

Understanding the Emergence of CO₂ Value Chains: An Agent-Based Evolutionary Game Approach Coupled with CO₂ Network Optimisation

Florian Verwilghen

Thesis submitted for the degree of
Master of Science in Engineering:
Energy

Supervisors

Prof. Dr., Ir. Erik Delarue
Dr., Ir. Jorge Andrés Moncada

Assessors

Prof. Dr., Ir. Geert Deconinck
Dr., Ir. Ruben Baetens

Assistant-supervisors

Ir. Flore Verbist
Dr., rer., pol. Veis Karbassi

© 2026 KU Leuven – Faculty of Engineering Science

Published by Florian Verwilghen,

Faculty of Engineering Science, Kasteelpark Arenberg 1 bus 2200, B-3001 Leuven

All rights reserved. No part of the publication may be reproduced in any form by print, photoprint, microfilm, electronic or any other means without written permission from the publisher. This publication contains the study work of a student in the context of the academic training and assessment. After this assessment no correction of the study work took place.

Preface

This thesis marks the culmination of a significant chapter in my academic journey. Its completion would not have been possible without the support, guidance, and encouragement of many people to whom I am deeply grateful.

First and foremost, I would like to express my heartfelt gratitude to my supervisor, Prof. Erik Delarue, for his valuable guidance, insightful advice, and constructive feedback throughout this project.

My sincere thanks go to Dr. Jorge A. Moncada, Flore Verbist, and Dr. Veis Karbassi for their dedicated support throughout this academic year. Their regular meetings, thoughtful feedback, and readiness to answer my questions greatly contributed to the development of this thesis. I am especially grateful to Flore Verbist for graciously allowing me to use a model she developed as part of this work.

I am also thankful to Prof. Geert Deconinck and Ruben Baetens for the time and effort they dedicated to reading and evaluating this thesis.

I owe a deep debt of gratitude to my family, and especially my parents, for their constant encouragement and unwavering belief in me throughout my studies. I am equally grateful to my brothers for their support and encouragement.

Finally, I wish to extend a special thank you to Oriane Valentin for her presence, her wholehearted support, and all she has given me over the past three years. Her help and encouragement have meant the world to me during this period.

Florian Verwilghen

Contents

Preface	i
Abstract	v
List of Figures and Tables	vi
List of Abbreviations and Symbols	viii
1 Introduction	1
1.1 Context and research problem	1
1.2 Research question, scope and contribution	2
2 Literature Review	4
2.1 CCS and CCU as industrial carbon management value chains	4
2.2 Barriers and feasibility conditions for CCS and CCU deployment	5
2.2.1 Plant-level techno-economic feasibility	5
2.2.2 Value-chain and coordination constraints	6
2.2.3 Institutional, financial and regulatory conditions	6
2.3 Infrastructure dependence, network effects and path dependence	7
2.3.1 Shared transport and storage networks	7
2.3.2 Economies of scale and first-mover disadvantage	8
2.3.3 Network effects, critical mass and path dependence	9
2.4 Policy design, carbon accounting and CCS-CCU competition	9
2.4.1 ETS design and carbon price incentives	9
2.4.2 Carbon accounting, permanence and the CCS-CCU level-playing field	10
2.4.3 Subsidies, risk-sharing and early-stage support	11
2.4.4 Conditions favouring CCS, CCU or mixed outcomes	11
2.5 Modelling approaches and literature gap	12
2.5.1 Centralised optimisation and CO ₂ network models	12
2.5.2 Evolutionary game and agent-based adoption models	12
2.5.3 Positioning of this thesis	13
3 Theory and Methods	14
3.1 Methodological foundations: EGT and ABM	14
3.1.1 From traditional to evolutionary game theory	14
3.1.2 Why an agent-based implementation?	16
3.1.3 Linking adoption dynamics with network optimisation	17

3.1.4	Methodological implications for this thesis	18
3.2	Model description using the ODD protocol	18
3.2.1	Purpose and general model logic	19
3.2.2	Entities, state variables, scales and input datasets	19
3.2.3	Process overview and scheduling	24
3.2.4	Submodels	27
3.3	Model assumptions and limitations	37
4	Results	39
4.1	Base scenario results under the reference ETS price	39
4.1.1	Adoption pathways across the three dimensions of the base scenarios	40
4.1.2	Identifying evolutionary stable strategies	42
4.1.3	Critical transition dynamics: tipping points, lock-in and critical mass	44
4.1.4	Tariff allocation and CCS coalition formation	49
4.1.5	CCU as a conditional alternative capture route	50
4.1.6	Decentralised ABM-EGT outcomes versus the centralised CCNO reference	52
4.2	ETS price sensitivity and equilibrium landscape	54
4.2.1	Adoption dynamics from a zero-capture initial condition	54
4.2.2	Number and type of stationary outcomes	56
5	Discussion	58
5.1	From simulation outcomes to transition mechanisms	58
5.2	Positioning the results in the existing literature	59
5.3	Reliability of the findings and limits of interpretation	60
6	Conclusion	62
6.1	Research problem and modelling objective	62
6.2	Answering the research questions	62
6.3	Contributions and implications	64
6.4	Future research directions	65
A	ODD design concepts	68
A.1	Basic principles	68
A.2	Emergence	69
A.3	Adaptation	69
A.4	Objectives	70
A.5	Learning	70
A.6	Prediction	71
A.7	Sensing	71
A.8	Interaction	72
A.9	Stochasticity	72
A.10	Collectives	73
A.11	Observation	73
B	Input data and parameter values	75

C Construction of the CO₂ Demand Database	78
C.1 Definition of total sectoral demand	79
C.2 Regionalisation to the study area	80
C.3 Spatial allocation of demand nodes	80
C.3.1 Greenhouse horticulture	81
C.3.2 Food and beverage demand	81
C.3.3 Industrial CCU demand estimate	81
C.4 Final structure of the demand database	82
D Convergence test	84
E Learning-by-doing for capture CAPEX	87
F Calibration of the reduced-form CCU delivery-cost function	88
G Reading simplex plots	90
H Seed-dependency study	92
I CCU delivery-flow maps	95
J Additional CCS network maps	96
Bibliography	98
Code of Conduct	104

Abstract

Decarbonising heavy industry does not end with capture technology; in many respects, it begins there. Even when firms can technically capture CO₂, those emissions must still be transported, stored, or utilised under economically viable conditions. Because these downstream activities often rely on shared infrastructure, economies of scale, and accessible demand, capture deployment also depends on whether sufficient coordination and critical mass can be achieved. Heavy-industry decarbonisation is therefore not only a plant-level technology problem, but a systemic value-chain problem in which individual firm decisions interact with collective infrastructure and market conditions.

This thesis examines the techno-economic, spatial, and policy conditions necessary for the emergence of a CO₂ value chain. Building on this coordination problem, it investigates how decentralised emitter decisions interact with shared transport and storage infrastructure, policy incentives, and downstream CO₂ demand.

To address this question, an agent-based evolutionary game model is developed in which heterogeneous industrial emitters repeatedly choose among a no-capture production route, carbon capture and storage (CCS), and carbon capture and utilisation (CCU). CCS payoffs include a transport and storage (T&S) tariff derived from a centralised CO₂ network optimisation model, allowing emitter-facing T&S costs to evolve with the group of emitters adopting CCS. CCU is modelled through spatially distributed demand nodes, finite utilisation demand, and delivery costs. The simulations evaluate how carbon-accounting rules, T&S tariff design, CCU availability, and the Emissions Trading System (ETS) price shape the diffusion of industrial carbon-management strategies.

The results indicate that CCS diffusion is governed by a critical-mass mechanism. When too few emitters adopt CCS, shared T&S infrastructure remains costly and the system can remain locked in low-capture or non-capture outcomes. Once enough captured volume enters CCS, emitter-facing T&S tariffs can decrease, making further CCS adoption more attractive. CCU offers an alternative capture pathway, but only when finite utilisation demand and delivery costs are favourable. Higher ETS prices, carbon dioxide removal (CDR) credits, and location-sensitive T&S tariff allocation lower the threshold required for widespread CCS diffusion.

In summary, the findings demonstrate that a carbon-management system that appears efficient from a centralised perspective may not emerge organically from decentralised firm-level decision-making.

List of Figures and Tables

List of Figures

3.1	Spatial distribution of industrial emitters and CO ₂ demand points in the trilateral study region, covering Belgium, the Netherlands and North Rhine-Westphalia. Green points represent emitters and red points CO ₂ demand points.	20
3.2	Process flowchart of the agent-based evolutionary model. Green boxes indicate begin and final steps, purple boxes indicate entity creation and non-static properties/state updates, blue boxes indicate system updates and output recording, and the yellow diamond indicates the convergence check. Red outlines indicate where the CCNO model is called to initialise or update the CCS T&S network.	25
3.3	ABM–CCNO coupling inside the iterative loop.	27
4.1	Adoption-share trajectories across the eight base scenarios.	40
4.2	CO ₂ flow trajectories by destination and carbon origin across the eight base scenarios.	41
4.3	Basin-of-attraction maps for the CCS-only scenarios on a one-dimensional simplex; Appendix G explains how to read this type of plot. Grey arrows show initial-to-final CCS-share movements.	42
4.4	Basin-of-attraction maps for the CCU-enabled scenarios on a triangular simplex; Appendix G explains how to read this type of plot. Grey arrows show initial-to-final movements, while red lines highlight example trajectories.	43
4.5	Dynamic comparison of S1 and S5 over the number of iterations.	45
4.6	Cluster-level build-up of captured CCS volume in S1 & S5.	48
4.7	Spatial location of the main clusters involved in the critical transition, indicated on the final T&S networks of S1 on the left and S5 on the right.	49
4.8	Realised T&S networks in S1 (left) and S3 (right).	50
4.9	Conditional role of CCU across the four CCU-enabled scenarios.	51
4.10	Spatial comparison between the standalone CCNO CDR reference and the ABM-EGT–CCNO outcomes S1 and S3 at ETS 150 EUR/tCO ₂	53
4.11	The upper row reports adoption shares, the middle row reports CO ₂ flows by route, and the lower row reports the average T&S cost.	55

4.12	Simplex representations for scenarios S2 and S6 across ETS prices. Each row compares the same ETS price, while the left and right columns correspond respectively to the CDR and No-CDR settings. Grey arrows show the movement from initial to final strategy distributions. Red trajectories highlight representative paths from selected initial conditions.	57
E.1	Learning-by-doing factor applied to the capture-specific CAPEX.	87
G.1	Example of a one-dimensional simplex plot for a two-strategy case.	90
G.2	Example of a triangular simplex plot for a three-strategy case.	91
H.1	Adoption-share trajectories in the CCS-only scenarios with CDR credits.	92
H.2	Adoption-share trajectories in the CCS-only scenarios without CDR credits.	93
H.3	Seed-dependency of adoption-share trajectories in the CCU-enabled scenarios with CDR credits.	94
H.4	Seed-dependency of adoption-share trajectories in the CCU-enabled scenarios without CDR credits.	94
I.1	CCU delivery-flow maps	95
J.1	Supplementary CCS network maps.	96
J.2	Supplementary CCS network map for the decentralised ABM-EGT-CCNO outcome S5 under the No-CDR setting.	97

List of Tables

3.1	Industrial large-scale emitters – properties	21
3.2	CO ₂ demanders – properties	22
3.3	Passive Carbon Network Operator agent – properties	22
3.4	Global environment – properties	23
4.1	Overview of the eight base scenarios at the reference ETS price of 150 EUR/tCO ₂	39
4.2	Key diagnostics of the CCS critical-transition sequence in S1 and S5.	47
4.3	Comparison between the standalone CCNO model and selected ABM-EGT-CCNO outcomes at ETS 150 EUR/tCO ₂	52
B.1	Main input data and parameter values used in the reported ABM-EGT simulations	76
C.1	Central 2050 CO ₂ demand used in the ABM-EGT model.	79
C.2	Spatial structure of the central CO ₂ demand database.	82
D.1	Convergence parameters used in the main logit simulations	86

List of Abbreviations and Symbols

Abbreviations

ABM	Agent-Based Modelling
ABM-EGT	Agent-Based Modelling–Evolutionary Game Theory framework
ABM-EGT–CCNO	Agent-based evolutionary-game framework coupled with the Centralised CO ₂ Network Optimisation model
AIDRES	Advancing Industrial Decarbonisation by Assessing the Future Use of Renewable Energies in Industrial Processes
CAPEX	Capital Expenditure
CC	Carbon Capture
No CC	No Carbon Capture
CCS	Carbon Capture and Storage
CCU	Carbon Capture and Utilisation
CCUS	Carbon Capture, Utilisation and Storage
CCNO	Centralised CO ₂ Network Optimisation
CDR	Carbon Dioxide Removal
CNO	Carbon Network Operator
CO ₂ StoP	European CO ₂ storage-potential database/project
EGT	Evolutionary Game Theory
ESS	Evolutionarily Stable Strategy
ETS	Emissions Trading System
EU	European Union
EU ETS	European Union Emissions Trading System
GDP	Gross Domestic Product
IPCC	Intergovernmental Panel on Climate Change
ktCO ₂	Kilotonnes of CO ₂
ktpa	Kilotonnes per annum
MEUR	Million euros
MtCO ₂	Megatonnes of CO ₂
NRW	North Rhine-Westphalia
ODD	Overview, Design concepts and Details
OPEX	Operating Expenditure
ROI	Region of Interest
T&S	Transport and Storage

tCO₂ Tonnes of CO₂

Symbols

Sets and indices

i	Industrial emitter index
d	CO ₂ demander index
c	Cluster index
p	Pipeline index
s	Strategy index or storage-site index, depending on the context
t	Iteration index
k	CO ₂ type index, with $k \in \{\text{fossil, bio}\}$
I	Set of industrial emitters
S_i	Feasible strategy set of emitter i
$s_i(t)$	Strategy adopted by emitter i at iteration t
s_i^{new}	New strategy selected by emitter i during revision
C_t^{CCS}	Set of emitters adopting CCS at iteration t
E_t	Set of active network pipelines at iteration t
S_t	Set of active storage sites at iteration t
$c(i)$	Cluster to which emitter i belongs
$S1-S8$	Base-scenario labels used in the $2 \times 2 \times 2$ scenario matrix

Variables

$\Pi_{i,t}^{\text{NO}}$	Payoff of the No CC strategy for emitter i at iteration t [EUR/yr]
$\Pi_{i,t}^{\text{CCS}}$	Payoff of the CCS strategy for emitter i at iteration t [EUR/yr]
$\Pi_{i,t}^{\text{CCU}}$	Payoff of the CCU strategy for emitter i at iteration t [EUR/yr]
$\Pi_i(s)$	Payoff of emitter i under strategy s [EUR/yr]
$C_i^{\text{route,NO}}$	Annual cost of the no-capture production route for emitter i [EUR/yr]
$C_{i,t}^{\text{route,CC}}$	Annual cost of the capture production route for emitter i at iteration t [EUR/yr]
$C_{i,t}^{\text{CCU}}$	Annual private cost of the CCU strategy for emitter i [EUR/yr]
C_t^{op}	Annual operator cost of the optimised T&S network [MEUR/yr]
$C_{c,t}^{\text{T\&S}}$	T&S cost allocated to cluster c at iteration t [MEUR/yr]
$C_{p,t}^{\text{pipe}}$	Pipeline cost associated with active pipeline p at iteration t [MEUR/yr]
$C_{p,t}^{\text{booster}}$	Booster-station cost associated with active pipeline p at iteration t [MEUR/yr]
$C_{s,t}^{\text{storage}}$	Storage cost associated with storage site s at iteration t [MEUR/yr]

LIST OF FIGURES AND TABLES

$C_{i,t}^{\text{access}}$	Annualised local access cost assigned to emitter i [MEUR/yr]
$R_{i,t}^{\text{CCU}}$	Net CCU revenue of emitter i at iteration t [EUR/yr]
$A_{i,t}^{\text{CCU}}$	CCU accounting adjustment of emitter i at iteration t [EUR/yr]
A_k	Accounting correction applied to CO ₂ type k [EUR/tCO ₂]
Q_i^{cap}	Captured CO ₂ volume of emitter i [tCO ₂ /yr]
Q_i^{fossil}	Fossil captured CO ₂ volume of emitter i [tCO ₂ /yr]
Q_i^{bio}	Biogenic captured CO ₂ volume of emitter i [tCO ₂ /yr]
$Q_{i,t}^{\text{sold}}$	CO ₂ volume sold by emitter i to CCU demand at iteration t [tCO ₂ /yr]
$Q_{i,t}^{\text{fb}}$	Fallback CO ₂ volume of emitter i redirected to CCS [tCO ₂ /yr]
Q_t^{CCS}	Total captured CO ₂ volume of the realised CCS coalition at iteration t [MtCO ₂ /yr]
$Q_{c,t}^{\text{CCS}}$	Captured CO ₂ volume of cluster c at iteration t [MtCO ₂ /yr]
$Q_{c,t}^{\text{pot}}$	Remaining potential captured volume in cluster c [MtCO ₂ /yr]
$Q_{i,k}^{\text{rem}}$	Remaining available CO ₂ supply of type k from emitter i [tCO ₂ /yr]
D_d^{rem}	Remaining CO ₂ demand of demander d [tCO ₂ /yr]
D_{total}	Total CO ₂ demand [MtCO ₂ /yr]
$D_{\text{greenhouse}}$	Greenhouse-horticulture CO ₂ demand [MtCO ₂ /yr]
$D_{\text{F\&B}}$	Food and beverage CO ₂ demand [MtCO ₂ /yr]
$D_{\text{industrial CCU}}$	Industrial CCU demand [MtCO ₂ /yr]
D_i	CO ₂ demand allocated to demand node i [MtCO ₂ /yr]
D_b	Total demand assigned to demand block b [MtCO ₂ /yr]
w_i	Allocation weight of demand node i
$q_{i,k,d}$	Quantity of CO ₂ type k allocated from emitter i to demander d [tCO ₂ /yr]
q_i^{eff}	Effective captured volume used to express access cost per tonne [MtCO ₂ /yr]
x_t	Vector of adoption shares at iteration t
x_t^{NO}	Share of emitters adopting No CC at iteration t
x_t^{CCS}	Share of emitters adopting CCS at iteration t
x_t^{CCU}	Share of emitters adopting CCU at iteration t
d_t	Short-term drift of adoption shares
d_t^L	Long-term drift of adoption shares
σ_t	Volatility of adoption shares over the long convergence window

Prices and tariffs

p^{carbon}	Carbon price [EUR/tCO ₂]
p^{ETS}	ETS price [EUR/tCO ₂]
p^{CDR}	CDR credit value [EUR/tCO ₂]
p^{CO_2}	CO ₂ commodity selling price [EUR/tCO ₂]
p^{biomass}	Dry-biomass reference price [EUR/t dry biomass]

$\tau_{i,t}^{\text{T\&S}}$	T&S tariff faced by emitter i at iteration t [EUR/tCO ₂]
τ^{boot}	Bootstrap T&S tariff [EUR/tCO ₂]
$\bar{\tau}_t^{\text{T\&S}}$	Average T&S cost at iteration t [EUR/tCO ₂]
$\tau_{c,t}^{\text{clus}}$	Realised T&S tariff of cluster c [EUR/tCO ₂]
$\tau_{i,t}^{\text{proxy}}$	Proxy entrant tariff for emitter i [EUR/tCO ₂]
$\tau_{i,t}^{\text{fb}}$	Fallback storage tariff applied to unsold CCU volume [EUR/tCO ₂]
$NR_{i,d}$	Net revenue per tonne from selling CO ₂ from emitter i to demander d [EUR/tCO ₂]
$V_{i,k,d}$	Marginal value of a candidate CCU transaction [EUR/tCO ₂]
$c_{i,d}^{\text{del,E}}$	Emitter-side delivery cost from emitter i to demander d [EUR/tCO ₂]
$d_{i,d}$	Distance between emitter i and demander d [km]

Parameters

$f(N_t)$	Learning factor applied to capture-specific CAPEX
N_t	Cumulative number of emitters that have adopted CCS or CCU up to iteration t [adopters]
N_{ref}	Reference deployment level for the learning curve [adopters]
LR	Learning rate [%]
b	Experience-curve exponent
f_{min}	Minimum learning multiplier
$\text{CAPEX}_{i,t}^{\text{cap}}$	Capture-specific CAPEX of emitter i after learning at iteration t [EUR]
$\text{CAPEX}_{i,0}^{\text{cap}}$	Initial capture-specific CAPEX of emitter i [EUR]
p_{revision}	Probability that an emitter receives a revision opportunity
$P_i(s)$	Probability that emitter i selects strategy s under the logit rule
β	Logit sensitivity parameter
Λ	Payoff scaling parameter [EUR/yr]
μ	Mutation probability
η	Maturity-weighting coefficient for entrant-tariff construction
$\omega_{c,t}$	Weight assigned to the realised cluster tariff in the hybrid entrant tariff
$\alpha_{c,p,t}$	Share of pipeline p attributed to cluster c at iteration t
$\alpha_{c,s,t}$	Share of storage site s attributed to cluster c at iteration t
$F_{p,t}$	Total flow on active pipeline p at iteration t [MtCO ₂ /yr]
$F_{c,p,t}$	Flow on pipeline p attributed to cluster c at iteration t [MtCO ₂ /yr]
q^{floor}	Minimum volume floor used in entrant-tariff calculations [MtCO ₂ /yr]
f_C^{biomass}	Carbon fraction of dry biomass [tC/t dry biomass]
W	Short convergence window [iterations]
W_L	Long convergence window [iterations]
$\varepsilon_{\text{mean}}$	Maximum accepted short-term drift

LIST OF FIGURES AND TABLES

$\varepsilon_{\text{long}}$	Maximum accepted long-term drift
ε_{vol}	Maximum accepted volatility
λ_{ROI}	GDP-based regional allocation factor for the study region
GDP_{BE}	Gross Domestic Product of Belgium [EUR/yr or USD/yr]
GDP_{NL}	Gross Domestic Product of the Netherlands [EUR/yr or USD/yr]
GDP_{NRW}	Gross Domestic Product of North Rhine-Westphalia [EUR/yr or USD/yr]
GDP_{EU}	Gross Domestic Product of the European Union [EUR/yr or USD/yr]

Chapter 1

Introduction

1.1 Context and research problem

Decarbonising heavy industry is one of the most difficult challenges in European climate policy. Energy efficiency, electrification, fuel switching, circularity and material efficiency can reduce a large share of industrial emissions, yet several energy-intensive sectors will still face residual emissions that are hard to eliminate. Sectors such as cement, steel, chemicals and refineries rely on high-temperature heat or produce CO₂ as a direct chemical by-product, and may therefore continue to face residual emissions even after other mitigation measures are applied. For these industries, carbon management has become part of the European climate-neutrality strategy. The European Commission expects captured CO₂ volumes to scale up substantially in the coming decades and frames CO₂ capture, transport, storage and utilisation as part of the decarbonisation portfolio for hard-to-abate industries [20].

