summary, all 103 streams, read left to right. Blocks are named as on the flowsheet. A dash marks a stream entering or leaving the battery limits; *n/a* marks a stream that carries no flow in the converged run, for which Aspen reports no state.

Stream		Mass flow (kg/h)	Molar flow (kmol/h)	Composition (kmol/h)
1-6	<i>product</i>	8.56	3.35	H2 3.322, ANE 0.017, AR 0.010
2-13	<i>product</i>	24 678	858.6	N2 780.4, O2 46.8, CO2 20.2, AR 9.364, H2O 1.486, ANE 0.350
2-3	<i>product</i>	2 683	20.3	GA 20.3
2-7	<i>product</i>	2 172	112.9	H2O 110.2, CHN 1.452, N2 0.785, O2 0.216, CO2 0.112, CHA 0.085, AR 0.012, ANE 0.006
3-8	<i>product</i>	45 320	1 428	H2O 991.9, HNO3 431.2, N2 1.188, N2O 0.959, O2 0.808, CHN 0.741, NO2 0.352, CHA 0.348, AA 0.272, ANE 0.131, MTHYLANE 0.050, AR 0.033, NO 0.013
4-25	<i>product</i>	463.8	13.0	H2O 7.077, HNO3 3.285, NO3- 1.179, H3O+ 1.179, AA 0.230, CHA 0.002, CHN 0.002
4-30	<i>product</i>	205.4	7.39	N2 2.839, O2 2.610, H2O 1.810, AR 0.050, CHN 0.044, CHA 0.033, HNO3 0.002
4-44	<i>product</i>	12 246	83.8	AA(S) 83.8
4-45	<i>product</i>	6 449	349.2	H2O 347.9, AA 1.243
4-5	<i>product</i>	8 650	461.6	H2O 456.6, CHN 2.695, HNO3 1.404, CHA 0.628, O2 0.193, N2 0.068, AR 0.002
5-15	<i>product</i>	10 585	363.3	N2 269.3, O2 91.1, AR 2.332, H2O 0.327, ANE 0.282, MTHYLANE 0.050, NO2 0.002
AIR-IN	<i>feed</i>	28 959	1 000	N2 781.2, O2 209.5, AR 9.300
AIR-IN-2	<i>feed</i>	7 526	259.9	N2 203.0, O2 54.4, AR 2.417
BEN-IN	<i>feed</i>	8 797	112.6	BEN 112.5, TOLUENE 0.100
H2-IN	<i>feed</i>	691.1	341.2	H2 341.1, AR 0.087
H2O-IN	<i>feed</i>	291.3	16.2	H2O 16.2
H2O-IN2	<i>feed</i>	6 264	347.7	H2O 347.7
HNO3-IN	<i>feed</i>	60 874	1 931	H2O 1 352, HNO3 579.6

Table G.2: Conventional route: composition of the 18 streams crossing the battery limits. Components below 0.001 kmol/h are omitted.