However, scaling up cannot be achieved through technology alone. Capture projects require integration with the necessary infrastructure and markets to manage CO₂ after it exits the emitting facility. This mutual dependence has been described as a “*decades-old chicken-and-egg problem*”, where emitters may delay carbon-capture investment until reliable transport and storage are available, while infrastructure developers await committed CO₂ volumes [16]. Consequently, industrial carbon management extends beyond plant-level technology or cost-competitiveness and relies on coordination, timing, and collective infrastructure development.

Two principal pathways in industrial carbon management are carbon capture and storage (CCS) and carbon capture and utilisation (CCU). CCS involves the permanent storage of captured CO₂ in geological formations, whereas CCU directs captured CO₂ to downstream applications such as fuels, chemicals, materials, or greenhouse horticulture. Both pathways rely on downstream storage or utilisation outlets and should therefore be conceptualised as value chains that connect emitters with infrastructure, markets, accounting rules and policy incentives [58, 61].

This value-chain perspective is central to the research problem addressed in this thesis, as exemplified by the Port of Antwerp. The planned Kairos@C project by BASF and Air Liquide was expected to provide the first large captured CO₂ volumes for a shared CO₂ network and export terminal to which other industrial

firms could later connect [17]. This points to a structural tension: industrial carbon management requires simultaneous commitment from multiple actors, each of whom assesses investments from an individual perspective.

Shared CO₂ transport and storage (T&S) infrastructure typically entails substantial fixed costs, with average costs declining as throughput increases. From a system perspective, this dynamic favours shared networks, industrial clusters, and early infrastructure deployment. For individual emitters, however, the associated uncertainties are significant. Such interdependence can result in first-mover disadvantages and path-dependent outcomes: early participation, emitter location and coalition composition can shape whether later CCS adoption becomes easier or remains unattractive [66, 27].

This interdependence constitutes the core problem addressed by this thesis. While CCS and CCU may appear attractive from a coordinated perspective, actual deployment is determined by decentralised decisions made by heterogeneous firms. These firms do not jointly optimise the system; instead, they respond to private costs, anticipated revenues, infrastructure access, carbon-accounting incentives, and expectations regarding the actions of other actors. Aggregate outcomes may stagnate, shift abruptly once a critical mass is reached, or follow a trajectory that diverges from a centrally coordinated, least-cost solution.

The complexity increases when CCS and CCU are considered concurrently. At the system level, both pathways can be complementary, but at the firm level, they may compete for the same CO₂ streams, infrastructure, subsidies, investment capital, and regulatory attention. The relative attractiveness of each pathway to an emitter depends on carbon prices, accounting rules, T&S tariffs, utilisation demand, and geographic proximity to storage sites and utilisation markets [61].

These interdependencies necessitate a modelling approach capable of representing both system-level infrastructure conditions and decentralised adoption decisions. Centralised techno-economic optimisation is effective for identifying cost-efficient CO₂ network configurations and infrastructure layouts, but typically assumes coordination. Evolutionary Game Theory (EGT)¹ and Agent-Based Modelling (ABM)² can capture heterogeneous actors, adaptive decision-making, and path-dependent diffusion. However, frameworks that link decentralised emitter adoption to spatially explicit shared CO₂ infrastructure costs remain limited. This thesis addresses this gap by combining an ABM-EGT adoption model with a spatially explicit centralised CO₂ network-optimisation layer.

1.2 Research question, scope and contribution

This thesis investigates how industrial carbon management strategies diffuse when large emitters choose among available carbon-management options while depending on shared CO₂ transport and storage infrastructure. The three strategies modelled

¹Evolutionary Game Theory studies how strategies evolve within a population over time. The concept is introduced in more detail in Chapter 3.

²Agent-Based Modelling is a computational modelling approach in which a system is represented through individual agents. This is further described in Chapter 3.

are continued operation without carbon capture (No CC), CCS, and CCU. Strategy evolution is represented through an agent-based evolutionary game framework, with CCS payoffs linked to a Centralised CO₂ Network Optimisation model (CCNO) and CCU constrained by geographically explicit demand nodes. This integrated modelling structure is termed the ABM-EGT-CCNO framework. Rather than providing forecasts, the model functions as a mechanism-oriented tool to examine how policy, infrastructure, and decentralised behaviour interact to generate diverse outcomes.

The main research question is:

Under which techno-economic, spatial, and policy conditions does a CO₂ value chain emerge, when industrial emitters' adoption decisions are shaped by interdependencies arising from shared transport and storage infrastructure whose costs depend on its utilisation?

This main question is divided into four sub-questions:

1. How do carbon-accounting rules, T&S tariff design, and CCU availability shape industrial carbon-management outcomes under a reference ETS price?
2. How do shared CO₂ transport and storage infrastructure conditions shape adoption dynamics, particularly the formation of critical mass, tipping points, and lock-in? ³
3. Which evolutionarily stable strategies (ESS)⁴ characterise the adoption dynamics of the system, and under what conditions do they emerge, persist, or lose stability?
4. To what extent do ABM-EGT adoption outcomes diverge from a centralised techno-economic optimisation pathway?

The modelled region covers: Belgium, the Netherlands and North Rhine-Westphalia. This area combines dense industrial clusters, large emitters, proximity to North Sea storage infrastructure and geographically distributed CO₂ utilisation demand.

This thesis contributes by modelling No CC, CCS and CCU as competing strategies, linking decentralised adoption to shared CO₂ infrastructure through the ABM-EGT-CCNO coupling, and representing emitter-facing T&S tariff structures. This enables analysis of critical-mass effects in the CCUS value chain and direct comparison with a centralised techno-economic benchmark.

Chapter 2 situates the research within the existing literature. Chapter 3 presents the model and its components. Chapter 4 reports the simulation results, and Chapter 5 draws out the implications for policy design, infrastructure feedbacks, critical mass and path dependence. Chapter 6 concludes by answering the research questions, summarising the contribution and outlining future research directions.

³In this thesis, critical mass refers to the point at which sufficient captured volume and spatially relevant participation make shared T&S conditions more favourable for further CCS adoption. Tipping points refer to rapid transitions between adoption regimes, while lock-in refers to persistence in a low-capture or pathway-specific outcome.

⁴ESS are stable long-run strategy-distribution attractors in the finite ABM-EGT simulations.

Chapter 2

Literature Review

Chapter 1 framed industrial carbon management as a coordination problem: CCS and CCU may be attractive from a coordinated system perspective, but their diffusion depends on decentralised firm-level decisions, shared CO₂ infrastructure, policy incentives and downstream market conditions. This chapter reviews the literature needed to frame this problem. It first discusses CCS and CCU as value chains, then turns to deployment barriers, shared CO₂ infrastructure, policy design and existing modelling approaches.

2.1 CCS and CCU as industrial carbon management value chains

The literature on industrial carbon management emphasises that CCS and CCU should not be analysed solely in terms of capture technologies. Instead, they must be considered within the context of the value chains necessary for their deployment [20, 33, 58]. Both CO₂ management methods require initial capture at the emitting facility, followed by conditioning and compression of the captured stream. Subsequently, a downstream management solution is required. For CCS, this involves transporting and permanently storing CO₂ in geological storage sites. For CCU, the captured CO₂ is used as a feedstock in other processes, such as the production of fuels, chemicals, materials, and other industrial products [33, 58].

Adopting a value chain perspective alters the assessment of project feasibility. Carbon capture projects cannot be evaluated solely based on cost or the technical performance of the capture unit. Even if capture is technically feasible at a specific site, implementation may be impossible or economically unviable without reliable transport routes, storage sites, hubs, terminals, or end-use markets. Thus, effective carbon management depends on a diverse array of complementary assets and services. In Europe, this downstream dependency is particularly pronounced due to the uneven distribution of industrial emission sources, transport corridors, and geological storage options [66, 20].

The value chain approach highlights that CCS and CCU deployment involves multiple stakeholders. The emitting company represents only one component of this chain. Transport operators, storage site developers, port or hub operators, technol-

ogy providers, regulators, and financial stakeholders all influence project feasibility. These actors often have differing motivations, time horizons, and risk exposures. As a result, industrial emitters may be reluctant to invest in capture without assurance of a downstream market, while transport or storage developers may hesitate to invest without guaranteed, sufficient, and reliable CO₂ volumes. This interdependence throughout the chain creates a coordination problem [58, 66].

This distinction among stakeholders is particularly relevant for CCU, as the entity capturing CO₂ is not necessarily the one utilising it. Captured CO₂ may be transferred to another industrial stakeholder for use as a feedstock. Thus, the profitability of a CCU project depends not only on the emitter’s capture costs but also on downstream demand, product markets, logistics, carbon accounting rules, and contractual agreements between the CO₂ supplier and user [33, 58].

From a system-level perspective, CCS and CCU can be considered complementary pathways, both contributing to industrial carbon management by requiring capture technologies and generating significant CO₂ streams that may justify infrastructure development. An industrial carbon management portfolio can thus include both storage and utilisation, with different pathways suited to various CO₂ streams [20, 33]. However, system-level complementarity does not imply that companies perceive CCS and CCU as complementary options. At the emitter level, captured CO₂ must be allocated to a downstream pathway, necessitating a choice between storage, utilisation, or both. Companies undergo competition between these strategies, especially given limited public subsidies, investment capital, infrastructure capacity, and favourable accounting treatments. Consequently, CCS and CCU may compete for the same CO₂ streams, investors, and political support, despite both offering system-level benefits [61, 33].

These concepts are central to the thesis. Although CCS and CCU may be desirable from a system-wide perspective, individual firms base their actions on profit forecasts, infrastructure access, downstream opportunities, and regulatory incentives. As a result, decentralised decisions can produce diffusion pathways that diverge from system-centered allocations.

2.2 Barriers and feasibility conditions for CCS and CCU deployment

2.2.1 Plant-level techno-economic feasibility

A first set of challenges relates to the technical and economic viability of CCS and CCU at the level of each industrial site. Carbon capture projects generally require capture units, process integration, as well as compression equipment and auxiliary facilities. All of this entails significant upfront investments. These costs generally vary from site to site. Furthermore, these costs depend on the industrial process, the concentration and purity of the CO₂ stream, as well as the extent to which the capture equipment can be integrated into the existing plant. Operating costs can

also reach high levels, as capture requires energy and may affect the efficiency or flexibility of the underlying production process [58, 33].

These initial financial barriers are particularly significant for industrial companies facing competition. The reason for this is that the additional investment and operating costs can undermine the commercial viability of capture. A capture project may therefore be technically feasible but unattractive if the expected value of savings from carbon payments, policy support, or utilisation revenues does not offset the additional costs and risks [58, 33].

Uncertainty exacerbates these technical and economic barriers. The profitability of CCS and CCU depends not only on current technology costs, but also on projections regarding future carbon prices, energy prices, technological progress, political support, and downstream market conditions. Companies must invest now, even if the future value of carbon management remains uncertain. Thus, although capture is technically feasible, investment might be postponed. This is why, as the literature emphasises, the feasibility of CCUS cannot be limited to a simple static comparison of cost per tonne. It must also be considered as an investment choice under conditions of uncertainty [58].

2.2.2 Value-chain and coordination constraints

The second set of challenges arises because CCS and CCU depend on stakeholders and infrastructure external to the emitting facility. Capture technology by itself is insufficient. A CCS project requires access to transport and storage infrastructure, while a CCU project depends on reliable downstream consumers, established product markets, logistics, or contractual agreements. Without these complementary elements, a project may not be feasible, even if capture is technically and economically viable at the plant level [33, 66].

These interdependencies generate coordination challenges across the carbon management value chain. Developers of transport and storage infrastructure often require credible forecasts of future CO₂ volumes before committing to investment. Conversely, emitters may hesitate to invest in capture technologies unless downstream access is assured. CO₂ commodity users also require a reliable supply before investing in utilisation processes, while emitters seek guarantees of long-term demand prior to proceeding with capture. As a result, deployment may be delayed not only by the absence of physical infrastructure but also by uncertainties related to forecasts, contracts, tariffs, storage access, purchase agreements, and the scale of future participation [58, 33].

2.2.3 Institutional, financial and regulatory conditions

A third category of barriers is institutional and financial. CCS and CCU projects are highly capital-intensive and often depend on the credibility of long-term policies. Permit procedures, liability related to long-term storage, monitoring and reporting requirements, lack of clarity regarding responsibilities throughout the value chain, and uncertainty regarding regulatory treatment undermine investor confidence and

increase transaction costs [58, 33]. These obstacles can be particularly difficult to overcome for early-stage projects. These projects face uncertainty not only regarding technology and pricing but also regarding future regulations governing the utilisation of captured CO₂.

This institutional uncertainty has a significant impact on financing. Due to the volatility of carbon prices and the potential dependence of carbon management revenues on public support programmes, as well as uncertain infrastructure tariffs or downstream demand, projects may be viewed as risky by private investors. The literature therefore often refers to the need for a range of support instruments. Examples of such instruments include investment subsidies, contracts for difference, government guarantees, regulated infrastructure models, or other risk-sharing mechanisms. Such instruments can reduce the risk for early adopters and help bridge the gap between the long-term benefits of the system and the risks associated with short-term private investments [58, 66].

Regulatory design is also very important, as it affects the relative competitiveness of CCS and CCU. Accounting rules determine how captured, stored, transferred, or utilised CO₂ is treated. This, in turn, influences the private value of each strategy. We will not elaborate on this issue here, as the competition between storage and utilisation can only be understood with a thorough understanding of carbon accounting and the design of the ETS. All of this is therefore analysed in more detail in Section 2.4.

Not all barriers to the development of carbon capture identified in the literature are explicitly considered in this thesis. Instead, the analysis focuses on the mechanisms that could lead to incentives for adoption: private profitability, dependence on shared infrastructure, coordination issues, carbon pricing, CDR credits, and policy-driven differences between the benefits of CCS and CCU.

2.3 Infrastructure dependence, network effects and path dependence

2.3.1 Shared transport and storage networks

As mentioned earlier, a key aspect of CCS is that it relies on infrastructure that is generally shared by multiple industrial emitters. Once CO₂ is captured at the plant level, it is transported and injected into a geological storage site. In practice, this may involve the development of pipelines, hubs, terminals, transport routes, and storage facilities for an entire group of users connected to the same system rather than for a single emitter [66, 33].

This aspect of shared infrastructure has several implications. First, the costs of CCS are determined, among other things, by the organisation, financing, and use of this transport and storage infrastructure. Second, geography becomes a crucial factor in feasibility. Indeed, emitters located near industrial centers, ports, pipeline corridors, or even connection points to storage infrastructure may have very different

opportunities compared to other emitters that might be isolated. Third, infrastructure creates interdependence among companies.

The literature typically identifies clusters as favourable sites for carbon capture deployment, as multiple large emitters in close proximity can share transport and storage infrastructure, increase initial CO₂ volumes, and reduce average costs [66, 33]. However, the presence of a cluster alone does not guarantee a functioning network; effective collaboration among emitters, infrastructure developers, storage operators, and public agencies is required.

Studies of source-sink matching confirm this spatial heterogeneity, demonstrating the uneven distribution of emitters and storage locations, as well as the dependence of potential transit routes on distance, capacity, and cost constraints [40, 52]. This is a critical consideration for this thesis, as the location of emitters and their access to the network can significantly affect their CCS payoffs.

2.3.2 Economies of scale and first-mover disadvantage

Economies of scale can be achieved with shared CO₂ infrastructure. T&S networks tend to have high fixed costs, but average prices could decrease with higher volumes of CO₂ being moved or stored. This creates a significant incentive to connect CO₂ flows from numerous industrial locations into a shared infrastructure system. Hence, infrastructure development at a suitable scale can lower unit transport and storage costs and enable wider deployment from a system perspective [66, 33].

However, this pricing structure creates a first-mover problem. During the initial phase of deployment, infrastructure is likely to be underutilised. As a result, the first emitters to join a network may face relatively high transit and storage costs, although these costs may decrease per tonne as more users participate. Early adopters may also experience uncertainty regarding future network development, tariff arrangements, and the timing of entry by other enterprises. Consequently, stakeholders may prefer to delay investment until others have acted, even though collective deployment would ultimately reduce costs for all participants [66, 27].

This tension is particularly crucial in a decentralised context. From a central planner’s perspective, early infrastructure investment may appear efficient, enabling future cost savings and increased aggregate capture. However, individual companies focus on their own projected expenses and revenues. If they anticipate high initial tariffs or uncertain access conditions, they may postpone investment. Although infrastructure developers are not explicitly modelled as agents in this thesis, the research indicates that they also face uncertainty and may be hesitant to develop transport and storage capacity without committed capture volumes. This dynamic creates a coordination challenge between investment in capture and infrastructure development.

Public support and risk-sharing mechanisms can mitigate the first-mover disadvantage by reducing the vulnerability of early projects and bridging the gap between short-term private risks and long-term system benefits [66]. This concept is further discussed in Section 2.4.

2.3.3 Network effects, critical mass and path dependence

The use of shared infrastructure implies that CCS could be sensitive to network effects. Network effects are defined here as a situation where the desirability of CCS for one emitter increases with the participation of other emitters, since higher shared CO₂ volumes can improve the utilisation of infrastructure and reduce average transport and storage costs. Golombek et al. show that investment in carbon capture and transport and storage services are interdependent: capture needs available T&S, and T&S profitability depends on the number of emitters using it [27]. This generates a self-reinforcing mechanism where adoption fuels infrastructure construction and lower infrastructure costs further fuels adoption.