G.2. Electrochemical route

Stream	From	To	T (°C)	P (bar)	Mass flow (kg/h)	Molar flow (kmol/h)	Stream	From	To	T (°C)	P (bar)	Mass flow (kg/h)	Molar flow (kmol/h)
1-1	M11	R1	122.4	50.0	1 540	672.8	4-14	S41	M41	17.4	1.00	365 740	12 676
1-2	R1	HX11	718.0	50.0	10 337	447.6	4-15	S41	CEN1	17.4	1.00	32 662	1 132
1-3	HX11	HX12	185.0	50.0	10 337	447.6	4-16	CEN1	V3	18.6	1.00	20 050	1 041
1-4	HX12	S11	35.0	50.0	855.5	334.9	4-17	CEN1	M44	18.6	1.00	12 612	90.8
1-5	S11	M11	35.0	50.0	855.5	334.9	4-18	V3	HX44	18.6	0.39	20 050	1 041
1-6	S11	—	n/a	n/a	0.00	0.00	4-2	COL2	HX41	99.6	1.00	73 900	4 063
2-0	HX12	R2	35.0	50.0	9 481	112.7	4-20	HX44	M42	77.0	0.39	6 920	312.4
2-1	R2	SP1	127.0	15.0	55 634	773.5	4-21	HX44	M43	86.0	0.60	74 270	4 123
2-10	HX22	SP2	20.0	15.0	25 190	864.4	4-22	HX44	M43	77.0	0.39	13 130	728.8
2-11	HX22	M21	20.0	15.0	50 191	933.8	4-23	M42	HX46	57.8	0.39	10 376	468.3
2-12	SP2	M22	20.0	15.0	521.6	6.20	4-24	HX46	CRYS2	15.0	0.39	10 376	468.3
2-13	SP2	—	20.0	15.0	24 668	858.2	4-25	CRYS2	S42	17.0	1.00	10 376	468.3
2-14	M22	HX23	71.8	1.00	42 382	626.9	4-26	S42	M42	17.0	1.00	3 455	155.9
2-15	HX23	M21	35.0	1.00	42 382	626.9	4-27	S42	CEN2	17.0	1.00	6 920	312.3

Stream	From	To	T (°C)	P (bar)	Mass flow (kg/h)	Molar flow (kmol/h)	Stream	From	To	T (°C)	P (bar)	Mass flow (kg/h)	Molar flow (kmol/h)
2-16	M21	R2	18.6	1.00	92 574	1 561	4-28	CEN2	CNDUTYS	18.0	1.00	6 541	309.8
2-17	CMPR1	HX24	224.3	4.00	28 959	1 000	4-29	CEN2	M44	18.0	1.00	379.5	2.60
2-18	HX24	CMPR2	35.0	4.00	28 959	1 000	4-3	V1	HX41	93.8	0.80	180 390	9 332
2-19	CMPR2	R2	229.3	15.0	28 959	1 000	4-30	M43	HX45	75.6	0.39	87 400	4 851
2-2	SP1	HX21	127.0	15.0	52 949	753.2	4-31	HX45	M46	75.6	0.39	87 400	4 851
2-20	COL1	M23	154.1	1.00	8 916	90.5	4-32	—	M31	92.8	1.00	71 901	3 961
2-3	SP1	—	127.0	15.0	2 684	20.3	4-33	HX47	M44	91.8	2.00	6 059	336.3
2-4	HX21	R3	40.0	15.0	52 949	753.2	4-34	M44	HX48	95.0	2.00	167 680	3 783
2-5	R3	DECANT2	100.0	15.0	52 949	824.3	4-35	HX48	CRYS4	105.0	2.00	167 680	3 783
2-6	DECANT2	COL1	106.4	15.0	50 777	711.2	4-36	CRYS4	S43	99.0	2.00	167 690	3 783
2-7	DECANT2	M23	106.4	15.0	2 173	113.2	4-37	S43	M44	99.0	2.00	148 630	3 353
2-8	COL1	M22	72.0	1.00	41 861	620.7	4-38	S43	M45	99.0	2.00	19 062	430.0
2-9	R2	HX22	127.0	15.0	75 381	1 798	4-39	M45	HX49	25.5	1.00	251 960	5 682
3-0	M23	HX31	95.8	1.00	11 089	203.7	4-4	HX41	V2	94.1	0.80	106 930	5 255
3-1	HX31	M31	25.0	1.00	11 089	203.7	4-40	HX49	CRYS3	13.0	1.00	251 960	5 682
3-10	R4	SP3	113.1	50.0	1 040	337.9	4-41	CRYS3	S45	17.0	1.00	251 960	5 682
3-11	SP3	—	113.1	50.0	232.1	7.26	4-42	S45	M45	17.0	1.00	232 890	5 252
3-12	SP3	HX33	113.1	50.0	807.9	330.6	4-43	S45	CEN3	17.0	1.00	19 069	430.0
3-13	HX33	—	45.0	50.0	661.0	322.5	4-44	CEN3	—	18.5	1.00	12 732	87.1
3-14	HX33	M34	45.0	50.0	146.9	8.15	4-45	CEN3	M46	18.5	1.00	6 337	342.9
3-16	UEDS1B	HX35	100.4	1.00	184 550	9 744	4-46	—	M33	81.5	1.00	165 970	9 204
3-17	HX35	M32	25.0	1.00	184 550	9 744	4-47	CNDUTYS	—	27.0	1.00	6 541	309.8
3-18	M33	UEDS1B	80.2	1.00	184 550	9 744	4-48	HX47	—	95.5	1.00	73 900	4 063
3-2	M31	M32	85.2	1.00	82 990	4 165	4-49	M46	—	81.4	1.00	167 200	9 272
3-3	M32	HX32	44.1	1.00	268 220	13 920	4-5	HX41	HX42	94.1	0.80	73 458	4 077
3-4	HX32	R4	25.0	1.00	268 220	13 920	4-6	HX41	HX47	96.8	1.00	73 900	4 063
3-5	R4	EDS1	113.1	50.0	267 180	13 619	4-7	V2	HX42	86.5	0.60	106 930	5 255
3-6	EDS1	M33	113.1	50.0	13 033	232.3	4-8	HX42	M41	88.2	0.60	32 663	1 132
3-7	EDS1	UEDS1A	113.1	50.0	254 140	13 387	4-9	HX42	M46	93.5	0.80	73 458	4 077
3-8	UEDS1A	M34	163.9	20.0	254 140	13 387	AIR-IN	—	CMPR1	25.0	1.00	28 959	1 000
4-0	M34	COL2	163.9	20.0	254 290	13 395	BEN-IN	—	R1	210.0	50.0	8 797	112.6
4-1	COL2	V1	100.0	1.00	180 390	9 332	H2-IN	—	M11	210.0	50.0	684.4	337.9
4-10	HX42	HX44	88.2	0.60	74 270	4 123	H2OIN1	—	M33	25.0	50.0	5 539	307.4
4-11	M41	HX43	24.4	0.60	398 400	13 808	H2OIN2	—	HX47	25.0	2.00	6 059	336.3
4-12	HX43	CRYS1	15.0	0.60	398 400	13 808	KOHIN	—	M32	25.0	50.0	682.6	12.2
4-13	CRYS1	S41	17.0	1.00	398 400	13 808							