This mechanism is in line with Arthur’s study of competing technologies under growing returns to adoption. In that analysis, early events can lead to numerous equilibria, path dependence and lock-in [4]. At the CCS level, increasing returns are generated because the incentive to adopt this technology is partly collective. The benefits for an individual emitter depend not only on its own capture costs, but also on the size and composition of the entire CCS coalition that uses the shared transport and storage network.

This feedback shows that the deployment may depend on reaching a critical mass. Transport and storage infrastructure could remain costly and/or uncertain if not enough emitters sign up. This would thus inhibit further investment in carbon capture. If adoption rises above a threshold, self-reinforcing adoption can occur. But if participation is below that threshold, the system may be trapped in a low-adoption state. Thus, there could be tipping points where the availability of infrastructure, participation, or even policy support encourages the system to wider spread [27].

Such feedback can result in path-dependent outcomes. The extent of initial adoption, infrastructure choices, and the level of government support may determine whether the system stabilises at a low-adoption equilibrium or achieves higher adoption with shared infrastructure. Therefore, the system’s trajectory cannot be assessed solely on the basis of costs. It also depends on the timing of emitters’ adoption decisions, firms’ expectations regarding future infrastructure and policy conditions, and the degree of coordination among stakeholders.

2.4 Policy design, carbon accounting and CCS-CCU competition

2.4.1 ETS design and carbon price incentives

The relative attractiveness of CCS and CCU towards emitters is strongly influenced by policy design. In the European context, it is the EU Emissions Trading System (ETS) that is of critical importance. This ETS creates an economic incentive to avoid, limit or manage CO₂ emissions. A higher carbon price pushes, in principle, industrial emitters towards carbon capture. However, since the regulatory environment might treat stored CO₂ and utilised CO₂ differently, the same carbon price

could create different incentives for CCS and CCU. The profitability of CCS and CCU depend thus not only on the level of the carbon price, but also on the way how stored and utilised CO₂ are treated by the regulatory framework [61, 20].

This is even more important when comparing CCS to CCU. CCS is governed by a fairly simple accounting principle: if the captured CO₂ is permanently stored and not counted as emitted, the emitter avoids the corresponding emissions obligation. As for CCU, the situation is more complicated. The reason for this is that captured CO₂ can be transferred to another sector, converted into a product, temporarily stored, or ultimately re-emitted. How responsibilities and benefits are distributed throughout the value chain therefore has a strong influence on the economic value of CCU. Verbist et al. show that the design of the emissions trading system (ETS) can influence the level playing field between CCS and CCU. This is particularly the case when emissions are transferred from one sector to another [61].

Several studies highlight the need for additional policy and risk-sharing measures, indicating that a carbon price alone might not be enough if firms are uncertain about infrastructure access, support schemes or accounting standards [27, 66, 58].

2.4.2 Carbon accounting, permanence and the CCS–CCU level-playing field

Carbon accounting accounts for one of the most important aspects of the CCS–CCU competition. Permanently stored CO₂ can be seen as eliminated from the emissions stream. The climate effect of utilisation depends on the fate of the carbon after it is introduced into the product system. Some utilisation pathways could retain CO₂ for extended periods, for example, in mineralised materials. Some may hold it just temporarily, like in some fuels or chemical goods, before it is re-released at the conclusion of the product life cycle. This implies that not all CCU paths have the same mitigation value [61, 33].

This has direct consequences on the incentives large-scale emitters perceive. Firms may have less incentive to invest in CCU if utilisation is not credited or recognised in the same way as permanent storage. This would be the case even if CCU could support valuable industrial value chains. And on the other hand, if temporary utilisation is rewarded too strongly, pathways that provide less durable climate benefits than geological storage could be too incentivised [61, 58].

The level-playing field between CCS and CCU is therefore partially a matter of regulatory design. This is defined by the crediting of avoided emissions, the assignment of responsibility for transferred CO₂ and the management of downstream re-emissions [61]. These regulations can generate incentives that favour one approach over another regardless of underlying technical performance. This is relevant for this thesis, as CCS and CCU are modelled as strategic alternatives. Their relative attractiveness is determined not just by costs and revenues but also by policy rules determining the valuation of carbon flows.

2.4.3 Subsidies, risk-sharing and early-stage support

In addition to carbon pricing and accounting, well-targeted support instruments can have a large impact on carbon capture deployment pathways. It is suggested in the literature that carbon pricing alone may not give sufficient incentives for large-scale CCS and CCU diffusion. This is especially true in the early phases of deployment. Projects are capital-intensive, with unknown future revenues and co-ordination challenges. Additional instruments could thus be necessary. Examples of such instruments are investment subsidies, contracts for difference, public guarantees, regulated infrastructure models, risk-sharing mechanisms, and, for eligible permanently stored CO₂ streams, CDR credits or negative-emission payments to reduce private risk and support initial deployment [58, 66, 33, 20].

These instruments may have different impacts on CCS and CCU. For instance, for CCS, de-risking transport and storage infrastructure can reduce the early T&S costs. Moreover, CDR credits can offer an additional revenue stream when stored CO₂ results in negative emissions, e.g. in the case of biogenic CO₂ storage. For CCU, support may instead be linked to downstream market creation. In those two cases, public intervention can affect the viability of carbon management and also which route is more attractive [58, 33, 66].

This consideration is particularly relevant in a decentralised framework, where emitters respond to private incentives rather than centralised directives. Support mechanisms can alter the timing of adoption, reduce first-mover risk, or enhance the profitability of specific options. Consequently, policy support influences not only static costs but also the temporal dynamics of adoption decisions.

2.4.4 Conditions favouring CCS, CCU or mixed outcomes

The literature suggests that no single carbon management approach is likely to be dominant in all circumstances. CCS may be more attractive if access to storage is robust and if transport networks can reach a sufficient scale. Furthermore, if policy-makers recognise permanent storage through avoided costs under the emissions trading system (ETS) with CDR credits, or payments for negative emissions, CCS would become even more attractive. As for CCU, it may be more attractive first and foremost when there are downstream markets for captured CO₂. But also when its use is well-regarded under accounting rules, or when storage is limited, costly, or uncertain. Outcomes combining elements of both may also arise, with variations across regions, sectors, or even at the company level. This could be driven by specific access to infrastructure, market opportunities, and the policy environment [61, 33, 20].

The key consequence for this thesis is that firms select between non-adoption, CCS and CCU, depending on shifting payoffs influenced by carbon prices, accounting rules, infrastructure availability, downstream demand and support mechanisms. Thus, policy influences not only the overall scale of deployment of industrial carbon management but also the division between storage and use. As a result, CCS-

dominated, CCU-dominated, mixed or low-adoption results can arise according to the policy and market contexts.

2.5 Modelling approaches and literature gap

2.5.1 Centralised optimisation and CO₂ network models

The first main part of the literature studies CCS and CCUS deployment through centralised techno-economic optimisation and network models. Typically, these models are developed from the perspective of a central planner aiming to optimise total system cost or infrastructure deployment, subject to technological, geographical, and policy constraints. This analytical approach is particularly valuable for investigating source-sink matching, transport routing, storage allocation, infrastructure sizing, and the spatial organisation of CO₂ networks. The relevance of this approach is demonstrated by the source-sink matching literature reviewed by Lin et al. and the broader review of CCUS network modelling and optimisation by Rahmanta et al. [40, 52]. Recent studies indicate a shift from basic point-to-point source-sink matching towards more complex network representations that incorporate hubs, shared infrastructure, and temporal evolution [40, 52]. The primary contribution of these studies is to delineate the characteristics of a cost-effective carbon management system under coordinated CCS and CCU deployment.

However, these approaches exhibit limitations in explaining decentralised adoption dynamics. Many centralised optimisation models assume coordinated deployment from a system-wide perspective and often fail to represent industrial emitters as independent actors capable of making strategic adoption decisions.

2.5.2 Evolutionary game and agent-based adoption models

A second strand of the literature focuses more directly on the dynamics of adoption and diffusion more directly. Evolutionary game theory models are particularly relevant when adoption depends on strategic interactions among actors with bounded rationality. Rather than assuming immediate adoption of globally optimal strategies, these models examine how strategies evolve as actors adjust their choices in response to relative payoffs, policy conditions, expectations, and the observed behaviour of others.

For instance, Zhao and Liu analyse CCS technology adoption using an evolutionary game model [67]. Their findings indicate that adoption outcomes are sensitive to policy support, supervision, and initial conditions, and that sufficiently strong incentives are required to achieve high adoption rates. Similarly, Nie et al. apply evolutionary game theory to CCU adoption, focusing on strategic interactions between local governments and coal-fired power plants [45]. Their analysis demonstrates that market and policy conditions, including the value of CO₂-derived products and hydrogen-related factors, influence the attractiveness of CCU adoption over time. These studies underscore that adoption is a dynamic process shaped by incen-

tives, expectations, and strategic adjustments, rather than a static cost-comparison exercise.

Agent-based models provide a related yet more flexible approach, as they enable heterogeneous agents to make decentralised decisions and generate aggregate diffusion patterns through their interactions. This approach is particularly valuable when firms differ in size, location, costs, technology choices, or sensitivity to policy incentives. For example, the CCS literature includes a multi-agent simulation by Yu et al. that examines CCS technology diffusion in the cement industry, analysing how deployment is influenced over time by policies such as subsidies, quota allocation, and firm heterogeneity [65]. These models are well-suited for investigating path dependence, lock-in, threshold effects, and the impact of early policy interventions.

These adoption-oriented approaches help explain why low-carbon technologies may fail to diffuse, even when they appear attractive in static techno-economic analyses. However, most of these models represent infrastructure in a simplified way and often analyse CCS or CCU in isolation, rather than as competing strategies available to the same industrial population.

2.5.3 Positioning of this thesis

The literature thus presents two complementary yet only partially integrated perspectives. Centralised optimisation and CO₂ network models are effective for analysing infrastructure, geography, and system costs, while evolutionary game and agent-based models are better suited for representing decentralised decision-making, heterogeneity, bounded rationality, and dynamic adoption. Few studies, however, integrate these dimensions within a single framework.

This situation reveals a significant research gap. Existing models typically either identify cost-effective carbon management systems from a planner-oriented perspective or examine technology adoption using simplified infrastructure representations. Few models consider No CC, CCS, and CCU as mutually exclusive options for the same industrial population, or analyse how decentralised industrial emitters adjust their strategies in response to shared CO₂ transport and storage costs that depend on collective CCS adoption.

The present thesis lies at this interface. It combines an agent-based adoption framework with a techno-economic representation of shared CO₂ transport and storage infrastructure. It compares decentralised adoption trajectories with a centralised techno-economic perspective and investigates whether individually rational decisions can converge towards, fall short of, or diverge from system-oriented deployment outcomes.

Chapter 3

Theory and Methods

3.1 Methodological foundations: EGT and ABM

The model developed in this thesis is based on Evolutionary Game Theory (EGT) and implemented through an Agent-Based Modelling (ABM) approach. EGT provides the theoretical framework for analysing how strategies evolve within a population, while ABM provides a computational way to represent this evolutionary process at the level of individual agents. The model is therefore an agent-based implementation of an evolutionary adoption framework.

This section introduces the methodological framework that underpins the model. It first establishes the relevance of game theory and EGT for addressing strategic adoption problems. Subsequently, it discusses the suitability of ABM for implementing evolutionary dynamics in a finite and heterogeneous population. The section also details the rationale for coupling the ABM-EGT adoption model with a centralised network-optimisation layer for CCS infrastructure. Finally, it derives the methodological implications for interpreting the model's results.

3.1.1 From traditional to evolutionary game theory

Game theory is employed to analyse situations involving strategic interactions among decision-makers. In such contexts, the outcome experienced by an actor depends not only on its own choices but also on the choices made by others. Izquierdo et al. introduce this concept as follows:

“Game theory is a discipline devoted to studying social interactions where individuals’ decisions are interdependent, i.e. situations where the outcome of the interaction for any individual generally depends not only on her own choices, but also on the choices made by every other individual.” [37]

A game can thus be conceptualised as an abstract representation of strategic interaction. A standard way to represent such an interaction is the normal form: a compact description of the players, their available strategies, and the payoffs associated with each possible combination of strategies [37, 59, 63]. Payoffs are assigned

to combinations of choices made by interacting players, rather than to isolated decisions. This interdependence of payoffs is particularly relevant for industrial carbon management, where firms individually select between CCS, CCU, and non-adoption. The relative attractiveness of each option may depend on shared infrastructure, carbon prices, downstream CO₂ demand, policy incentives, and the decisions of other emitters.

In traditional non-cooperative game theory, the main objective is usually to find solutions consistent with individual payoff maximisation. The central concept is the Nash equilibrium which Van der Laan and Tieman call the “cornerstone” concept for predicting the outcome of a non-cooperative game [59]. A strategy profile is a Nash equilibrium if each player’s strategy is a best response to the strategies chosen by the others. That implies that no player can obtain a higher payoff by changing its strategy unilaterally when the other players keep their strategies unchanged [37, 59].

A Nash equilibrium serves as a static reference point: once this configuration is achieved, no individual player has an incentive to unilaterally change strategy while others maintain theirs [37, 59]. However, this concept does not address the process by which such a configuration is attained, nor does it specify which outcome will emerge when multiple configurations satisfy the no-unilateral-deviation condition. This distinction is significant for industrial carbon management, where the central question concerns not only the individual stability of No CC, CCS, or CCU outcomes, but also the mechanisms by which decentralised emitters transition toward these outcomes through repeated strategy revisions.

Evolutionary game theory brings this dynamic dimension into traditional game-theoretic analysis. Rather than focusing only on static equilibrium conditions, it studies how the distribution of strategies changes within a population across a sequence of revision steps. Van der Laan and Tieman explain this contrast clearly: traditional game theory focuses on static Nash equilibria, while evolutionary game theory models a process in which the number of individuals using a strategy changes across successive revisions according to that strategy’s success [59].

This idea originally came from biology, where strategies were interpreted as inherited traits and payoffs as reproductive fitness. In socio-economic applications, however, strategies are choices, behaviours or routines that agents may revise during their lifetime. Strategy change can therefore result from learning, imitation, comparison between alternatives, or responses to changing economic and policy conditions. Izquierdo et al. describe this socio-economic interpretation as follows:

“In many contexts, it seems more natural to assume that players are capable of adapting their behavior within their lifetime, occasionally revising their strategy in a way that tends to favor strategies leading to higher payoffs over strategies leading to lower payoffs.” [37]

This interpretation is adopted in the present thesis. Industrial firms are not presumed to coordinate instantaneously on a system-wide optimum. Rather, they are modelled as actors that periodically reconsider their current carbon-management

strategies. The central concept is adaptive strategy revision: emitters compare the relative attractiveness of available options and may shift toward strategies that offer higher private payoffs.

In this thesis, the evolutionary adoption logic is built around four elements: a finite population of industrial emitters; three possible strategies, namely continued operation without carbon capture (No CC), carbon capture and storage (CCS), and carbon capture and utilisation (CCU); a payoff structure describing the private attractiveness of each strategy; and a revision rule describing how emitters update their strategy across revision steps.

The payoff structure connects the economic and policy environment to adoption dynamics, while the revision rule translates payoff differences into strategy changes. EGT is thus well-suited for technology adoption problems, as it conceptualises diffusion as a dynamic process rather than a static cost comparison. Initial conditions, payoff differences, stochasticity, and revision rules can determine whether the system remains locked in one configuration, converges to another, or transitions toward broader adoption. In this thesis, EGT provides the theoretical foundation for analysing No CC, CCS, and CCU as competing strategies. The model is not intended to solve an analytical equilibrium game; instead, it simulates repeated strategy revisions in a finite population and examines the aggregate adoption patterns that become approximately stationary.

3.1.2 Why an agent-based implementation?

The preceding subsection introduced EGT as the theoretical framework for describing the evolution of strategies within a population. Agent-based modelling operationalises this framework into a computational model in which individual emitters are explicitly represented, and aggregate adoption patterns emerge from their repeated strategy revisions.

An agent-based implementation is appropriate in this thesis because the adoption process combines several features. First, the population of industrial emitters is finite and heterogeneous. Emitters differ in their emissions, capture costs, geographical location, access to CO₂ transport and storage infrastructure, and opportunities for CO₂ utilisation. Representing emitters individually makes it possible to keep differences between sites explicit. In a representative-agent formulation, emitters would be aggregated into an average industrial actor, which would hide differences in location, emissions, capture feasibility and access to CO₂ infrastructure. Railsback and Grimm explain this distinction by stating that ABMs represent a system’s individual components and their behaviours, rather than describing the system only through variables defined at the system level [53]. In this thesis, these individual components are industrial emitters, each with their own characteristics, current strategy and strategy-revision rule.

Second, ABM is suitable because adoption is modelled as a sequence of individual revision decisions, rather than as a single aggregate system-wide choice. Emitters do not make a one-time decision at the start of the simulation; instead, they receive

revision opportunities during the iterative process, evaluate their feasible strategies, and update their choices according to either deterministic best response or stochastic choice. This approach is important because the aggregate trajectory can depend on which emitters revise, the timing of revisions, the initial distribution of strategies, and the stochasticity of the revision rule. The ABM implementation thus enables path dependence and the emergence of multiple approximately stationary outcomes from repeated individual decisions.

Third, the outcome of interest is not only located at the level of individual emitters. The thesis is also concerned with system-level patterns: the aggregate shares of No CC, CCS and CCU, the emergence or absence of a CCS coalition, the amount of captured CO₂, and the approximate stationarity of these aggregate patterns. This is where the concept of emergence becomes important. Railsback and Grimm describe this logic as follows:

“Using ABMs lets us address problems that concern emergence: system dynamics that arise from how the system’s individual components interact with and respond to each other and their environment.” [53]

This idea is directly connected to the model developed in this thesis. Individual emitters revise their strategies one by one, but the accumulation of these revisions changes the aggregate distribution of No CC, CCS and CCU. These aggregate outcomes are not imposed directly by the modeller. They emerge from repeated agent-level decisions under the payoff structure and revision rule defined in the model.

Finally, ABM is advantageous because it facilitates feedback between individual and aggregate levels. In this thesis, the number and composition of CCS adopters influence the utilisation and cost conditions of shared CO₂ transport and storage infrastructure. These system-level conditions subsequently affect the CCS payoff encountered by emitters in later revision opportunities. The model therefore examines both how firms influence the system and how the evolving system, in turn, affects firms’ subsequent decisions. This bidirectional relationship between decentralised decisions and aggregate outcomes is a primary justification for employing an agent-based implementation.

3.1.3 Linking adoption dynamics with network optimisation

The ABM-EGT framework describes how industrial emitters revise their carbon management strategies in a decentralised way. However, CCS adoption cannot be evaluated only at the level of individual capture costs. Its attractiveness also depends on the shared CO₂ transport and storage infrastructure that becomes available when several emitters adopt CCS.

This creates a methodological challenge. On the one hand, a decentralised adoption model is needed to represent heterogeneous emitters, payoff-based revision and path-dependent diffusion. On the other hand, the cost of CCS depends on a system level infrastructure problem: given a set of CCS adopters and captured CO₂

volumes, the corresponding transport and storage network must be evaluated. This thesis therefore links the ABM-EGT adoption model to a Centralised CO₂ Network Optimisation model, hereafter referred to as the CCNO model.

Adoption remains determined by the ABM-EGT model through emitter-level payoff evaluation and strategy revision. As for the CCNO model, its role is to evaluate the transport and storage network associated with the CCS coalition that emerges from the ABM. Its outputs provide network-cost information that can be translated into CCS transport and storage cost signals for later payoff evaluations.

This coupling establishes a feedback loop between decentralised adoption and shared infrastructure. Emitter choices determine the realised CCS coalition, which in turn defines the network conditions evaluated by the CCNO model. The resulting cost signals produced by the CCNO model subsequently influence the attractiveness of CCS in later revision steps. This methodological linkage connects decentralised industrial adoption dynamics with the techno-economic logic of CO₂ network deployment.