Table G.3: Electrochemical route: Aspen Plus stream summary, all 99 streams, read left to right. Blocks are named as on the flowsheet. A dash marks a stream entering or leaving the battery limits; *n/a* marks a stream that carries no flow in the converged run, for which Aspen reports no state.

Stream	Mass flow (kg/h)	Molar flow (kmol/h)	Composition (kmol/h)
1-6	<i>product</i>	0.00	0.00 none above 0.001
2-13	<i>product</i>	24 668	858.2 N2 780.2, O2 46.7, CO2 20.2, AR 9.373, H2O 1.399, ANE 0.326
2-3	<i>product</i>	2 684	20.3 GA 20.3
3-11	<i>product</i>	232.1	7.26 H2O 2.392, CO2 2.313, O2 1.907, N2 0.535, MTHYLANE 0.096, CHN 0.016, AR 0.006
3-13	<i>product</i>	661.0	322.5 H2 321.8, H2O 0.683
4-32	<i>feed</i>	71 901	3 961 H2O 3 953, CHN 6.381, CO2 0.863, O2 0.767, N2 0.036
4-44	<i>product</i>	12 732	87.1 AA(S) 87.1
4-46	<i>feed</i>	165 970	9 204 H2O 9 203, AA 1.196, KOH 0.061, CHN 0.032, GA 0.017
4-47	<i>product</i>	6 541	309.8 H2O 293.4, KOH 12.2, GA 3.183, AA 1.045
4-48	<i>product</i>	73 900	4 063 H2O 4 047, CHN 6.553, O2 4.428, CO2 4.359, N2 0.479, H2 0.029, AR 0.008, MTHYLANE 0.004

Stream		Mass flow	Molar flow	Composition (kmol/h)
4-49	product	167 200	9 272	H2O 9 270, AA 1.220, KOH 0.061, CHN 0.033, GA 0.017
AIR-IN	feed	28 959	1 000	N2 781.2, O2 209.5, AR 9.300
BEN-IN	feed	8 797	112.6	BEN 112.5, TOLUENE 0.100
H2-IN	feed	684.4	337.9	H2 337.8, AR 0.086
H2OIN1	feed	5 539	307.4	H2O 307.4
H2OIN2	feed	6 059	336.3	H2O 336.3
KOHIN	feed	682.6	12.2	KOH 12.2

Table G.4: Electrochemical route: composition of the 17 streams crossing the battery limits. Components below 0.001 kmol/h are omitted.

G.3. Biomass route to t3HDA

Stream	From	To	T (°C)	P (bar)	Mass flow (kg/h)	Molar flow (kmol/h)	Stream	From	To	T (°C)	P (bar)	Mass flow (kg/h)	Molar flow (kmol/h)
1-1	HX11	HX12	81.6	1.00	280 000	14 746	2-18	CEN1	M23	2.9	1.00	13 632	123.6
1-10	CB1	R1	25.0	1.00	604 620	31 915	2-19	M23	HX22	22.0	2.00	48 033	2 033
1-11	SP1	—	80.0	1.00	5 698	227.2	2-2	CORRECT1	CORRECT3	29.8	1.00	295 850	15 728
1-12	SP1	HX14	80.0	1.00	2 311	248.8	2-20	HX22	M24	92.0	2.00	48 033	2 033
1-13	HX14	CMPR1	25.0	1.00	361.6	140.6	2-21	M24	HX23	95.9	2.00	240 170	10 169
1-14	CMPR1	HX15	204.2	3.50	361.6	140.6	2-22	HX23	CRY3	100.0	2.00	240 170	10 170
1-15	HX15	CMPR2	25.0	3.50	302.9	137.4	2-23	CRY3	BCRY3	97.0	2.00	240 170	10 170
1-16	CMPR2	HX16	171.8	10.0	302.9	137.4	2-24	BCRY3	S22	97.0	2.00	240 170	10 170
1-17	HX16	CMPR3	25.0	10.0	287.8	136.5	2-25	S22	CB2	97.0	2.00	192 130	8 136
1-18	HX14	M11	25.0	1.00	1 949	108.2	2-26	CB2	M24	97.0	2.00	192 130	8 136
1-19	HX15	M11	25.0	3.50	58.7	3.26	2-27	S22	HX22	97.0	2.00	48 033	2 034
1-2	HX12	F1	95.0	1.00	280 000	14 749	2-28	HX22	M25	29.6	2.00	48 033	2 034
1-20	HX16	M11	25.0	10.0	15.1	0.84	2-29	M25	HX24	18.8	2.00	332 170	14 066
1-21	M11	—	25.0	1.00	2 023	112.3	2-3	CORRECT3	HX21	25.0	1.00	295 850	15 728
1-22	CMPR3	—	180.2	30.0	287.8	136.5	2-30	HX24	CRY2	15.0	2.00	332 170	14 066
1-3	F1	—	95.0	1.00	17 547	865.9	2-31	CRY2	S23	17.0	2.00	332 170	14 066
1-4	F1	HX11	95.0	1.00	262 450	13 884	2-32	S23	CB3	17.0	2.00	284 140	12 032
1-5	HX11	R1	30.0	1.00	262 450	13 884	2-33	CB3	M25	17.0	2.00	284 140	12 032
1-6	R1	SP1	80.0	1.00	8 009	476.1	2-34	S23	CEN2	17.0	2.00	48 033	2 034
1-7	R1	HX13	80.0	1.00	859 060	45 346	2-35	CEN2	—	18.1	1.00	12 334	84.4
1-8	HX13	S11	25.0	1.00	859 060	45 346	2-36	CEN2	—	18.1	1.00	35 699	1 950
1-9	S11	CB1	25.0	1.00	604 620	31 915	2-4	HX21	M21	15.0	1.00	295 850	15 728
2-0	S11	—	25.0	1.00	254 440	13 431	2-5	M21	COOLER1	19.9	1.00	309 280	15 957
2-1	—	CORRECT1	25.0	1.00	254 630	13 433	2-6	COOLER1	CRY1	2.0	1.00	309 280	15 962
2-10	F2	—	n/a	1.00	0.00	0.00	2-7	CRY1	S21	2.0	1.00	309 280	15 962
2-11	F2	CEN1	2.0	1.00	309 280	15 962	2-8	S21	F2	2.0	1.00	309 280	15 962
2-12	CEN1	M22	2.9	1.00	295 650	15 838	2-9	S21	M21	n/a	n/a	0.00	0.00
2-13	M22	CNDUTYS	3.0	1.00	297 640	15 864	BROTH-IN	—	HX11	25.0	1.00	280 000	14 742
2-14	CNDUTYS	CLR-ELEC	4.1	1.00	297 640	15 864	H2OIN	—	M23	25.0	2.00	34 401	1 910
2-15	CLR-ELEC	CORRECT2	1.5	1.00	297 640	15 865	H2SO4IN	—	M21	25.0	1.00	13 436	331.1
2-16	CORRECT2	CLR-WTR	80.0	1.00	297 640	15 849	LIMEIN	—	M22	25.0	1.00	1 989	26.8
2-17	CLR-WTR	—	59.3	1.00	297 640	15 851	NH3IN	—	CORRECT1	25.0	1.00	41 216	2 297