3.1.4 Methodological implications for this thesis

Collectively, the EGT, ABM, and CCNO components constitute a mechanism-oriented adoption framework rather than a deployment forecast. EGT supplies the logic of payoff-based strategy revision, ABM operationalises this logic for a finite and heterogeneous population of industrial emitters, and the CCNO coupling introduces infrastructure feedback, whereby collective CCS adoption influences subsequent CCS payoffs. These components enable the analysis of how decentralised choices generate aggregate adoption patterns, critical-mass effects, lock-in, path dependence, and approximately stationary outcomes.

This approach is directly aligned with the thesis objective. The aim is not to determine which industrial carbon management pathway would be cost-efficient from a coordinated system perspective, but rather to analyse whether such a pathway can emerge when firms act individually, encounter heterogeneous payoff conditions, and interact via shared CO₂ infrastructure. The framework thus enables a comparison between coordinated techno-economic deployment logic and decentralised adoption logic.

3.2 Model description using the ODD protocol

This section describes the implemented model using the ODD protocol, a standard format for making agent-based models transparent, understandable and reproducible [29, 30, 53]. ODD stands for *Overview, Design concepts and Details*. The main text follows this logic, but focuses on the elements needed to understand the model: the model's purpose, agents and state variables, process scheduling and submodels. The eleven ODD design concepts are reported separately in Appendix A, and detailed numerical inputs and parameter values are summarised in Appendix B.

3.2.1 Purpose and general model logic

The model serves three interconnected purposes. First, it analyses the diffusion of industrial carbon management strategies among large-scale industrial emitters. Each emitter may select among three strategies: continued operation without carbon capture (No CC), carbon capture and storage (CCS), or carbon capture and utilisation (CCU). The objective is to provide insight into how expected private profitability, shared CO₂ infrastructure, policy assumptions, and market conditions influence aggregate adoption outcomes.

Second, the model is employed to investigate the dynamics underlying these outcomes. Simulations are conducted from various initial distributions of No CC, CCS, and CCU, enabling assessment of whether the system converges to a single final state or to multiple stable adoption regimes. This approach also facilitates analysis of critical-mass effects by testing whether self-reinforcing diffusion requires a sufficiently large or strategically positioned group of capture adopters. The model thus examines both final adoption shares and the transition dynamics leading to storage-dominated, utilisation-dominated, mixed, or persistent low-adoption outcomes.

Third, the model adopts a decentralised modelling perspective, in contrast to a centralised techno-economic approach. While a centralised optimisation method can identify a cost-efficient deployment pathway for CO₂ transport and storage infrastructure under coordinated system-level decisions, the agent-based evolutionary model examines whether such a pathway can emerge when emitters independently adjust their strategies based on individual payoffs and local conditions. This distinction is central to the research question, as a pathway that appears optimal from a system perspective may not be realised through decentralised adoption.

3.2.2 Entities, state variables, scales and input datasets

Agents

The model contains three types of agents: *industrial emitters*, *CO₂ demanders*, and a passive *Carbon Network Operator* (CNO). Industrial emitters are the strategic and adaptive agents: they represent large-scale industrial CO₂ emitting sites and can revise their carbon-management strategy. CO₂ demanders are passive demand agents for the CCU pathway. They do not revise strategies, but constrain how much captured CO₂ can be allocated to utilisation. The CNO is a passive infrastructure agent for the CCS pathway. It represents the shared T&S infrastructure layer associated to the group of emitters adopting CCS, namely the realised CCS coalition.

Each industrial emitter represents one large-scale industrial site located in the trilateral study region shown in Figure 3.1. In the current regional configuration, the model creates 227 industrial emitter agents. The emitter population is heterogeneous and includes several industrial sectors, such as chemicals, glass, cement, light-liquid fuel production/refineries, steel and fertilisers. This sectoral distinction




FIGURE 3.1: Spatial distribution of industrial emitters and CO₂ demand points in the trilateral study region, covering Belgium, the Netherlands and North Rhine-Westphalia. Green points represent emitters and red points CO₂ demand points.

matters because emitters differ in their production routes, emissions profile, capture feasibility and costs.

The ABM emitter agents are created from scenario-specific preprocessed emitter input files. These files are generated from the same techno-economic industrial dataset used to prepare the centralised CCNO model inputs. Upstream of this pre-processing, the dataset is based on AIDRES-related industrial data¹ and contains information on industrial production sites, including geographical coordinates, sector, product type, production quantity, emissions, production routes, CAPEX and OPEX [21]. The ABM therefore does not use the full industrial database directly; it uses the preprocessed route-cost files obtained before the simulation starts, following the same logic as in the CCNO input preparation [62].

For each emitter, this preprocessing selects two representative production routes. The first route is the lowest-cost available route without carbon capture and provides the techno-economic basis for the No CC strategy. The second route is the lowest-cost available route with carbon capture, when such a route exists, and provides the common upstream basis for both CCS and CCU. If no feasible capture route exists for an emitter, capture-based strategies are removed from its feasible strategy set. The distinction between CCS and CCU therefore does not lie in the upstream capture route, but in the downstream management of the captured CO₂: permanent geological storage for CCS, or allocation to utilisation demand for CCU.

The emitter input files are scenario-specific because the route selection and route costs are computed under the carbon price and carbon-accounting assumptions of

¹AIDRES stands for *Advancing Industrial Decarbonisation by Assessing the Future Use of Renewable Energies in Industrial Processes*. It is a European Commission Joint Research Centre dataset describing industrial sites and decarbonisation routes through site-level techno-economic data.

the run. For example, in scenarios with CDR credits, the value of eligible negative emissions is already reflected in the preprocessed capture-route costs and in the route selected for each emitter.

Table 3.1 summarises the principal properties of industrial emitter agents. Static properties refer to fixed characteristics that remain constant throughout the simulation. In contrast, non-static properties may change during the simulation, either due to strategy revisions by the agent or as a result of learning effects that alter the capture-route CAPEX. The behavioural strategy describes how emitters update their strategies during the evolutionary process.

TABLE 3.1: Industrial large-scale emitters – properties

Industrial large-scale emitters		
Static	Non-static	Behavioural strategy
Emitter ID, AIDRES site ID and node ID	Current strategy: <i>No CC</i> , <i>CCS</i> or <i>CCU</i>	Emitters revise their strategy only when they receive a revision opportunity. Revising emitters calculate their payoffs and choose their strategy based on those payoffs.
Geographical location [lat., lon.]	Capture-route CAPEX after learning	
Sector and sector ID		
Production capacity [ktpa]		
Capture feasibility flag [yes/no]		
No-capture route CAPEX and OPEX ²		
Capture-route CAPEX and OPEX ³		
Potential captured CO ₂ volume [tCO ₂ /yr]		
Residual emissions under the capture route [tCO ₂ /yr]		
No-capture route emissions [tCO ₂ /yr]		

The second type of agent represents CO₂ demand points for utilisation. These agents are passive downstream outlets for the CCU pathway. They are created from the spatial CO₂ demand database described in Appendix C. This database represents a 2050-oriented demand scenario for the trilateral region. The database integrates three demand categories: greenhouse horticulture, conventional food and beverage demand, and industrial CCU demand. Details regarding the construction and spatial allocation of this demand are provided in Appendix C. The main model utilises only the resulting demand nodes.

In the reference demand scenario, 25 geographically explicit CO₂ demanders are established. Each demander is characterised by a fixed location, sectoral classification, and an initial annual demand, which represents the maximum quantity of CO₂ that can be absorbed by the node in a single model run. This demand does not have to be fully met: unmet CO₂ demand is allowed and does not create a penalty in the model. The demanders are external outlets, as the model does not account

³Cost variables are reported in the units of the preprocessed route-cost input files and are converted into annual private costs in the payoff submodel. Transport and storage costs for CCS, as well as potential revenues from CCU, are not included in these route-cost terms and are added separately in the payoff calculation.

3. THEORY AND METHODS

for emitters utilising their own captured CO₂ internally. The principal properties of these demanders are summarised in Table 3.2.

TABLE 3.2: CO₂ demanders – properties

CO ₂ demanders			
Static	Non-static		Behavioural strategy
Country and subsector	Remaining CO ₂ demand [tCO ₂ /yr]		No strategy revision. Demanders are passive external demand nodes in the CCU allocation process.
Hub or cluster name	Allocated CO ₂ received through CCU matching [tCO ₂ /yr]		
Geographical location [lat., lon.]			
Initial annual CO ₂ demand [tCO ₂ /yr]			

The third agent type is the Carbon Network Operator (CNO), which serves as a passive infrastructure agent within the carbon capture and storage (CCS) pathway. The CNO represents the shared CO₂ transport and storage infrastructure associated with the realised CCS coalition. It does not possess a carbon-management strategy, revision rule, or adoption decision.

After a CCS coalition is formed, the CNO employs the CCNO model to evaluate the corresponding transport and storage (T&S) network for participating CCS emitting sites and their captured CO₂ volumes. This evaluation yields the activated network topology, pipeline flows, storage utilisation, and total annual operator cost. These outputs inform the construction of T&S tariff signals, which are subsequently used in emitter payoff calculations. Thus, the CNO functions as the passive infrastructure interface between decentralised emitter adoption and techno-economic network evaluation. The main properties of the CNO are summarised in Table 3.3.

TABLE 3.3: Passive Carbon Network Operator agent – properties

Passive Carbon Network Operator agent		
Static properties	Non-static properties	Role in the model
Candidate T&S network, routing nodes and pipeline	Current participating CCS sites	No strategy revision. The CNO is a passive infrastructure agent that updates the T&S network state associated with the realised CCS coalition and provides operator-cost information used to construct tariff signals.
Candidate terminals, offshore connectors and storage sites	Activated network topology	
Investment-cost assumptions for pipelines and storage assets	Pipeline flows and capacity use [MtCO ₂ /yr]	
OPEX assumptions for pipelines, booster stations and storage	Storage use and allocated storage volumes [MtCO ₂ /yr]	
Network-cost and operator-cost parameters	Total annual network operator cost [MEUR/yr]	
Tariff-construction mode	Average and differentiated T&S tariff indicators [EUR/tCO ₂]	

Environment

All three agent types operate within a unified global model environment. This environment encompasses scenario assumptions, aggregate system variables, and run-level information that influence emitter payoffs. It is not modelled as an additional strategic agent; it does not revise strategies or make adoption decisions. Rather, it provides the economic, policy, and system-level context in which industrial emitters evaluate No CC, CCS, and CCU options.

The environment stores three types of information. First, it stores scenario-level assumptions that remain constant throughout a simulation run, including the carbon price, CO₂ commodity price assumptions, whether carbon dioxide removal (CDR) credits are included and route-accounting assumptions, CCU transport assumptions, capture-learning parameters, CCS tariff mode, and the CCU demand scenario.

Second, the environment stores run-level behavioural inputs, such as initial strategy shares, revision probability, revision rule, stochastic parameters, random seed, and convergence settings. These inputs are not specific to individual emitters but instead define the initialisation and updating of the evolutionary adoption process. The corresponding numerical values are provided in Appendix B.

Third, the environment tracks aggregate variables that are updated throughout the simulation. These variables include adoption shares, captured volumes, CCU sales, fallback flows from CCU to CCS, total emissions, and the most recent T&S tariff signals received from the CNO for subsequent payoff evaluations.

TABLE 3.4: Global environment – properties

Global environment	
Static	Non-static
Scenario identifiers	Adoption shares of <i>No CC</i> , <i>CCS</i> and <i>CCU</i> [-]
Carbon price p_{carbon} [EUR/tCO ₂]	Total CO ₂ captured through CCS [tCO ₂ /yr]
CO ₂ commodity price and pricing mode [EUR/tCO ₂ ; mode]	Total CO ₂ captured through CCU [tCO ₂ /yr]
CDR credit and route-accounting assumptions	Total CO ₂ not captured [tCO ₂ /yr]
Learning rate for capture CAPEX and minimum CAPEX fraction	Total net emissions [tCO ₂ /yr]
Transport cost assumption for CCU allocation [EUR/tCO ₂ /km]	Biogenic and fossil captured CO ₂ volumes [tCO ₂ /yr]
T&S modelling mode and tariff method	Fallback CO ₂ volumes from CCU to CCS [tCO ₂ /yr]
CCU demand scenario or demand horizon	Latest T&S tariff signals used in payoff evaluation [EUR/tCO ₂]

The environment serves as the interface where agent-level states are aggregated into system-level quantities. When emitters revise their strategies, the aggregate shares of No CC, CCS, and CCU are updated accordingly. These shares influence total captured volumes, the allocation of CO₂ to demanders, and the composition of the CCS coalition communicated to the CNO/CCNO infrastructure layer. Tariff signals generated by the passive infrastructure agent are stored in the environment and subsequently affect emitter payoff evaluations during later revision opportunities.

The model environment is spatially explicit, with a scope corresponding to the trilateral study region comprising Belgium, the Netherlands, and the German province of North Rhine-Westphalia. Industrial emitters and CO₂ demanders are represented by latitude and longitude coordinates. This spatial data is utilised to estimate distances, identify CCU allocation opportunities, and associate emitters with infrastructure-related conditions. The model does not employ grid cells or a continuous spatial landscape; instead, space is represented through geographically located agents and their infrastructure relationships.

The temporal scale of the model requires careful interpretation. Iterations correspond to adaptive revision steps rather than calendar years; thus, the model does not simulate a chronological deployment pathway from the present to a future date. Nevertheless, inputs and outputs are annualised, with route costs, captured volumes, demand quantities, emissions, and tariffs all expressed on an annual basis. In the reference CCU-demand scenario, the demand database reflects a 2050-oriented input. Consequently, the model examines adoption patterns that may arise under a specified annualised scenario environment, rather than simulating a real-time transition from the present to 2050.

3.2.3 Process overview and scheduling

This subsection outlines the sequence in which model processes are executed during a simulation run, from initialisation to convergence. The objective is to present the scheduling logic of the agent-based model (ABM); detailed algorithms for payoff calculation, CCS T&S tariff construction, CCU allocation, learning, strategy revision, and convergence are provided in the relevant submodels in Section 3.2.4 and in Appendix D. The complete process sequence is summarised in Figure 3.2.

The model distinguishes between two categories of processes: entity processes and observer processes. Entity processes alter the state of the model; for example, emitter agents may revise their strategies, CO₂ demanders may receive updated CCU allocations, the passive CNO agent may update the T&S network state for the realised CCS coalition, and aggregate system variables may be recalculated. In contrast, observer processes do not influence agent decisions but solely record information generated during the simulation. These outputs include graphs of adoption shares, switching events, captured CO₂ volumes, emissions, tariffs, costs, and final Excel/CSV files. This distinction clarifies which steps drive the simulated dynamics and which are intended for analysis and reporting.

Before the iterative loop starts, each simulation run is initiated from a new model state. No agent state, coalition, network result, or historical variable is transferred from previous runs. The model first loads run-specific parameters, including the route-cost input set, carbon price, carbon-accounting setting, CCS tariff mode, CCU price and transport assumptions, learning parameters, revision rule, convergence settings, random seed, and initial strategy shares.

Subsequently, industrial emitter and CO₂ demand data are loaded, and the global environment and agent populations are instantiated. In the current regional config-




FIGURE 3.2: Process flowchart of the agent-based evolutionary model. Green boxes indicate begin and final steps, purple boxes indicate entity creation and non-static properties/state updates, blue boxes indicate system updates and output recording, and the yellow diamond indicates the convergence check. Red outlines indicate where the CCNO model is called to initialise or update the CCS T&S network.

uration, this results in 227 industrial emitter agents and 25 geographically explicit CO₂ demanders. Dynamic variables, including captured volumes, CCU sales, emissions, and remaining demand, are initialised prior to any strategy revision.

Initial strategies are assigned according to the selected initial shares. The emitter list is randomly reordered using the run random seed⁴, so that the same seed always reproduces the same initial allocation. The required number of emitters is then assigned to CCS and CCU, and the remaining emitters start in No CC. For example, if the initial shares of CCS and CCU are both zero, all emitters start in No CC.

The model then computes and records the initial system state as iteration 0. This state includes initial strategy shares, captured volumes, emission indicators, CCU demand status, learning factor, and available T&S tariff information. The

⁴A random seed is a fixed value used to initialise randomness, making a stochastic simulation run reproducible.

initial T&S reference is established according to the CCS tariff-construction logic outlined in Section 3.2.4.3.

Following initialisation, the model enters the main iteration loop. At the start of iteration t , the model utilises the realised state computed at the end of iteration $t - 1$. As a result, the realised CCU allocation, aggregate system quantities, and the most recent T&S tariff signals are already stored in the environment.

The first step in each iteration is to retrieve the latest T&S tariff signals and update the capture learning factor. Tariff signals are incorporated into both CCS and CCU payoff calculations. For CCU, these signals determine the fallback storage cost, which is the T&S cost applied to captured CO₂ that remains unsold due to saturated demand or the absence of profitable utilisation matches, resulting in redirection to CCS storage. The learning factor adjusts the capture-route CAPEX used in both CCS and CCU payoff calculations.

Subsequent steps involve emitter strategy revision. Only emitters granted a revision opportunity compute payoffs for their feasible strategies; others retain their current strategy. For CCS, the payoff calculation utilises the most recent applicable T&S tariff signal stored at the end of the previous realised state. For CCU, the payoff incorporates a hypothetical allocation in which the revising emitter is temporarily included among current CCU adopters. This approach enables evaluation of potential sales, revenues, accounting adjustments, and fallback quantities without altering the realised CCU market state.

Revised strategies are initially stored temporarily and are applied synchronously only after all revising emitters have completed their decisions. This procedure ensures that all revising emitters assess their options under identical pre-revision system conditions. The observer process records the computed payoffs and the strategies selected.

Once revised strategies have been applied, the realised system state is refreshed. The updated CCS adopters define the realised CCS coalition:

$$C_t^{\text{CCS}} = \{i \in I \mid s_i(t) = \text{CCS}\}. \quad (3.1)$$

Here, I is the set of industrial emitters and $s_i(t)$ is the strategy of emitter i at iteration t . The set C_t^{CCS} therefore contains all emitters currently adopting CCS.

This realised CCS coalition serves as the primary interface between strategic emitter agents and the passive CNO agent. When the T&S network state is updated, the agent-based model transmits the realised CCS coalition, captured CO₂ volumes, emitter-node mapping, and scenario or network data to the CNO. The CNO then employs the CCNO model to evaluate the cost-efficient T&S network for that coalition, including the activated network topology, CO₂ flows, storage utilisation, and network costs. This ABM-CCNO coupling is shown in Figure 3.3.

The network state provided by the passive CNO agent is subsequently forwarded to the CCS T&S tariff-construction submodel, as described in Section 3.2.4.3. This submodel translates network operator costs and flow information into T&S tariff




FIGURE 3.3: ABM–CCNO coupling inside the iterative loop.

signals, which are stored in the environment for use in the next payoff evaluation. This sequence ensures that emitters revise their strategies using the latest available tariff information, while the outcomes of their collective revisions influence the network-cost signals encountered in subsequent revision opportunities.

The CCU market is then cleared for the realised set of CCU emitters, determining actual CO_2 sales, delivery costs, revenues, accounting adjustments, and fallback flows. Fallback flows refer to unsold captured CO_2 volumes that are redirected to CCS storage. If, after market clearing, a CCU adopter sells no CO_2 in the realised allocation, it is reclassified to either CCS or No CC, depending on which alternative yields a higher payoff.

The model then aggregates the current system quantities, including adoption shares, captured volumes, CCU sales, CCS storage, fallback flows, emissions, tariff indicators, and adoption events. These quantities form the realised state at the end of the iteration and are recorded by the observer process. The convergence test is then applied to the aggregate adoption shares, as described in Appendix D. This test focuses on the stationarity of aggregate shares rather than on the complete absence of individual switching. If the stationarity criterion is satisfied, the simulation stops and final results, figures, maps and Excel/CSV files are exported. Otherwise, a new iteration begins from the realised state that has just been computed.