Table G.5: Biomass route to t3HDA: Aspen Plus stream summary, all 64 streams, read left to right. Blocks are named as on the flowsheet. A dash marks a stream entering or leaving the battery limits; *n/a* marks a stream that carries no flow in the converged run, for which Aspen reports no state.

Stream		Mass flow (kg/h)	Molar flow (kmol/h)	Composition (kmol/h)
1-11	product	5 698	227.2	O2 113.3, H2O 112.3, CO2 0.831, NH3 0.797, NH4+ 0.001
1-21	product	2 023	112.3	H2O 111.9, NH3 0.416, OH- 0.004, NH4+ 0.004, H2 0.002
1-22	product	287.8	136.5	H2 135.7, H2O 0.433, NH3 0.378
1-3	product	17 547	865.9	H2O 727.5, O2 125.5, CO2 7.610, NH3 5.145, N2 0.047
2-0	product	254 440	13 431	H2O 13 318, T3HDA 90.6, NH3 11.1, GLUCOSE 4.740, NH4+ 2.813, MA 1.674, HCO3- 1.193, CO3- 0.710, H2 0.248, NH2COO- 0.174, O2 0.080, OH- 0.026
2-1	feed	254 630	13 433	H2O 13 319, T3A(S) 78.7, NH4+ 14.1, T3A- 11.5, GLUCOSE 4.740, CO2 2.071, MAA- 1.609, T3A-2 0.462, H2 0.248, O2 0.080, MAA-2 0.065, H+ 0.056, HCO3- 0.006
2-10	product	0.00	0.00	none above 0.001
2-17	product	297 640	15 851	H2O 15 500, NH4+ 186.8, SO4- 73.9, HSO4- 46.3, LIME 26.8, H+ 7.431, GLUCOSE 4.730, T3A(AQ) 3.425, CO2 2.072, H2 0.248, O2 0.080, MAA- 0.020, T3A- 0.020
2-35	product	12 334	84.4	T3A(S) 84.4
2-36	product	35 699	1 950	H2O 1 944, T3A(AQ) 2.321, MAA(AQ) 1.181, H+ 0.829, MAA- 0.472, T3A- 0.472, NH4+ 0.410, HSO4- 0.233, SO4- 0.030, GLUCOSE 0.010, CO2 0.005
BROTH-IN	feed	280 000	14 742	H2O 14 489, O2 125.6, MA 92.3, NH4+ 11.7, HCO3- 8.370, NH3 8.336, GLUCOSE 4.740, CO3- 1.250, NH2COO- 0.824, CO2 0.076, N2 0.047, OH- 0.006
H2OIN	feed	34 401	1 910	H2O 1 910
H2SO4IN	feed	13 436	331.1	H+ 120.7, HSO4- 120.2, H2O 89.9, SO4- 0.260, H2SO4 0.001
LIMEIN	feed	1 989	26.8	LIME 26.8
NH3IN	feed	41 216	2 297	H2O 2 124, NH3 172.9, OH- 0.202, NH4+ 0.202

Table G.6: Biomass route to r3HDA: composition of the 15 streams crossing the battery limits. Components below 0.001 kmol/h are omitted.

G.4. Biomass route to adipic acid

Stream	From	To	T (°C)	P (bar)	Mass flow (kg/h)	Molar flow (kmol/h)	Stream	From	To	T (°C)	P (bar)	Mass flow (kg/h)	Molar flow (kmol/h)
1-1	HX11	HX12	81.6	1.00	280 000	14 746	2-25	S22	CB2	97.0	2.00	192 130	8 136
1-10	CB1	R1	25.0	1.00	622 530	32 862	2-26	CB2	M24	97.0	2.00	192 130	8 136
1-11	SP1	—	79.1	1.00	5 562	219.8	2-27	S22	HX22	97.0	2.00	48 031	2 034
1-12	SP1	H2-SPLIT	79.1	1.00	2 177	241.4	2-28	HX22	M25	29.6	2.00	48 031	2 034
1-13	HX14	CMPR1	25.0	1.00	121.2	47.1	2-29	M25	HX24	18.8	2.00	332 160	14 065
1-14	CMPR1	HX15	204.2	3.50	121.2	47.1	2-3	CORRECT3	HX21	25.0	1.00	296 100	15 742
1-15	HX15	CMPR2	25.0	3.50	101.5	46.0	2-30	HX24	CRY2	15.0	2.00	332 160	14 065
1-16	CMPR2	HX16	171.8	10.0	101.5	46.0	2-31	CRY2	S23	17.0	2.00	332 160	14 065
1-17	HX16	CMPR3	25.0	10.0	96.4	45.8	2-32	S23	CB3	17.0	2.00	284 130	12 031
1-18	HX14	M11	25.0	1.00	608.9	33.8	2-33	CB3	M25	17.0	2.00	284 130	12 031
1-19	HX15	M11	25.0	3.50	19.7	1.09	2-34	S23	CEN2	17.0	2.00	48 031	2 034
1-2	HX12	F1	95.0	1.00	280 000	14 749	2-35	CEN2	M-REDIS	18.1	1.00	12 334	84.4
1-20	HX16	M11	25.0	10.0	5.04	0.28	2-36	CEN2	—	18.1	1.00	35 698	1 949
1-21	M11	—	25.0	1.00	633.6	35.2	2-4	HX21	M21	15.0	1.00	296 100	15 742
1-22	CMPR3	—	180.2	30.0	96.4	45.8	2-5	M21	COOLER1	19.9	1.00	309 540	15 971
1-3	F1	—	95.0	1.00	17 547	865.9	2-6	COOLER1	CRY1	2.0	1.00	309 540	15 976
1-4	F1	HX11	95.0	1.00	262 450	13 884	2-7	CRY1	S21	2.0	1.00	309 540	15 976
1-5	HX11	R1	30.0	1.00	262 450	13 884	2-8	S21	F2	2.0	1.00	309 540	15 976
1-6	R1	SP1	79.1	1.00	7 740	461.2	2-9	S21	M21	n/a	n/a	0.00	0.00
1-7	R1	HX13	79.1	1.00	877 240	46 308	AA-COLD	COOL-AA	AA-CRY5	17.0	2.00	49 061	2 123
1-8	HX13	S11	25.0	1.00	877 240	46 308	AA-COLD2	COOLER2	AA-CRY52	2.0	1.00	49 061	2 123
1-9	S11	CB1	25.0	1.00	622 530	32 862	AA-LIQ	FAKE-H2	COOL-AA	70.0	2.00	49 061	2 123
2-0	S11	—	25.0	1.00	254 710	13 446	AA-PROD	R-POLISH	FAKE-H2	70.0	2.00	49 531	2 229