3.2.4 Submodels

3.2.4.1 Learning-by-doing for capture CAPEX

The model includes a learning-by-doing effect for carbon capture technology. This represents the idea that investment costs decrease as cumulative deployment increases. The formulation follows a standard one-factor experience curve, commonly used to represent cost reductions from cumulative capacity or experience [55, 54]. Learning is applied only to the capture-specific CAPEX component of the capture route. It is not applied to OPEX, non-capture route costs or T&S costs [23].

The learning factor is computed as:

$$f(N_t) = \max \left(f_{\min}, \left(\frac{\max(N_t, N_{\text{ref}})}{N_{\text{ref}}} \right)^b \right), \quad b = \frac{\log(1 - LR)}{\log(2)}. \quad (3.2)$$

Here, $f(N_t)$ is the learning factor applied to capture-specific CAPEX, LR is the learning rate, and b is the experience-curve exponent implied by that learning rate. The learning rate represents the fractional CAPEX reduction obtained after each doubling of cumulative capture adoption. In this thesis, cumulative experience is measured by N_t . This is the cumulative number of emitters that have adopted a capture-based strategy, either CCS or CCU, up to iteration t . This adoption-experience measure is used because the ABM tracks discrete emitter adoption rather than installed capture capacity. The updated capture-specific CAPEX is then:

$$CAPEX_{i,t}^{cap} = f(N_t) CAPEX_{i,0}^{cap}. \quad (3.3)$$

The model uses $LR = 5\%$, $N_{\text{ref}} = 1$, and $f_{\text{min}} = 0.75$. The learning rate is treated as a conservative assumption, since capture learning estimates are uncertain and technology-specific [55, 54]. The parameter $N_{\text{ref}} = 1$ normalises learning to the first capture adopter, while $f_{\text{min}} = 0.75$ prevents CAPEX reductions beyond 25% [23]. With these parameters, each doubling of cumulative capture adopters reduces the capture-specific CAPEX by 5%, until the plateau of 75% of the initial value is reached after 49 cumulative adopters, as illustrated in Figure E.1.

3.2.4.2 Payoff calculation

Following the evolutionary-game logic used in this thesis, agents compare strategies through payoffs before applying the strategy revision rule [37]. In the model, a payoff is defined as the opposite of the total annual private cost of a strategy. A higher payoff therefore means a lower annual cost.

For a revising emitter i at iteration t , the three strategy payoffs are defined as:

$$\begin{aligned} \Pi_{i,t}^{NO} &= - \underbrace{C_i^{route,NO}}_{\text{no-capture route cost}}, \\ \Pi_{i,t}^{CCS} &= - \left(\underbrace{C_{i,t}^{route,CC}}_{\text{capture route cost}} + \underbrace{Q_i^{cap} \tau_{i,t}^{T\&S}}_{\text{T\&S cost for stored CO}_2} \right), \\ \Pi_{i,t}^{CCU} &= - \left(\underbrace{C_{i,t}^{route,CC}}_{\text{capture route cost}} + \underbrace{Q_{i,t}^{fb} \tau_{i,t}^{fb}}_{\text{fallback CCS cost}} + \underbrace{A_{i,t}^{CCU}}_{\text{CCU accounting adjustment}} - \underbrace{R_{i,t}^{CCU}}_{\text{net CCU revenue}} \right). \end{aligned} \quad (3.4)$$

The *No CC* payoff relies only on the best no-capture production route cost. The CCS payoff uses the best capture production route and adds the cost of transporting and storing the emitter's full captured CO₂ volume. The CCU payoff uses the same capture production route, but allows part of the captured CO₂ to be sold to utilisation demand. Any unsold captured CO₂ becomes a fallback flow and is charged the fallback storage cost.

During the iterative loop, payoffs are computed only for emitters that receive a revision opportunity. For each revising emitter, the model evaluates the feasible strategies among No CC, CCS and CCU. If no technically feasible capture route exists, the capture-based strategies are removed from the feasible set. The payoff calculation uses the latest realised system state available at the time of revision. In particular, the CCS tariff is not recomputed by solving the CCNO model for each candidate strategy. Instead, it is taken from the T&S tariff information stored at the end of the previous realised state.

The route-cost terms are imported from the preprocessed techno-economic input data:

$$C_i^{route,NO} = \text{Totex_noCC_ptpa}_i, \quad C_{i,t}^{route,CC} = \text{Totex_CC_ptpa}_{i,t}. \quad (3.5)$$

They represent, respectively, the best available production route without carbon capture and the best available production route with carbon capture. The capture-route cost can evolve during the simulation if learning reduces the capture-specific CAPEX. These route costs are scenario-specific and already include the ETS and carbon-accounting assumptions. The ABM therefore does not add a second ETS or CDR term to the base route costs.

For CCS, Q_i^{cap} is the captured CO₂ volume and $\tau_{i,t}^{T\&S}$ is the applicable network-access tariff faced by emitter i at iteration t . This tariff is taken from the latest realised T&S network state returned by the CNO/CCNO layer. Its detailed construction from network outputs is described in Section 3.2.4.3.

For CCU, the captured CO₂ volume is split into a sold quantity and a fallback quantity (e.g. if only a part of the captured CO₂ is enough to meet the CO₂ commodity demand, and the rest must therefore be handled in another way):

$$Q_i^{\text{cap}} = Q_{i,t}^{\text{sold}} + Q_{i,t}^{\text{fb}}. \quad (3.6)$$

In the CCU payoff (Equation 3.4), the term $Q_{i,t}^{\text{fb}} \tau_{i,t}^{\text{fb}}$ represents the storage cost of the unsold captured CO₂. The term $R_{i,t}^{\text{CCU}}$ is the net revenue from CO₂ sales after delivery costs. The term $A_{i,t}^{\text{CCU}}$ is the accounting adjustment applied because CCU demand is treated as non-permanent in the standard model configuration. These CCU-specific quantities are furthermore defined in Section 3.2.4.4.

The payoff submodel therefore provides the payoff vector used by the strategy revision rule. This vector does not determine adoption on its own. It is passed to the strategy-revision submodel, where either deterministic best response or stochastic logit choice converts payoff differences into strategy updates. These revision rules are described in Section 3.2.4.5.

3.2.4.3 CCS T&S tariff construction from network outputs

As explained in the process overview, the CCS payoff requires a network-access tariff, $\tau_{i,t}^{T\&S}$, for each revising emitter. This tariff is not fixed exogenously for the whole

simulation. It is updated from the realised CCS coalition and the corresponding T&S network outputs produced by the passive CNO agent using the CCNO model. This submodel explains how network operator costs and flow information are converted into tariff signals stored for later payoff evaluation.

If the realised CCS coalition is too small for a meaningful shared-network evaluation, the model uses a fixed bootstrap tariff for all technically feasible emitters:

$$\tau_{i,t}^{T\&S} = \tau^{boot}, \quad \tau^{boot} = 35 \text{ EUR/tCO}_2. \quad (3.7)$$

This bootstrap tariff is a temporary default T&S tariff used only before a meaningful shared T&S network can be evaluated. It allows CCS payoffs to be computed even when the early CCS coalition is too small to generate a reliable network-based tariff.

When network outputs are available, the model first computes the average T&S cost:

$$\bar{\tau}_t^{T\&S} = \frac{C_t^{op}}{Q_t^{CCS}}, \quad (3.8)$$

where C_t^{op} is the annual operator cost of the optimised network and Q_t^{CCS} is the captured CO₂ volume of the realised CCS coalition. Because costs are expressed in MEUR/year and volumes in MtCO₂/year, the ratio directly gives a value in EUR/tCO₂. This indicator is used directly in the average-price mode and also serves as a reference quantity in the flow-share tariff construction.

Endogenous average-price tariff mode

In the average-price mode, all CCS-feasible emitters face the same network-access tariff, equal to the latest average T&S cost:

$$\tau_{i,t}^{T\&S} = \bar{\tau}_t^{T\&S}, \quad \forall i \text{ technically feasible for CCS.} \quad (3.9)$$

This mode preserves the endogenous feedback between the realised CCS coalition, the optimised T&S network, total captured volume and average network cost, while removing heterogeneity in tariff allocation across clusters and emitters.

Flow-share tariff mode

The purpose of the flow-share tariff mode is to allocate network operator costs to clusters according to their use of the shared T&S infrastructure. For each active cluster c , the ABM computes the share of pipeline, booster and storage use attributed to that cluster. If $F_{p,t}$ is the total flow on active pipeline p and $F_{c,p,t}$ the flow attributed to cluster c , then:

$$\alpha_{c,p,t} = \frac{F_{c,p,t}}{F_{p,t}}. \quad (3.10)$$

An equivalent flow-share allocation is used for storage to compute $\alpha_{c,s,t}$. The T&S cost allocated to cluster c is therefore:

$$C_{c,t}^{T\&S} = \sum_{p \in \mathcal{E}_t} \alpha_{c,p,t} \left(C_{p,t}^{pipe} + C_{p,t}^{booster} \right) + \sum_{s \in \mathcal{S}_t} \alpha_{c,s,t} C_{s,t}^{storage}. \quad (3.11)$$

Here, $\mathcal{E}_t$ is the set of active network pipelines and $\mathcal{S}_t$ is the set of active storage sites. The allocation is constructed so that allocated cluster costs recover the total annual network operator cost. The realised cluster tariff is then:

$$\tau_{c,t}^{clus} = \frac{C_{c,t}^{T\&S}}{Q_{c,t}^{CCS}}, \quad Q_{c,t}^{CCS} > 0. \quad (3.12)$$

Once cluster tariffs have been constructed at the end of iteration t , the emitter-level tariff used in iteration $t + 1$ depends on the emitter's position relative to the last realised CCS network.

In iteration $t + 1$, a cluster is active if it contains positive injected CO₂ flow. It is treated as mature if its realised CCS volume exceeds 2 MtCO₂/year or if it contains at least three CCS emitters. These thresholds, the volume floor q^{floor} and the weighting parameter η are modelling parameters chosen to stabilise entrant-tariff calculations and distinguish « newborn » from more developed clusters.

1. **Emitter already in the realised CCS coalition:** If emitter i is already part of the realised CCS coalition, it receives the realised tariff of its cluster:

$$\tau_{i,t}^{T\&S} = \tau_{c(i),t}^{clus}. \quad (3.13)$$

2. **Potential entrant in an inactive cluster:** If the emitter is not yet in CCS and its cluster has no realised CCS volume, no cluster tariff exists. The model therefore uses an average-plus-access tariff:

$$\tau_{i,t}^{T\&S} = \bar{\tau}_t^{T\&S} + \frac{C_{i,t}^{access}}{q_i^{eff}}, \quad q_i^{eff} = \max\left(\frac{Q_i^{cap}}{10^6}, q^{floor}\right). \quad (3.14)$$

Here, $C_{i,t}^{access}$ is the annualised local access cost assigned to emitter i at iteration t , and q_i^{eff} is the effective captured volume used to express this cost per tonne. Since Q_i^{cap} is expressed in tCO₂/yr, the division by 10⁶ converts it into MtCO₂/yr. The volume floor q^{floor} avoids unstable per-tonne access costs for very small emitters. This access cost is defined as:

$$C_{i,t}^{access} = C_{i,t}^{loc,pipe} + C_{i,t}^{loc,booster} + C_{i,t}^{loc,fix}. \quad (3.15)$$

The first two terms approximate the local pipeline and booster costs needed to connect the inactive cluster to the active CCS backbone. The term $C_{i,t}^{loc,fix}$ represents a fixed local connection cost. In the current configuration, $q^{floor} = 0.01$ MtCO₂/year.

3. **Potential entrant in an active and mature cluster:** If the emitter is not yet in CCS but belongs to an active and mature cluster, the model assumes that local infrastructure is sufficiently developed for the entrant to access the network under conditions close to those of current CCS adopters:

$$\tau_{i,t}^{T\&S} = \tau_{c(i),t}^{clus}. \quad (3.16)$$

4. **Potential entrant in an active but immature cluster:** If the cluster is active but not yet mature, the realised cluster tariff exists but may still be unstable because it is based on limited captured volume or few adopters. The model therefore uses a hybrid tariff:

$$\tau_{i,t}^{T\&S} = \omega_{c,t} \tau_{c,t}^{clus} + (1 - \omega_{c,t}) \tau_{i,t}^{proxy}, \quad (3.17)$$

with

$$\tau_{i,t}^{proxy} = \bar{\tau}_t^{T\&S} + \frac{C_{i,t}^{loc,fix}}{q_i^{eff}}, \quad \omega_{c,t} = \frac{Q_{c,t}^{CCS}}{Q_{c,t}^{CCS} + \eta Q_{c,t}^{pot}}. \quad (3.18)$$

Here, $\tau_{c,t}^{clus}$ is the realised tariff of cluster c , while $\tau_{i,t}^{proxy}$ is the proxy entrant tariff used when the realised cluster tariff is still considered immature. It combines the average T&S cost with a local fixed access surcharge for emitter i , expressed per tonne using q_i^{eff} . The weight $\omega_{c,t}$ determines how much the entrant tariff relies on the realised cluster tariff rather than on the proxy tariff. It increases with the realised CCS volume $Q_{c,t}^{CCS}$ in the cluster and decreases with the remaining potential captured volume $Q_{c,t}^{pot}$. The parameter $\eta = 0.5$ controls how quickly the tariff moves from the proxy value towards the realised cluster tariff as the cluster develops.

At the end of the update, the ABM stores the average T&S cost, cluster tariffs, emitter tariffs, realised CCS coalition and network diagnostics. These values define the T&S information available in the next payoff evaluation.

Overall, this submodel translates realised T&S network outputs into emitter-facing cost signals. In the average-price mode, CCS attractiveness evolves through a uniform endogenous average tariff. In the flow-share mode, it can also vary with coalition size, cluster maturity, spatial access and realised network use.

3.2.4.4 CCU demand allocation and revenues

In the model, the difference between CCU and CCS lies in the downstream destination of the captured CO₂. In CCU, a part of the captured volume can be sold to geographically explicit CO₂ demanders, while any unsold volume is redirected to CCS as fallback storage.

For each emitter i , captured CO₂ is split into fossil and biogenic components:

$$Q_i^{cap} = Q_i^{fossil} + Q_i^{bio}. \quad (3.19)$$

This distinction is needed because fossil and biogenic CO₂ receive different carbon-accounting corrections when sold to non-permanent CCU demand. CO₂ demanders are represented as passive demand nodes with a fixed annual demand and a geographical location. Demanders do not have their own payoff function, and the model does not represent demander-specific production routes or willingness-to-pay.

In the current model, the CO₂ commodity price is treated as an exogenous price rather than as an endogenous market-clearing price. It is represented through a

biomass-substitution value from the demander perspective. One tonne of CO₂ is valued according to the cost of obtaining an equivalent amount of biogenic carbon from dry biomass. The commodity price is therefore computed as:

$$p^{CO_2} = \frac{p^{biomass}}{f_C^{biomass} \left(\frac{44}{12}\right)} = \frac{200}{0.5 \left(\frac{44}{12}\right)} \approx 109 \text{ EUR/tCO}_2. \quad (3.20)$$

Here, $f_C^{biomass} = 0.5$ tC/t dry biomass is used as a generic dry-biomass carbon fraction, consistent with default biomass carbon-accounting values in IPCC guidance [34]. The factor 44/12 converts carbon mass into CO₂ mass [35]. The biomass price $p^{biomass} = 200$ EUR/t dry biomass is used as a rounded central assumption.

For CCU, sold CO₂ is assumed to be delivered by truck separately to the utilisation site rather than transported through the CCS T&S network. The emitter-side delivery cost, expressed in EUR/tCO₂, is represented by the reduced-form function:

$$c_{i,d}^{del,E} = \frac{1}{2} (20 + 0.15d_{i,d}) = 10 + 0.075d_{i,d}, \quad (3.21)$$

where $d_{i,d}$ is the distance between emitter i and demander d in km. The factor 1/2 reflects the assumption that the delivery-chain cost is shared equally between the upstream emitter and the downstream CO₂ demander. Only the emitter-side share enters the CCU payoff, since demanders are passive agents. The construction and calibration of this reduced-form delivery-cost function are detailed in Appendix F.

Before allocating CCU volumes, the model forms a set of candidate transactions between available captured CO₂ blocks and all demanders. During payoff evaluation, this set includes the current CCU emitters and the revising emitter evaluating a switch to CCU. This allows the model to test whether the emitter could sell its captured CO₂ to one or several demanders, depending on distance, remaining demand and profitability. For each emitter-demander pair, the net revenue received by the emitter per tonne sold is:

$$NR_{i,d} = p^{CO_2} - c_{i,d}^{del,E}. \quad (3.22)$$

In the standard configuration, CCU demand is treated as non-permanent. Since the imported capture-route costs are built under a storage-oriented accounting logic, selling captured CO₂ to non-permanent CCU requires an accounting correction. For CO₂ type $k \in \{fossil, bio\}$, this correction is:

$$A_k = \begin{cases} p^{ETS}, & \text{if } k = fossil, \\ p^{CDR}, & \text{if } k = bio \text{ and CDR credits are active,} \\ 0, & \text{if } k = bio \text{ and CDR credits are not active.} \end{cases} \quad (3.23)$$

For fossil CO₂, the correction reflects the loss of the ETS benefit associated with permanent storage. For biogenic CO₂, it applies only in the CDR scenario, where non-permanent CCU means losing the negative-emission credit that would have been obtained under storage.

For each candidate transaction involving emitter i , CO₂ type k and demander d , the model then computes a marginal value to estimate the value of that transaction:

$$V_{i,k,d} = NR_{i,d} + \tau_i^{fb} - A_k, \quad (3.24)$$

where τ_i^{fb} is the fallback T&S tariff applied to unsold captured CO₂ redirected to CCS. It is taken from the current CCS T&S tariff signal available for emitter i , namely the bootstrap tariff, the uniform average tariff, or the emitter/cluster-specific tariff depending on the current network state and tariff-construction mode. This term appears because selling one tonne to CCU avoids sending that tonne to fallback T&S. A transaction is kept for CCU allocation only if:

$$V_{i,k,d} > 0. \quad (3.25)$$

Captured CO₂ is then allocated by matching the remaining supply of candidate CCU emitters to the remaining demand of CO₂ demanders. This matching is performed over all positive candidate transactions, ranked by decreasing marginal value. For each selected emitter-demander-CO₂-type transaction, the allocated quantity is:

$$q_{i,k,d} = \min(Q_{i,k}^{rem}, D_d^{rem}), \quad (3.26)$$

where $Q_{i,k}^{rem}$ is the remaining available supply of CO₂ type k (fossil or biogenic) from emitter i , and D_d^{rem} is the remaining demand of demander d . After each match, both remaining supply and remaining demand are updated. As a result, one emitter can sell to several demanders, and one demander can receive CO₂ from several emitters, until either the emitter's supply or the demander's demand is exhausted. This greedy matching is not a full market-optimisation problem, but it avoids allocating demand only according to the order in which emitters can revise their strategy.

For each emitter, the matching produces a sold quantity and a fallback quantity:

$$Q_i^{cap} = Q_i^{sold} + Q_i^{fb}. \quad (3.27)$$

The sold quantity may be distributed across several demanders and across fossil and biogenic CO₂ blocks. The total net CCU revenue of emitter i is therefore obtained by summing the net revenue of all transactions assigned to that emitter:

$$R_i^{CCU} = \sum_{d,k} q_{i,k,d} (p^{CO_2} - c_{i,d}^{del,E}). \quad (3.28)$$

The associated accounting adjustment is computed analogously:

$$A_i^{CCU} = \sum_{d,k} q_{i,k,d} A_k. \quad (3.29)$$

The annual private cost of the CCU strategy is then:

$$C_{i,t}^{CCU} = C_{i,t}^{route,CC} + Q_{i,t}^{fb} \tau_{i,t}^{fb} + A_{i,t}^{CCU} - R_{i,t}^{CCU}. \quad (3.30)$$

This expression shows that CCU uses the same capture-route cost as CCS, adds fallback T&S costs only for unsold captured CO₂, applies the accounting correction for non-permanent utilisation, and subtracts the net revenue from CO₂ sales.

The submodel is used differently during payoff evaluation and realised-state updating. During strategy revision, a hypothetical CCU matching is computed with the current CCU emitters plus the revising candidate. This is used only to evaluate the candidate payoff. After strategy revisions have been applied, the CCU market is cleared again for the realised set of CCU emitters, and the resulting sold quantities, fallback quantities, revenues and accounting adjustments are stored. If an emitter chooses CCU but sells no CO₂ in the realised matching (due to this market-clearing), it is reclassified to either CCS or No CC, depending on which alternative gives the higher payoff. This avoids treating CCU as adopted when no utilisation outlet is actually obtained.

Overall, the CCU submodel represents a finite and spatially distributed CO₂ utilisation market. It should be interpreted as a simplified spatial allocation mechanism rather than a full downstream market-clearing model: demanders are passive, the CO₂ commodity price is exogenous, delivery costs are represented through a reduced-form trucking function, and matching is greedy rather than optimised jointly.

3.2.4.5 Strategy revision rules

The strategy revision rule specifies the process by which emitter agents update their carbon-management strategies during the simulation. This rule connects payoff calculations to the evolution of adoption shares. At each iteration, each emitter is given a revision opportunity with probability $p_{revision}$. Emitters that do not receive this opportunity retain their current strategy for that iteration. This approach reflects the reality that industrial investment decisions are reconsidered at discrete intervals rather than continuously.

Each emitter i chooses only among its feasible strategy set $\mathcal{S}_i \subseteq \{No\ CC, CCS, CCU\}$. If no technically feasible capture route exists, CCS and CCU are removed from $\mathcal{S}_i$, because both require carbon capture. The revision rule is therefore applied only to feasible strategies. For each revising emitter, the model computes the payoff of all strategies in $\mathcal{S}_i$ and converts these payoff values into a new strategy using one of two revision modes: deterministic best response or stochastic logit choice.

The first revision mode is a deterministic best-response rule. Under this rule, a revising emitter selects simply the feasible strategy with the highest payoff:

$$s_i^{new} \in \arg \max_{s \in \mathcal{S}_i} \Pi_i(s), \quad (3.31)$$

where $\Pi_i(s)$ is the payoff of emitter i under strategy s . If several strategies have exactly the same highest payoff, one of them is selected randomly. This rule represents the closest model equivalent to strict individual cost optimisation: the revising emitter chooses the strategy with the highest computed payoff.

The second revision mode employs a stochastic logit rule, also known as a noisy best-response or quantal response rule. This approach relaxes the assumption that a revising emitter always selects the strategy with the highest computed payoff. Instead, strategies with higher payoffs are assigned higher choice probabilities, although lower-payoff strategies remain possible selections. This logic aligns with stochastic evolutionary dynamics, where logit choice rules model payoff-dependent errors or noisy best responses [7, 32]. It is also consistent with the quantal response framework of McKelvey and Palfrey, in which decision-makers respond probabilistically to payoff differences rather than always choosing the best response [44]. Within this thesis, this approach offers a more realistic behavioural representation than a purely deterministic rule, as emitters may encounter imperfect information, uncertain future costs and revenues, internal decision frictions, or bounded rationality.

In the model, the logit rule is implemented with an intensity-of-choice parameter β and a payoff scaling parameter Λ . For a revising emitter i , the probability of selecting strategy s is:

$$P_i(s) = \frac{\exp\left(\beta \frac{\Pi_i(s)}{\Lambda}\right)}{\sum_{r \in \mathcal{S}_i} \exp\left(\beta \frac{\Pi_i(r)}{\Lambda}\right)}, \quad s \in \mathcal{S}_i. \quad (3.32)$$

This equation has the same structure as the standard logit choice rule: the probability of choosing a strategy increases with its payoff, but all feasible strategies can still be selected with positive probability [32, 44]. The scaling parameter Λ is included because payoffs are expressed in EUR/year and can therefore be numerically large. In the implementation, $\Lambda = 10^6$, so that payoff differences enter the logit rule in million EUR/year.

The parameter β controls the strength of payoff sensitivity. When β is close to zero, payoff differences have little influence and choices become close to random. When β is large, the highest-payoff strategy becomes much more likely to be selected, and the rule approaches deterministic best-response behaviour. This corresponds to the usual interpretation of logit models: lower noise, or higher choice intensity, produces more payoff-sensitive behaviour, while higher noise produces more random choice [32, 44].

For the main logit simulations, the model uses $\beta = 5$ and $\Lambda = 10^6$. This value of β is not taken as a universal value from the literature, because the numerical meaning of the intensity parameter depends on the scale of the payoffs. The literature motivates the logit functional form and the interpretation of the noise or sensitivity parameter, but it does not prescribe a single value valid across applications. In this thesis, $\beta = 5$ is therefore treated as a calibrated behavioural sensitivity parameter. With the chosen scaling, a payoff advantage of EUR 1 million/year increases the odds of selecting one strategy over another by a factor of $\exp(5)$.

The logit implementation also includes a mutation probability μ . If mutation is active, a revising emitter enters a random-exploration step with probability μ and selects one strategy randomly from its feasible strategy set $\mathcal{S}_i$. With probability

$1 - \mu$, the emitter follows the logit rule described above. Mutation therefore represents occasional payoff-independent exploration, a common feature in stochastic evolutionary models where agents may sometimes deviate from payoff-based choice [32, 37]. In the main logit runs, $\mu = 0.001$, so mutation remains rare: 0.1% of revising emitters choose randomly, while the others follow the payoff-dependent rule.

Strategy updates are applied synchronously. During each iteration, revised strategies are stored temporarily and implemented only after all revising emitters have made their decisions. This procedure ensures that all revising emitters evaluate their feasible strategies under identical pre-revision system conditions. The updated strategy distribution becomes effective for the realised system update at the end of the iteration.

3.3 Model assumptions and limitations

The model is intended as a mechanism-oriented adoption model rather than a forecast of future CCS and CCU deployment. Its purpose is to analyse how decentralised emitter decisions, shared T&S infrastructure, carbon-accounting rules, CDR incentives, and CCU market assumptions create different adoption outcomes. The following assumptions are not exhaustive but focus on those most directly shaping payoff rankings, infrastructure feedbacks, CCU allocation, and result interpretation.

First, industrial technology choices are preprocessed prior to the start of the agent-based model (ABM). For each emitter, the input data specify a representative no-capture route and, where feasible, a representative capture route. The ABM thus compares No CC, CCS, and CCU based on these preprocessed routes, rather than dynamically re-optimising all possible industrial process choices during the simulation. This approach maintains model tractability while preserving a detailed techno-economic cost basis.

Second, carbon-accounting and policy incentives are embedded within predefined scenarios. The route-cost inputs incorporate the ETS, CDR, and accounting assumptions relevant to the selected scenario. The ABM does not simulate a dynamic carbon market or strategic anticipation of future policy changes. Therefore, results should be interpreted as adoption outcomes under fixed policy and accounting scenarios, rather than as outcomes of an endogenous policy or carbon-price trajectory.

Third, CCS infrastructure feedback is modelled through a passive regional CNO/CCNO layer. The model assumes a harmonised regional T&S framework for Belgium, the Netherlands, and North Rhine-Westphalia. It does not represent separate national network operators, distinct cross-border access rules, or separate national tariff layers for cross-border CO₂ transport and storage. This limitation is significant because T&S tariffs directly influence CCS payoffs, critical-mass thresholds, and spatial participation patterns. Consequently, the results describe adoption dynamics under a shared regional infrastructure-cost recovery logic.

Fourth, the T&S tariff construction is a simplified representation of network-cost

allocation. The CCNO evaluates the network associated with the realised CCS coalition, and the ABM translates these network outputs into tariff signals. For potential entrants, tariffs are determined by the latest realised network state, local access cost assumptions, and cluster-maturity rules, rather than by a full re-optimisation for every possible candidate coalition. The bootstrap tariff, mature-cluster thresholds, volume floor, and entrant-tariff weighting parameter are modelling assumptions introduced to stabilise tariff construction in early or sparse CCS coalitions. These assumptions are critical, as they influence the conditions under which CCS adoption becomes sufficiently attractive to reach critical mass.

Fifth, the CCU market is modelled as a simplified spatial allocation mechanism. CO₂ demanders are represented as passive external demand nodes with fixed annual demand and location. The CO₂ commodity price is exogenous, and matching is conducted using a greedy allocation rule rather than a comprehensive downstream market-clearing model. Additionally, captured CO₂ not sold to CCU demand is redirected to CCS for fallback storage. These assumptions are significant because the size, location, and value of CCU demand strongly influence whether CCU can compete with CCS or remains demand-constrained.

Sixth, the model simplifies CO₂ stream quality and CCU permanence. Fossil and biogenic CO₂ are distinguished for ETS and CDR accounting purposes, but CO₂ purity, impurities, stream-specific conditioning costs, and demander-specific quality requirements are not explicitly represented. In the standard configuration, CCU demand is treated as non-permanent, so permanent and non-permanent CCU pathways are not compared endogenously. This limitation primarily affects the interpretation of CCU results: the model captures spatial and economic competition for finite CCU demand, but not the full technical heterogeneity of downstream utilisation routes.

Finally, learning-by-doing is represented in a stylised manner. Learning effects are applied only to the capture-specific CAPEX component and are driven by cumulative capture adoption. No learning is applied to T&S infrastructure, storage, CCU logistics, or other site-level costs. Similarly, the CCNO model selects infrastructure from a predefined candidate network and storage representation. As a result, cost and spatial conclusions should be interpreted as conditional on the candidate network, cost assumptions, and learning specification employed in the model.

Collectively, these assumptions define the principal limitations of the model. While they do not invalidate the results, they delineate the scope of the model's claims. The model is well suited to analysing interactions among decentralised adoption, shared T&S infrastructure, CCS–CCU competition, and policy-driven payoff differences. However, it is less suited to predicting precise firm-level investment decisions, the institutional implementation of cross-border CO₂ infrastructure, or the comprehensive downstream market formation of CCU value chains.

Chapter 4

Results

This chapter is divided into two parts. Section 4.1 presents the results of eight different scenarios simulated under a single reference ETS price. In that first part, we examine the impact of carbon accounting, the tariff type for CO₂ transport and storage network usage, as well as the inclusion of CCU. Section 4.2 then tests the sensitivity of these adoption outcomes to the ETS price. The purpose of this sensitivity analysis is to examine whether the identified capture outcomes and infrastructure-threshold dynamics remain stable when the carbon-price incentive becomes weaker or stronger.

4.1 Base scenario results under the reference ETS price

The base-scenario analysis uses a reference ETS price of 150 EUR/tCO₂. All simulations in this section start from a zero-capture initial condition: no emitter is initially assigned to CCS or CCU. This starting point allows for analysing whether capture strategies can emerge without assuming an initial coalition of adopters. Furthermore, all these simulations use the stochastic logit revision mode.

The scenario matrix varies three binary dimensions: carbon accounting, T&S tariff design and CCU availability. In the CDR-credit cases, eligible negative-emission values are already embedded in the scenario-specific preprocessed emitter routes, while they are absent in the No-CDR cases. T&S costs are translated into emitter-facing tariffs either through the flowshare method or through average pricing, and CCU is either excluded or available as an additional capture strategy. Together, these settings define the $2 \times 2 \times 2$ base-scenario design summarised in Table 4.1.

TABLE 4.1: Overview of the eight base scenarios at the reference ETS price of 150 EUR/tCO₂.

Carbon accounting	Flowshare tariff		Average-price tariff	
	CCS only	CCS + CCU	CCS only	CCS + CCU
CDR credits	S1	S2	S3	S4
No CDR	S5	S6	S7	S8

4.1.1 Adoption pathways across the three dimensions of the base scenarios

This subsection first gives an overview of the adoption pathways generated by the eight base scenarios at the reference ETS price of 150 EUR/tCO₂. The purpose is to first identify the main outcomes produced by the model.

Figures 4.1 and 4.2 should be read together. The adoption trajectories, shown in Figure 4.1, indicate which strategies become dominant in the emitter population, expressed as the percentage share of emitters. Figure 4.2 shows the corresponding CO₂ flow trajectories by destination and carbon origin. Taken together, these two views are used below to distinguish four final outcomes.




FIGURE 4.1: Adoption-share trajectories across the eight base scenarios.

First, the CDR-credit scenarios, S1 to S4, converge to high-volume CCS-dominated outcomes. In these cases, final CCS adoption reaches about 63–73% of the emitter population. Geological storage reaches approximately 106.7–115.3 MtCO₂/yr at convergence and is strongly dominated by biogenic CO₂, with final biogenic shares around 78.5–80.5%. The associated average T&S costs are low, between 19.0 and 23.5 EUR/tCO₂. This combination of high stored volumes and low T&S costs suggests that CDR credits help support a mature CO₂ transport and storage value chain, in which large biogenic CO₂ volumes enable economies of scale in the shared T&S network.

Second, the No-CDR scenarios, S5 and S6, form lower-volume fossil CCS-dominated outcomes. In adoption-share terms, these scenarios may initially look relatively close to the CDR-credit scenarios: final CCS adoption still reaches about 54–57%. However, final geological storage is only about 36.5 MtCO₂/yr, roughly one third of the storage volume observed in S1–S4, and the stored volume is entirely fossil. This reflects the absence of a reward for biogenic storage in the No-CDR setting: without CDR credits, bio-CCS routes are not selected as attractive first-best capture routes, so CCS emergence is driven by fossil capture options instead.¹ S5 and S6 therefore

¹This is linked to the scenario-specific input datasheets and preprocessing step. Each emitter

4.1. Base scenario results under the reference ETS price




FIGURE 4.2: CO₂ flow trajectories by destination and carbon origin across the eight base scenarios.

still show CCS diffusion, but through a lower-volume fossil storage pathway rather than the high-volume biogenic pathway observed under CDR credits.

Third, scenario S7 corresponds to a near-complete no-capture lock-in. At convergence, 98.7% of emitters remain in No CC, while only 1.3% are still assigned to CCS, so storage is practically zero. This small residual CCS share should not be interpreted as the emergence of a viable storage pathway. It results from a few early CCS entries enabled by the temporary bootstrap tariff used before a T&S network exists. Once the network is updated and actual T&S costs are internalised, these marginal CCS positions are no longer attractive; they simply remain in the final state because convergence is reached before those emitters revise again. S7 therefore represents a failed CCS take-off, with no mature shared T&S network emerging.

Fourth, the No-CDR, average-price scenario with CCU, S8, forms a CCU-dominated capture outcome. CCS remains marginal: final CCS adoption reaches only 4.0%, direct CCS storage is about 1.7 MtCO₂/yr, and total geological storage including fallback is about 2.0 MtCO₂/yr. The average T&S cost remains outside the low-cost region, indicating that the CCS infrastructure feedback does not activate. By contrast, 45.4% of emitters adopt CCU and 28.9 MtCO₂/yr are delivered to utilisation demand, essentially exhausting the available modelled demand. S8 therefore shows that, when storage does not become self-sustaining, decentralised capture adoption can still emerge through utilisation. In this case, utilisation acts as a substitute downstream outlet for captured CO₂, avoiding dependence on shared T&S scale economies and relying instead on accessible demand.

These four outcomes demonstrate that similar levels of CCS adoption can correspond to significantly different stored volumes, carbon origins, and infrastructure states. The following subsection examines whether these outcomes remain stable attractors when the initial distribution of strategies is varied.

is represented by two selected route options: the first-best route without carbon capture and the first-best route with carbon capture under the scenario-specific accounting assumptions.

4.1.2 Identifying evolutionary stable strategies

The previous subsection classified outcomes obtained from a zero-capture initial condition. This is useful for studying endogenous emergence. However, this approach does not indicate whether these outcomes are robust to different starting points. Additional simulations are therefore run from different initial distributions of No CC, CCS, and CCU. The objective is to distinguish outcomes that behave as robust attractors from those that depend on the initial coalition of capture adopters. In this thesis, evolutionarily stable strategies (ESS) are defined as stable long-run strategy-distribution attractors in the ABM-EGT simulation. An ESS is considered distinct if the final states form a clearly separated region in the simplex and remain robust to seed checks.

Figure 4.3 first shows the CCS-only scenarios. Since only No CC and CCS are available, the simplex reduces to a one-dimensional line between 100% No CC and 100% CCS initial adoption. S1 and S3 both converge towards a single high-volume CCS stationary outcome. This indicates one robust stationary attractor, and therefore outcome, in each case. The final state is largely independent of the initial No CC–CCS distribution. The slightly lower final point visible in S1 is not counted as a distinct stationary outcome, because seed checks show that it does not persist as a separate basin.




FIGURE 4.3: Basin-of-attraction maps for the CCS-only scenarios on a one-dimensional simplex; Appendix G explains how to read this type of plot. Grey arrows show initial-to-final CCS-share movements.

The No-CDR CCS-only scenarios are more sensitive to the initial CCS coalition. In S5, all final states remain CCS-dominated, but the line separates into two storage-volume basins: one lower-volume CCS outcome and one higher-volume CCS outcome. The difference is therefore not the dominant strategy itself, but the amount of CCS storage supported by the final coalition. The isolated point between

these two regions is not counted as an additional stationary outcome because seed checks show that it is not robust. With a different stochastic seed, the corresponding initial condition converges either to the lower-volume or higher-volume basin.

S7 shows an even clearer separation between two attractors. Low initial CCS shares lead back to a low-capture stationary outcome close to the no-capture corner, whereas higher initial CCS shares can sustain a more CCS-oriented outcome. Under average pricing and without CDR credits, the final state is therefore path-dependent: the system's long-run outcome depends on whether the initial CCS coalition is large enough to keep the storage pathway viable.

Figure 4.4 extends the analysis to the CCU-enabled scenarios. With three available strategies, initial conditions are represented on a triangular simplex. These simulations test whether CCU availability creates a separate stationary outcome with persistent utilisation, or whether final states remain storage-dominated.




FIGURE 4.4: Basin-of-attraction maps for the CCU-enabled scenarios on a triangular simplex; Appendix G explains how to read this type of plot. Grey arrows show initial-to-final movements, while red lines highlight example trajectories.

In S2 and S4, final states converge towards high-volume CCS outcomes. CCU

can appear during the transition, but it does not form a separate final basin under CDR credits. S6 behaves similarly at a smaller storage scale: final states remain concentrated around a lower-volume CCS outcome, with only residual CCU. These three scenarios therefore, do not show clear path dependence at the resolution of the initial-condition grid used here.

S8, the No-CDR scenario with average pricing and CCU, is the main exception. The simplex shows two clearly separated attraction regions, divided by a visible separatrix between the 40% and 50% initial CCS-share lines. One basin leads to the CCU-dominated stationary outcome observed in the base run starting from zero capture. The other basin is storage-oriented and is reached only when the initial condition contains enough CCS to support the storage pathway. S8 therefore shows the clearest case of two distinct ESS-like stationary outcomes: one centred on utilisation and one centred on storage.

Overall, the equilibrium stability analysis shows that, under the reference ETS price, the scenarios differ in their equilibrium structure. With CDR credits, the high-volume CCS outcome is generally a unique and robust attractor. Without CDR credits, the system becomes more sensitive to the initial coalition: S5 and S7 split into distinct stationary outcomes, while S8 displays a genuine separatrix between CCU-dominated and storage-oriented basins.

4.1.3 Critical transition dynamics: tipping points, lock-in and critical mass

After studying the main adoption pathways and their dependence on initial conditions, this subsection zooms into the transition process itself. This transition-level analysis exploits the strength of the ABM-EGT framework. The model records the sequence of individual revision events through which temporary CCS entries, failed attempts and coalition build-up appear before convergence. This section investigates how the system moves from no-capture lock-in towards stable CCS adoption, and how tipping points and critical-mass effects appear in the trajectories. Scenarios S1 and S5 are the clearest pair for illustration purposes. Both are CCS-only and use the flowshare-based T&S tariff, starting from zero capture. The only difference is the presence of CDR credits.