Stream	From	To	T (°C)	P (bar)	Mass flow (kg/h)	Molar flow (kmol/h)	Stream	From	To	T (°C)	P (bar)	Mass flow (kg/h)	Molar flow (kmol/h)
2-1	—	CORRECT1	25.0	1.00	254 900	13 448	AA-SOLID	AA-CEN	—	2.0	1.00	11 757	80.4
2-10	F2	—	n/a	1.00	0.00	0.00	BROTH-IN	—	HX11	25.0	1.00	280 000	14 742
2-11	F2	CEN1	2.0	1.00	309 540	15 976	CRY-OUT2	AA-CRYS2	AA-CEN	2.0	1.00	49 061	2 123
2-12	CEN1	M22	2.9	1.00	295 910	15 852	CRY-OUT	AA-CRYS	COOLER2	17.0	1.00	49 061	2 123
2-13	M22	CNDUTYS	3.0	1.00	297 900	15 878	DI-WATER	—	M-REDIS	25.0	1.00	35 750	1 984
2-14	CNDUTYS	CLR-ELEC	4.1	1.00	297 900	15 878	H2-COMP	CMPR4	R-POLISH	178.1	2.00	1 447	160.5
2-15	CLR-ELEC	CORRECT2	1.2	1.00	297 900	15 879	H2-FEED	H2-SPLIT	CMPR4	79.1	1.00	1 447	160.5
2-16	CORRECT2	CLR-WTR	80.0	1.00	297 900	15 863	H2-RXN	FAKE-H2	—	70.0	2.00	469.7	106.1
2-17	CLR-WTR	—	56.8	1.00	297 900	15 866	H2-SALE	H2-SPLIT	HX14	79.1	1.00	730.1	80.9
2-18	CEN1	M23	2.9	1.00	13 632	123.6	H2OIN	—	M23	25.0	2.00	34 400	1 909
2-19	M23	HX22	22.0	2.00	48 031	2 033	H2SO4IN	—	M21	25.0	1.00	13 440	331.2
2-2	CORRECT1	CORRECT3	29.8	1.00	296 100	15 742	LIMEIN	—	M22	25.0	1.00	1 990	26.9
2-20	HX22	M24	92.0	2.00	48 031	2 033	ML-AA	AA-CEN	—	2.0	1.00	37 304	2 043
2-21	M24	HX23	95.9	2.00	240 160	10 169	NH3IN	—	CORRECT1	25.0	1.00	41 202	2 297
2-22	HX23	CRYS3	100.0	2.00	240 160	10 169	T3HDA-AQ	M-REDIS	HX-POLIS	24.2	1.00	48 084	2 069
2-23	CRYS3	BCRYS3	97.0	2.00	240 160	10 169	T3HDA-HO	HX-POLIS	R-POLISH	70.0	1.00	48 084	2 069
2-24	BCRYS3	S22	97.0	2.00	240 160	10 169							

Table G.7: Biomass route to adipic acid: Aspen Plus stream summary, all 79 streams, read left to right. Blocks are named as on the flowsheet. A dash marks a stream entering or leaving the battery limits; *n/a* marks a stream that carries no flow in the converged run, for which Aspen reports no state.