Figure 4.5 compares the dynamic evolution of S1 and S5. For each scenario, the figure is organised into three stacked panels. The upper panel tracks the average T&S cost together with the range of emitter-specific tariffs, the middle panel shows the adoption shares of No CC and CCS, and the lower panel reports the corresponding CO₂ flows by route and carbon origin.

A common three-phase pattern can be identified in both scenarios. The red markers along the x-axes of the panels in Figure 4.5 indicate the phase intervals discussed below.

Phase 1 – Early stage/origin: At the beginning of both simulations, CCS remains marginal. In the upper panels of Figure 4.5, this phase is characterised by very high and unstable T&S costs, together with a wide dispersion of emitter-specific

4.1. Base scenario results under the reference ETS price




FIGURE 4.5: Dynamic comparison of S1 and S5 over the number of iterations.

tariffs. In the middle panels, the system remains dominated by No CC adoption. This corresponds to a weak-infrastructure regime: CCS exists as an option, but it is not yet supported by a sufficiently used T&S system. The high costs in the upper panels reflect this under-utilisation: the network cost is spread over too little captured CO₂ for CCS to become self-reinforcing.

Phase 2 – Development/diffusion stage: The second phase corresponds to the period during which the system moves out of the low-adoption phase. This transition is visible through three simultaneous changes: CCS adoption increases, captured CCS volumes rise sharply, and average T&S costs fall from very high and unstable values to a much lower range. In the illustrated run, this break occurs after fewer stochastic revision steps in S1 than in S5. This precise timing is specific to the stochastic revision sequence generated by the selected seed. Other seeds can shift the break earlier or later, as shown in Appendix H. The comparison therefore concerns the number of revision opportunities needed to assemble a viable CCS coalition. In S1, this coalition forms more easily than in S5, for the carbon-accounting reasons explained below.

Phase 3 – Maturity/saturation stage: After the transition, both scenarios reach a more stable regime. Average T&S costs remain low, tariff dispersion is reduced, and CCS becomes the dominant strategy. However, the transition remains partial. There is a significant proportion of emitters that remain in No CC. The final state is therefore more accurately described as a mature partial CCS outcome rather than full CCS adoption. The two scenarios differ in the scale and composition of this outcome: S1 achieves a much larger, biogenic-dominated stored volume, whereas S5 stabilises at a smaller, entirely fossil CCS volume.

Carbon-accounting interpretation of the transition

The difference between S1 and S5 starts before the iterative adoption process. For each scenario, the preprocessing step selects a no-capture route and, when feasible, a capture route for each emitter under the scenario-specific accounting assumptions. The capture options passed to the ABM can therefore differ between the CDR and No-CDR runs, both in captured volume and in fossil-biogenic composition.

In S1, CDR credits increase the value of storing biogenic CO₂. This makes several large biogenic capture routes economically more favourable from the beginning of the simulation, even while the shared T&S network is still poorly used and expensive. When one of these emitters receives a revision opportunity and enters CCS, it can add a large captured volume to the coalition. This directly helps lower the average T&S cost faced in later revisions and makes subsequent CCS entries easier to sustain.

In S5, the additional value of biogenic storage is absent. The early CCS build-up has to come from capture routes that are already attractive under the No-CDR accounting setting. This changes the nature of the first successful entries. Instead of being supported by an additional value for biogenic storage, CCS adoption depends on fossil routes for which avoided ETS costs are high enough to offset capture costs and initially high T&S costs. Such favourable opportunities exist, as shown by the later transition in S5. But they are less readily available during the early revision sequence. The stochastic process therefore needs more attempts before it generates an entry sequence that moves the system out of the low-adoption regime.

The transition moments observed in S1 and S5 remain seed-dependent, since the order of revision opportunities and logit choices changes across stochastic runs. This is shown in Appendix H. The transition can occur earlier or later depending on the seed, but the S5 trajectories generally shift later than the S1 trajectories. The important comparison is therefore the relative difficulty of reaching the transition. S1 tends to require fewer favourable revision events because large biogenic storage routes are already made attractive by the CDR-credit setting. S5 generally needs a more specific sequence of fossil entrants before the CCS coalition can reduce T&S costs sufficiently. The diagnostics below explains this mechanism by showing which emitters, volumes and cost changes correspond to the transition in the illustrated runs.

Transition diagnostics and critical-mass mechanism

Table 4.2 links the three-phase pattern in Figure 4.5 to the main numerical changes observed in S1 and S5. For each scenario, the first iteration interval captures the critical transition from Phase 1 to Phase 2. This is the tipping-point window: CCS moves from a marginal position to a state where captured volumes rise sharply and average T&S costs fall strongly. The second interval follows this initial break and captures the build-up from Phase 2 towards Phase 3, where CCS adoption and stored volumes continue to increase until the system approaches the mature partial-CCS regime described above. This table should be read together with the adoption-share, CO₂-flow and T&S costs panels in Figure 4.5.

TABLE 4.2: Key diagnostics of the CCS critical-transition sequence in S1 and S5.

Run	Iteration segment	CCS adopters	CCS share [%]	Captured volume [MtCO ₂ /yr]	Avg. T&S cost [EUR/tCO ₂]
S1	30–38	2 → 22	0.9 → 9.7	0.02 → 31.78	3198.5 → 24.3
S1	38–100	22 → 136	9.7 → 59.9	31.78 → 106.76	24.3 → 18.6
S5	205–241	2 → 54	0.9 → 23.8	0.02 → 20.34	8842.5 → 26.4
S5	241–380	54 → 130	23.8 → 57.3	20.34 → 36.56	26.4 → 22.6

In S1, the first interval shows that the tipping point occurs before CCS becomes widespread in terms of adoption share. Between iterations 30 and 38, the CCS share rises from 0.9% to 9.7%, while captured CCS volume increases from almost zero to 31.78 MtCO₂/yr and the average T&S cost falls from 3198.5 to 24.3 EUR/tCO₂. This break is driven by a few high-impact entrants, including a biogenic refinery emitter of about 7.51 MtCO₂/yr, a fossil steel emitter of about 6.69 MtCO₂/yr, and another biogenic refinery emitter of about 4.51 MtCO₂/yr. The key point is therefore not the number of adopters alone, but whether the entering coalition brings enough volume, in favourable locations, to change shared T&S utilisation.

The early trajectory of S5 clarifies this distinction. Before the main transition, several temporary CCS increases appear in the adoption-share panel, but the corresponding CO₂ flows remain too small to become tipping points. Around iteration 92, for example, CCS adoption briefly rises, yet the associated captured volume is only about 0.56 MtCO₂/yr. The largest entrant, in those bumps, is a cement emitter of about 0.42 MtCO₂/yr, which faces a T&S tariff above 700 EUR/tCO₂ after the realised network update and returns to No CC. These early bumps are therefore attempted CCS entries, but they do not create the self-reinforcing infrastructure regime needed for CCS to remain attractive.

The successful transition in S5 occurs only later, when the entering coalition has a different composition. During the first transition interval in S5, several larger fossil-industrial emitters enter CCS, including steel emitters of about 3.99 and 3.11 MtCO₂/yr, chemical emitters above 1 MtCO₂/yr, and refinery emitters close to or above 1 MtCO₂/yr. By iteration 241, the CCS coalition reaches 20.34 MtCO₂/yr and the average T&S cost falls to 26.4 EUR/tCO₂. Compared with the earlier failed attempts, the successful transition involves emitters that add more captured volume and improve the use of the shared T&S network.

This comparison demonstrates that critical mass is not defined by a fixed adoption-share threshold. Temporary CCS entries may fail if they add too little useful volume to the network. Conversely, a relatively small share of emitters can trigger the tipping point if their captured volume is sufficient to lower T&S costs and stabilise further CCS adoption. The tipping point is therefore best understood first as a volume-weighted infrastructure threshold. The spatial structure of this threshold is analysed next through the cluster-level and network evidence.

Spatial location of main clusters

Figures 4.6 and 4.7 provide the cluster-level and spatial evidence behind this mechanism. Figure 4.6 decomposes captured CCS volume by cluster during the transition window. Here, clusters refer to groups of nearby emitters that share similar access conditions to the candidate T&S network. They are used as spatial groupings to identify where captured volumes enter the network during the transition.




FIGURE 4.6: Cluster-level build-up of captured CCS volume in S1 & S5.

The increase in captured CCS volume occurs through discrete cluster-level steps rather than a smooth system-wide rise. This confirms that the transition is driven by specific groups of emitters. The ABM-EGT output is valuable in this context because it links the aggregate tipping point to the underlying sequence of emitter and cluster entries. In S1, the transition and subsequent build-up involve several strategic clusters, including C20, C22, C424, C92, and C88. In S5, the same mechanism appears at a smaller scale and is entirely fossil-led. The early bumps fail because the entering clusters do not add enough useful volume in a network configuration that sufficiently improves T&S cost conditions. The latter transition occurs when larger fossil-industrial emitters enter clusters capable of supporting a viable T&S network.

Figure 4.7 shows the spatial dimension of this mechanism. Strategic clusters combine large potential captured volumes, concentrated CCS demand and favourable access to the candidate T&S network. Their role is to add CO₂ volume at locations that improve the utilisation of shared T&S assets and help unlock economies of scale for later adopters.

The cluster evidence therefore adds the spatial explanation to the volume-weighted threshold identified above. Overall, the tipping point observed in S1 and S5 is best interpreted as a volume-weighted and spatially structured infrastructure threshold. Small CCS coalitions can appear temporarily, but they fail if they do not reduce T&S costs durably. A true transition occurs only when high-volume and spatially coherent emitters enter CCS and create conditions under which additional adopters can remain in the coalition. The following subsections build on this mechanism to analyse the impact tariff design and CCU availability have on these findings.




FIGURE 4.7: Spatial location of the main clusters involved in the critical transition, indicated on the final T&S networks of S1 on the left and S5 on the right.

4.1.4 Tariff allocation and CCS coalition formation

The tariff-design dimension is important because T&S infrastructure costs are not only system-level costs. In practice, they must be translated into charges paid by network users. The ABM-EGT framework makes this cost-allocation step explicit. After each realised CCS coalition is evaluated by the passive CNO, the resulting annual T&S operator cost is converted into emitter-facing tariffs. These tariffs then enter the CCS payoff and influence later strategy revisions. This allows the model to test how different tariff-allocation rules affect decentralised CCS adoption, which is an important advantage of representing adoption decisions at the emitter level.

In both tariff modes, the total annual T&S operator cost returned by the CNO is recovered through tariffs. The difference between them lies in how this cost is distributed. Under the average-price tariff, the total cost is spread uniformly over the CCS coalition. Under the flowshare tariff, costs are allocated more locally according to cluster-level use of pipelines, storage and access infrastructure.

As can be seen in Figure 4.1 and Figure 4.2, under CDR credits, both tariff modes lead to high-volume CCS-dominated outcomes: final CCS adoption remains in the 63–73% range and stored volumes exceed 100 MtCO₂/yr. However, the spatial composition of the CCS coalition differs. This is visible when comparing S1 and S3. In the flowshare case, emitters located far from the main T&S backbone, or in smaller peripheral clusters, are less likely to remain in CCS because they internalise more of the connection and network-use burden they create. In the average-price case, this burden is allocated across the whole CCS coalition. As a result, some peripheral emitters become able to adopt CCS because their individual tariff no longer reflects their full local infrastructure impact. This spatial difference is illustrated in Figure 4.8, which compares the realised T&S networks in S1 and S3.

The flowshare tariff acts as a spatial selection mechanism: it favours emitters whose captured volumes and locations make efficient use of the shared network. The




FIGURE 4.8: Realised T&S networks in S1 (left) and S3 (right).

average-price tariff weakens this selection by socialising connection costs across all CCS adopters. Remote emitters can then distribute part of their high connection and network-use costs across the whole CCS coalition, including central clusters that use the network more efficiently.

The same mechanism has a different consequence when CDR credits are removed. Under the flowshare tariff, S5 and S6 still activate a CCS regime, although only at a lower storage scale of about 36.5 MtCO₂/yr. This suggests that favourable early adopters can remain in CCS long enough for captured volume to accumulate and for the infrastructure feedback to become self-reinforcing, leading to the critical-mass effect. Under the average-price tariff, by contrast, this spatial tariff signal is removed. Early adopters no longer benefit from being well located, because they face the same average tariff as all CCS-feasible emitters. In a weaker No-CDR setting, where CCS payoffs are already less favourable, this uniform tariff can make the initial CCS coalition too fragile to survive the high-cost phase. This explains why S7 remains locked in quasi-total non-adoption without reaching the tipping point needed for a mature CCS pathway, and why S8 develops a CCU-dominated capture outcome instead of a mature CCS-dominated outcome.

4.1.5 CCU as a conditional alternative capture route

Across the scenarios with utilisation available, CCU does not play a fixed role at the reference ETS price. Its function depends on whether the CCS pathway activates a self-reinforcing infrastructure dynamic and on the relative downstream costs of the two capture routes. CCS depends on the endogenous T&S tariff, while CCU depends on distance-based delivery costs to finite CO₂ demand nodes. Figure 4.9 compares the four CCU-enabled scenarios and shows that CCU can appear as a temporary dominant strategy, remain residual, or become the dominant capture route.

In the CDR scenarios S2 and S4, CCU is not part of the final utilisation outcome, but it can play an important transitional role. In S2, CCU appears early and then

rapidly disappears as CCS captured volume increases and the T&S cost falls. In S4, the transient CCU phase is much more pronounced: utilisation temporarily absorbs almost the total available CO₂ demand while the average-price CCS pathway is still immature. This indicates that CCU is competitive mainly while the CCS/T&S route still faces immature and unstable network costs. Once the tipping point for CCS adoption is reached, the average T&S cost falls into the low-cost range observed in the CDR scenarios, and storage becomes the dominant route. Final CCU sales then fall to zero, while direct CCS storage stabilises at about 113.3 MtCO₂/yr in S2 and 106.7 MtCO₂/yr in S4.




FIGURE 4.9: Conditional role of CCU across the four CCU-enabled scenarios.

The No-CDR cases show that utilisation becomes more important precisely when the CCS route is harder to stabilise. In S6, the flowshare tariff still allows a CCS coalition to survive and eventually activate the infrastructure feedback. This happens at a much lower scale than in the CDR cases: CCS reaches about 36.5 MtCO₂/yr, while CCU remains residual at about 3.1 MtCO₂/yr sold. The contrast with S8 can be read as a relative-cost result. In S8, the combination of No CDR credits and the average-price tariff prevents CCS from crossing the low-cost infrastructure threshold. CCS storage remains therefore limited to about 1.7 MtCO₂/yr. Utilisation deliveries remain attractive for emitters that can access demand nodes at sufficiently low delivery cost, so CCU becomes the main capture route, with 45.4% final adoption and 28.9 MtCO₂/yr delivered to utilisation demand.

The failure of CCS emergence does not necessarily imply the failure of capture adoption. When the CCS/T&S route does not achieve a low-cost regime, utilisation can serve as an alternative outlet if delivery costs to demand nodes remain sufficiently low. CCU operates under a different downstream-cost logic than CCS: it is constrained by spatial matching and finite demand, whereas CCS depends on whether the shared T&S tariff falls with scale. A supplementary flow map for S8 is provided in Appendix I, in which the flows between emitters and demanders are

shown. Overall, CCU could be interpreted as conditional on the CCS infrastructure pathway: when the CCS feedback activates, CCU becomes residual; when the feedback fails, it can support a utilisation-centred capture outcome.

4.1.6 Decentralised ABM-EGT outcomes versus the centralised CCNO reference

This subsection compares the decentralised ABM-EGT model coupled to the CCNO with the standalone CCNO model. The comparison is restricted to the CCS-only cases at the reference ETS price of 150 EUR/tCO₂. This avoids mixing the CCNO comparison with CCU route competition, which is only represented in the ABM-EGT framework. The comparison therefore focuses on the CCS-only outcomes under the CDR and No-CDR accounting settings.

Before comparing the results, it is useful to clarify what the standalone CCNO model does. The standalone CCNO model is a centralised techno-economic optimisation model. It decides which emitters adopt a CCS route, which storage sites are opened, and which pipelines are built and sized. The model takes the same industrial emitter data as the ABM, together with candidate CO₂ transport pipelines and storage options. Its goal is to minimise the total system cost. In this model, emitter adoption is internalised by the optimisation. An emitter is connected when the benefit of switching from its No-CC route to its CC route is sufficient to justify the capture, transport and storage costs created by its inclusion into CCS.

This differs from the ABM-EGT-CCNO model. In the coupled model, emitters do not enter CCS through a central planner/operator. They revise their strategy in a decentralised way according to their private payoff. The ABM then sends to the CCNO model only the CCS coalition that emerges from these decentralised adoption decisions. The CCNO model is therefore used inside the simulation as a network-evaluation layer. It computes the T&S network and costs for the realised CCS coalition, after which these costs are translated into emitter-facing tariff signals for later revisions. The comparison is therefore between a coordinated least-cost reference and adoption pathways generated by decentralised stochastic decisions.

TABLE 4.3: Comparison between the standalone CCNO model and selected ABM-EGT-CCNO outcomes at ETS 150 EUR/tCO₂.

Indicator	CCNO CDR	ABM S1	ABM S3	CCNO No-CDR	ABM S5	ABM S7
CCS adopters	168	143	165	147	130	3
CCS adopters [%]	74.0	63.0	72.7	64.8	57.3	1.3
CCS stored [MtCO ₂ /yr]	106.6	108.3	115.3	37.0	36.5	≈ 0
Average T&S cost [EUR/tCO ₂]	19.8	19.0	23.5	28.3	22.7	/
Total system cost [MEUR/yr]	84 236.58	84 538.66	84 585.76	89 053.45	89 180.34	91 600.49
Gap vs. optimum [MEUR/yr]	0.00	+302.08	+349.19	0.00	+126.85	+2 547.00

Table 4.3 shows that the decentralised ABM-EGT-CCNO model can reach storage volumes close to the standalone CCNO reference, but through a different selection process. In the CDR case, S1 stores 108.3 MtCO₂/yr with 143 CCS adopters, while the standalone CCNO reference stores 106.6 MtCO₂/yr with 168 connected

emitters. S3 also reaches a high-volume CCS outcome, with 115.3 MtCO₂/yr stored and 165 CCS adopters. In the No-CDR case, S5 stores 36.5 MtCO₂/yr with 130 CCS adopters, compared with 37.0 MtCO₂/yr and 147 connected emitters in the standalone CCNO reference. S7, in contrast, remains close to no-capture lock-in and does not form a mature CCS network.

The comparison of these results is important. As explained above, the standalone CCNO model chooses emitters, pipelines and storage jointly in one coordinated optimisation. The ABM-EGT-CCNO model only sends to the CCNO the coalition that emerges from decentralised adoption. The final storage volume can therefore be similar while the set of connected emitters and the realised network differ. In both accounting cases where a mature CCS pathway emerges, the ABM outcome is more selective in terms of connected emitters than the standalone benchmark.

The total system-cost comparison reinforces this interpretation and makes the tariff-design effect explicit. In the CDR case, both decentralised outcomes remain close to the centralised reference, with S1 slightly closer than S3. The difference is much stronger without CDR credits: S5 remains close to the CCNO reference, with a gap of only 126.85 MEUR/yr, whereas S7 remains locked in and leaves a gap of 2 547.00 MEUR/yr. This suggests that uniform T&S pricing can weaken the spatial cost signal needed for efficient decentralised CCS coalition formation.

Figure 4.10 confirms this difference spatially for the CDR setting. The standalone CCNO reference builds a coordinated network for the emitters selected by the optimisation. The ABM-EGT-CCNO maps instead reflect the coalitions that «survive» the decentralised revision process. S1 and S3 both reach high-volume CCS outcomes, but their realised network structures differ because the emitters face different tariff-allocation rules during the ABM process. Larger versions of the main CCS network maps, including the No-CDR reference and S5 outcome, are provided in Appendix J.




FIGURE 4.10: Spatial comparison between the standalone CCNO CDR reference and the ABM-EGT-CCNO outcomes S1 and S3 at ETS 150 EUR/tCO₂.