Stream		Mass flow (kg/h)	Molar flow (kmol/h)	Composition (kmol/h)
1-11	product	5 562	219.8	O2 113.3, H2O 105.0, CO2 0.783, NH3 0.748, NH4+ 0.001
1-21	product	633.6	35.2	H2O 35.0, NH3 0.127, OH- 0.001, NH4+ 0.001
1-22	product	96.4	45.8	H2 45.5, H2O 0.145, NH3 0.123
1-3	product	17 547	865.9	H2O 727.5, O2 125.5, CO2 7.610, NH3 5.145, N2 0.047
2-0	product	254 710	13 446	H2O 13 332, T3HDA 90.6, NH3 11.2, GLUCOSE 4.740, NH4+ 2.870, MA 1.674, HCO3- 1.225, CO3- 0.720, H2 0.260, NH2COO- 0.179, O2 0.084, OH- 0.026
2-1	feed	254 900	13 448	H2O 13 334, T3A(S) 78.6, NH4+ 14.2, T3A- 11.6, GLUCOSE 4.740, CO2 2.118, MAA- 1.609, T3A-2 0.468, H2 0.260, O2 0.084, MAA-2 0.065, H+ 0.055, HCO3- 0.006
2-10	product	0.00	0.00	none above 0.001
2-17	product	297 900	15 866	H2O 15 514, NH4+ 186.8, SO4- 74.4, HSO4- 45.9, LIME 26.9, H+ 7.844, GLUCOSE 4.730, T3A(AQ) 3.429, CO2 2.120, H2 0.260, O2 0.083, MAA- 0.020, T3A- 0.020
2-36	product	35 698	1 949	H2O 1 944, T3A(AQ) 2.321, MAA(AQ) 1.181, H+ 0.829, MAA- 0.472, T3A- 0.472, NH4+ 0.410, HSO4- 0.233, SO4- 0.030, GLUCOSE 0.010, CO2 0.005
AA-SOLID	product	11 757	80.4	AA(S) 80.4
BROTH-IN	feed	280 000	14 742	H2O 14 489, O2 125.6, MA 92.3, NH4+ 11.7, HCO3- 8.370, NH3 8.336, GLUCOSE 4.740, CO3- 1.250, NH2COO- 0.824, CO2 0.076, N2 0.047, OH- 0.006
DI-WATER	feed	35 750	1 984	H2O 1 984
H2-RXN	product	469.7	106.1	H2 90.1, H2O 16.0
H2OIN	feed	34 400	1 909	H2O 1 909
H2SO4IN	feed	13 440	331.2	H+ 120.8, HSO4- 120.2, H2O 90.0, SO4- 0.261, H2SO4 0.001
LIMEIN	feed	1 990	26.9	LIME 26.9
ML-AA	product	37 304	2 043	H2O 2 038, AA 3.636, NH4+ 0.498, T3A-2 0.188, T3A- 0.122, H2 0.050
NH3IN	feed	41 202	2 297	H2O 2 123, NH3 172.8, OH- 0.202, NH4+ 0.202

Table G.8: Biomass route to adipic acid: composition of the 18 streams crossing the battery limits. Components below 0.001 kmol/h are omitted.

Acknowledgement of generative AI use

I acknowledge the use of Claude (Anthropic, <https://claude.ai>) to support the development of this thesis by refining language, assisting with programming and debugging, and providing feedback. It was used in the following ways:

- **Refining language.** Drafting and rewriting passages of the results, conclusions and appendices from my own notes, arguments and dictated text, and editing for concision and consistency. All claims, their framing and their interpretation are my own, and every draft was reviewed and revised by me.
- **Programming and debugging.** Assisting with the Python used for the techno-economic calculations, the scenario sweeps and the figures, and with checking intermediate values against the process model.
- **Providing feedback.** Checking internal consistency between the reported numbers, the tables and the figures.

The process models, the cost and emission calculations, and the results and conclusions drawn from them are my own work. No text was included without being read, checked and edited by me, and no source was cited that I had not consulted myself.