Overall, this comparison highlights the distinct purposes of the two models. The standalone CCNO model identifies outcomes when emitter adoption, network construction, and storage use are jointly optimised. In contrast, the ABM-EGT-CCNO

model examines whether a comparable outcome can emerge from decentralised adoption decisions. The value of the ABM-EGT approach thus provides an alternative perspective on the same techno-economic system. Instead of assuming that emitter connection decisions are internalised by a central planner/operator, it asks what system outcome can emerge when decision power is distributed across industrial emitters themselves. This makes it possible to study whether a coordinated storage opportunity identified by the CCNO model can also be achieved through decentralised, stochastic, and payoff-driven adoption decisions. In addition, it allows the study of conditions under which the decentralised process diverges from the centralised model.

4.2 ETS price sensitivity and equilibrium landscape

The previous section analysed the reference ETS price of 150 EUR/tCO₂. This section examines whether the adoption pathways and stationary outcomes identified previously remain consistent under varying ETS prices. Rather than repeating the full eight-scenario matrix, the sensitivity analysis focuses on the flowshare-tariff and CCS+CCU configuration, which represents the most comprehensive decentralised setting. In this configuration, emitters face differentiated T&S tariffs, CCS can generate infrastructure feedbacks, and CCU competes as an alternative capture route. The ETS price is varied between 75, 150, and 300 EUR/tCO₂.

Two complementary questions are studied. First, from a zero-capture initial condition, does a higher ETS price make it easier for the system to escape low adoption and activate the CCS infrastructure feedback? Second, when initial conditions are varied, does the ETS price change the number and type of stationary outcomes reached by the evolutionary process?

4.2.1 Adoption dynamics from a zero-capture initial condition

This subsection analyses ETS-price sensitivity using the same zero-capture initial condition as in the base-scenario analysis. The scenario corresponds to S2, which includes CDR credits, flowshare-based T&S tariffs, and CCU availability. The only parameter varied is the ETS price. The objective is to compare final adoption shares and to assess how the ETS price influences the conditions under which CCS reaches critical mass or CCU becomes the primary capture route.

Figure 4.11 demonstrates that the ETS price shifts the dynamic pathway followed by the system. At 75 EUR/tCO₂, the ETS-driven willingness to pay for avoided emissions is insufficient to activate the shared T&S infrastructure. CCS adoption remains marginal, stored volumes are low, and the average T&S cost never enters the low-cost region. Nevertheless, the system does not remain in a complete no-capture lock-in. CCU emerges as the dominant capture route, achieving approximately 62% of final adoption and nearly fulfilling the total available utilisation demand, around 29 MtCO₂/yr. This high site-level adoption share aligns with the limited CCU volume because many CCU adopters are small emitters; the median CCU adopter sells only about 19 ktCO₂/yr.

4.2. ETS price sensitivity and equilibrium landscape




FIGURE 4.11: The upper row reports adoption shares, the middle row reports CO₂ flows by route, and the lower row reports the average T&S cost.

As discussed earlier, at the ETS price of 150 EUR/tCO₂, the model produces a different trajectory. CCU appears during the early revision sequence, but it remains temporary. Once enough captured volumes, in favourable locations enters CCS, the average T&S cost falls and the system crosses the critical infrastructure threshold described in Section 4.1.3. CCS then becomes self-reinforcing and stabilises at about 66% adoption, with roughly 113 MtCO₂/yr of direct geological storage. CCU sales fall to zero at convergence, showing that utilisation can temporarily absorb captured CO₂ before being displaced once the CCS infrastructure feedback becomes mature.

At 300 EUR/tCO₂, the low-adoption phase is strongly compressed in iteration number. CCS adoption rises quickly, the average T&S cost enters the low-cost region early, and the system converges toward a high-volume CCS outcome. Final CCS adoption reaches approximately 73%, with stored volume rising to around 126 MtCO₂/yr, while CCU remains negligible. The higher ETS price, therefore, makes the critical-mass condition easier to assemble: more emitters find capture privately attractive during early revision opportunities, increasing the likelihood of forming a sufficiently robust CCS coalition.

Overall, the ETS price influences not only the final level of capture adoption but also the transition mechanism. At low ETS price, CCS fails to reach critical mass and CCU becomes the primary capture strategy. At the reference price, the system undergoes a temporary CCU phase before shifting toward CCS. At high ETS price, CCS becomes sufficiently attractive for rapid activation of the infrastructure feedback, leaving minimal opportunity for CCU to persist. Thus, the ETS price functions as a threshold-shifting parameter in the decentralised adoption process.

4.2.2 Number and type of stationary outcomes

The second part of the ETS sensitivity analysis examines whether the zero-capture trajectories are robust to different initial conditions. The analysis focuses on scenarios S2 and S6, both of which include CCU and employ the flowshare tariff, but differ in the presence of CDR credits. This approach maintains the equilibrium-landscape analysis on the most decentralised tariff representation while enabling assessment of the ETS price with and without the additional activation effect of CDR credits.

According to the definition introduced in Section 4.1.2, final states are considered stationary outcomes when they form stable attractors within the finite stochastic ABM-EGT process. The objective is to determine the number of such attractors in each ETS-price setting and to classify their type as CCS-dominated, CCU-including, or mixed.

Figure 4.12 shows that the ETS price changes the type of dominant outcome more than it creates many competing basins. Within the tested initial-condition grid, each panel mainly converges towards one compact final-state region. The equilibrium landscape is therefore not characterised by multiple separated attractors, as in S8 in the base-scenario analysis. Instead, the ETS price shifts the dominant stationary outcome from a CCU-including capture pathway at low ETS price towards CCS-dominated outcomes at higher ETS prices.

At 75 EUR/tCO₂, both S2 and S6 converge towards final states with limited CCS and a huge CCU component. This is consistent with the zero-capture trajectory in Figure 4.11: the ETS price is too low for CCS to reliably activate the shared T&S infrastructure feedback, but high enough for CCU demand to support capture adoption. The dominant ESS-like outcome is therefore not a mature CCS outcome, but a CCU-including capture outcome combined with persistent non-adoption.

At 150 EUR/tCO₂, the final states move towards the CCS side of the simplex. This confirms the reference-price result of Section 4.1.2: in both S2 and S6, CCU fails to form an independent final attractor under the flowshare tariff. The CDR case converges to a high-volume CCS outcome, while the No-CDR case converges to a lower-volume CCS outcome, but both are CCS-dominated in strategic terms.

At 300 EUR/tCO₂, the CCS-dominated attractor becomes more pronounced. Final states concentrate near high CCS shares, and CCU ceases to exist as a stable final strategy. The higher ETS price diminishes the influence of the initial strategy distribution, as more emitters find CCS attractive early in the process, enabling infrastructure feedback activation from a broader range of starting points. The remaining No CC share indicates that higher ETS prices do not guarantee full adoption, but they do enhance the robustness of the CCS-dominated outcome.

The simplex analysis thus complements the zero-capture trajectories. Figure 4.11 shows the evolution from the most conservative starting point, while Figure 4.12 shows whether final states depend on initial strategy shares. Together, they show that the ETS price acts as a threshold-shifting parameter. Low prices favour a CCU-including outcome because CCS remains below critical mass, the reference price leads to a CCS-dominated outcome, and high prices make this CCS attractor more robust.

4.2. ETS price sensitivity and equilibrium landscape




FIGURE 4.12: Simplex representations for scenarios S2 and S6 across ETS prices. Each row compares the same ETS price, while the left and right columns correspond respectively to the CDR and No-CDR settings. Grey arrows show the movement from initial to final strategy distributions. Red trajectories highlight representative paths from selected initial conditions.

Chapter 5

Discussion

5.1 From simulation outcomes to transition mechanisms

Across the simulations, carbon-management diffusion appears to depend on threshold effects. CCS becomes more attractive when the realised coalition of adopters brings enough captured volume to improve the use of shared T&S infrastructure. Before this point, transport and storage costs can remain high and isolated CCS entries may disappear or remain marginal. Once the coalition is large enough, average T&S costs decrease and later CCS adoption becomes easier to sustain. This feedback explains why similar techno-economic settings can lead either to low-capture outcomes or to CCS-dominated regimes.

The concept of critical mass is useful here, but must be interpreted carefully. General theories of critical mass describe how adoption can become self-reinforcing after a threshold is crossed [41]. In this model, this threshold is not mainly a fixed share of adopters. It is volume-weighted and spatially structured. A small group of large, well-located emitters can affect T&S costs more than a larger group of smaller or poorly located emitters. The relevant question is whether the entering coalition adds enough captured volume in locations that improve shared-network utilisation.

The ESS mapping adds a stability perspective to this mechanism. It shows that some outcomes are stable strategy-distribution attractors of the ABM-EGT process. Under CDR credits, the high-volume CCS outcome behaves as a robust ESS over the tested initial conditions. Without CDR credits, the stability structure becomes more conditional. S5 and S7 split into distinct attraction regions, while S8 shows the clearest separation between a CCU-dominated ESS and a storage-oriented ESS.

Policy and market assumptions shift this threshold, including through tariff design. Average pricing may appear non-discriminatory, but by socialising network costs it can weaken the spatial signal that favours well-located, high-volume early adopters. Flowshare pricing preserves this signal and therefore helps decentralised coalitions move closer to the centralised least-cost benchmark. The limited persistence of CCU in several CCS-dominated scenarios does not mean that CCU is

unimportant. It becomes more relevant when CCS fails to reach a self-reinforcing T&S configuration, or when utilisation demand can be served under favourable conditions.

The comparison with the standalone CCNO model adds the same message, but from another angle. The optimisation model identifies what can be selected when emitter connections and network design are coordinated, while the ABM-EGT-CCNO model tests whether similar outcomes can emerge from payoff-based strategy revision. Similar stored volumes can sometimes be reached with different emitters and networks, whereas in other cases a centralised opportunity fails to emerge from the decentralised process. The central mechanism is therefore the interaction between individual strategy revision and collective system conditions.

5.2 Positioning the results in the existing literature

The findings are consistent with literature that frames CCS and CCU as value chains rather than isolated capture technologies. Capture feasibility at plant level is not sufficient for deployment, because captured CO₂ also requires downstream transport, storage or utilisation outlets, accounting rules and coordination across actors [58, 20]. This value-chain dependence is visible in the simulations: a technically feasible capture route does not automatically lead to adoption. Depending on whether downstream conditions become attractive, the same system can remain close to No CC, develop a CCS-dominated outcome, or a CCU-dominated outcome. The model thus supports the view that industrial carbon management depends on the joint development of capture, infrastructure and market conditions.

The findings also support studies that identify shared CO₂ infrastructure as a source of economies of scale, first-mover disadvantages and coordination problems [66, 27]. In the model, early CCS adoption can remain fragile because the T&S network is still poorly utilised and costly. CCS becomes easier to sustain only when the realised coalition adds enough useful volume to improve network utilisation and reduce average T&S costs. This gives a dynamic interpretation of the coordination problem discussed in the infrastructure literature. Some early adoption attempts fail to change system conditions enough to survive, while other coalitions can shift the system towards a lower-cost infrastructure regime. This is also where the critical-mass concept becomes specific to shared CO₂ infrastructure: the threshold is not simply a fixed percentage of adopters, but depends on the captured volume and spatial position of these adopters.

The results also relate to the literature on CCS-CCU competition and carbon accounting. Verbist et al. show that ETS design and accounting rules can affect the level playing field between storage and utilisation [61]. The simulations are consistent with this view, while adding a transition-dynamics perspective. Accounting assumptions and CDR credits affect not only the final comparison between CCS and CCU payoffs, but also which emitters become plausible early adopters and whether the CCS coalition can cross the infrastructure threshold. This explains why CCU

becomes more persistent when CCS fails to develop a mature T&S pathway, while it becomes transitional or residual when CCS reaches a low-cost storage regime.

This CCS–CCU balance should also be interpreted in light of the regional scope of the case study. The trilateral region is not neutral with respect to storage opportunities. It combines dense industrial emitters with proximity to North Sea storage options, while European storage potential is spatially uneven, as shown by CO2StoP’s mapping of geological storage locations and capacities [51]. The North Sea is also identified as the largest CO₂ storage resource in Europe [2]. This regional characteristic can make CCS particularly favourable in this region of interest. However, the results should not be reduced to geography alone. ETS price levels, CDR credits, tariff allocation, and CCU demand all influence whether CCS becomes self-reinforcing or whether CCU remains a more persistent outlet. The limited persistence of CCU in several high-volume CCS scenarios should therefore not be interpreted as evidence that CCU is generally unimportant. It is conditional on a storage-favourable regional setting and on the demand, accounting and tariff assumptions represented in the model.

Finally, the comparison with centralised CO₂ network optimisation studies should be interpreted carefully. These studies identify cost-efficient source–sink connections and infrastructure layouts under coordinated deployment assumptions [40, 52]. The ABM-EGT–CCNO results do not always reproduce the same outcomes, but this does not necessarily contradict them. It reflects a difference in the question being asked. Centralised models show what can be selected when emitter connection and network design are optimised together. The coupled ABM-EGT–CCNO model asks whether similar outcomes can emerge when emitters revise strategies individually according to private payoffs and evolving tariff signals. The results suggest that centralised feasibility and decentralised emergence are related but distinct problems. More broadly, the study supports existing findings on value-chain dependence, infrastructure coordination and accounting effects, while showing how these mechanisms can influence adoption in a decentralised setting.

5.3 Reliability of the findings and limits of interpretation

The results should be interpreted with the strengths and limitations of the model in mind. Its main strength lies in linking decentralised adoption decision-making with shared infrastructure feedbacks. This makes it possible to analyse failed early CCS coalitions, threshold effects, and path dependence. Concomitantly, the model remains a simplified representation of industrial carbon-management adoption. Its results are more reliable as qualitative insights into adoption mechanisms than as exact numerical estimates of future deployment.

The seed-dependency checks reported in Appendix H support this interpretation. Changing the stochastic seed affects transition timing, trajectory smoothness and sometimes the precise final shares, as expected in a logit-based ABM-EGT model.

However, the broad scenario classification remains relatively stable in the tested runs: CDR-credit scenarios remain associated with high-volume CCS outcomes, No-CDR flowshare scenarios remain more fragile but can still activate CCS, and the No-CDR average-price setting remains the most difficult configuration for CCS emergence. The main mechanisms do not appear to be artefacts of one random sequence, although exact tipping-point timing should not be over-interpreted.

Several limitations define the scope of the conclusions. Industrial technology choices are preprocessed before the ABM starts: each emitter is represented as an existing site with a no-capture route and, when feasible, a capture route selected under scenario-specific assumptions. The model thus studies adoption among existing emitters in the trilateral region. It does not represent plant relocation, closure, entry of new industrial sites, or competition with production outside the region. The results should be read as adoption mechanisms conditional on the existing spatial distribution of industrial emitters.

The policy environment is fixed by scenario design. The model does not simulate an endogenous ETS trajectory, regulatory change, strategic anticipation of future policy, or uncertainty over support schemes. It can compare adoption regimes under selected policy conditions, but it cannot predict which policy environment will materialise. Similarly, the T&S tariff representation captures the feedback between CCS coalitions and infrastructure costs, but not the full institutional process of CO₂ network planning, financing, regulation and cross-border tariff design. The flowshare and average-price results should be taken as evidence that tariff allocation can influence decentralised adoption dynamics under the tested assumptions.

The CCU submodel is stylised as well. CO₂ demanders are passive nodes with fixed annual demand, the commodity price is exogenous, delivery costs are represented by a reduced-form trucking function, and matching is simplified. The model also treats CCU demand as non-permanent in the standard configuration and neglect detailed CO₂ quality requirements or downstream production choices. The CCU results are useful for analysing how finite and spatially distributed utilisation demand can affect adoption, but not for forecasting future CCU market size or the detailed ranking of utilisation pathways.

Given these strengths and limitations, the study supports a mechanism-level answer to the main research question. In the model, a CO₂ value chain emerges when policy incentives, route-cost inputs, tariff allocation and regional infrastructure conditions allow decentralised emitters to keep capture attractive through a viable downstream route. For CCS, this requires a spatially coherent coalition with enough captured volume to reduce shared T&S costs and make later adoption easier. For CCU, it requires finite utilisation demand that can be matched with captured volumes at a sufficiently attractive net value. CDR credits, higher ETS prices and flowshare-based tariff allocation make the CCS condition easier to reach in the tested setting. These conclusions are strongest as qualitative insights into adoption mechanisms, rather than as precise forecasts of firm-level investments, network layouts or deployment timing.

Chapter 6

Conclusion

6.1 Research problem and modelling objective

This thesis started from the observation that industrial carbon management combines a technology challenge with a coordination problem. Carbon capture and storage (CCS) and carbon capture and utilisation (CCU) require capture at the industrial site, while their feasibility also depends on downstream transport and storage (T&S) access, utilisation demand and the economic rules governing these value chains. An emitter may be technically able to capture CO₂, yet still lack an attractive downstream option for transporting, storing or selling it.

This thesis aims to study how such a CO₂ value chain can develop when industrial emitters make adoption decisions under shared infrastructure and market interdependencies. The analysis covered a trilateral industrial region comprising Belgium, the Netherlands and North Rhine-Westphalia, represented through a fixed population of emitters, a candidate CO₂ transport and storage system, and geographically distributed CO₂ utilisation demand. These inputs constitute a structured modelling environment for analysing adoption mechanisms rather than a direct representation of actual infrastructure plans.

To address this objective, the thesis developed an agent-based evolutionary-game model in which heterogeneous emitters revise their strategy between continued operation without carbon capture (No CC), CCS and CCU. CCS payoffs were coupled to a Centralised CO₂ Network Optimisation (CCNO) model, so that T&S costs could evolve with the realised CCS coalition. CCU was represented through spatial demand nodes, finite utilisation demand, and delivery costs. This design makes it possible to study phenomena such as critical mass, coalition formation, tariff allocation, carbon-accounting incentives, CCU demand constraints, evolutionarily stable strategy distributions, and the differences between decentralised adoption and coordinated techno-economic deployment.

6.2 Answering the research questions

For the first sub-question, concerning carbon-accounting rules, T&S tariff design and CCU availability, the results show that these dimensions mainly affect the conditions under which capture-based strategies can take hold and persist. CDR credits make

high-volume storage routes more attractive, especially when large biogenic CO₂ streams are available. This helps a CCS coalition reach the point where shared T&S costs fall enough to support further adoption more readily than without CDR credits. Tariff design matters because it changes how shared infrastructure costs are allocated across emitters. In the tested scenarios, the flowshare tariff preserves a stronger link between network use, location, and emitter-facing costs. It also preserves the cost advantage of well-positioned early adopters during the high-cost phase. CCU follows a different logic: it becomes most relevant when captured CO₂ can be matched with finite utilisation demand at a positive net value, and when the CCS pathway has not reached a mature T&S configuration.

For the second sub-question, concerning shared T&S infrastructure, critical mass, tipping points and lock-in, the main insight is that critical mass is defined by whether adopting emitters bring enough captured volume in favourable locations to lower shared T&S costs. CCS can remain unattractive when the coalition is too small or poorly positioned, because T&S costs remain high and early entries may disappear. This explains why some temporary CCS entries fail, while other coalitions trigger a fall in average T&S costs and make later adoption easier. Tipping points are the moments when such a cost-reducing coalition forms, while lock-in occurs when it does not and the system remains close to No CC or shifts towards utilisation-centred capture. The ETS sensitivity analysis supports this interpretation: a higher ETS price makes the threshold easier to cross, while a lower ETS price leaves more room for utilisation-centred capture outcomes.

For the third sub-question, concerning ESS and their conditions of emergence, the simulations identify several stable strategy-distribution attractors. With CDR credits, the system generally converges towards a high-volume CCS-dominated ESS across the tested initial conditions. This indicates that the storage pathway is robust under these assumptions. Without CDR credits, the ESS structure becomes more sensitive to tariff design and initial coalition composition. Flowshare pricing can still support fossil CCS-dominated ESS, while the No-CDR average-price setting can lead to low-capture or CCU-centered ESS. The clearest case of multiple ESS appears in S8, the No-CDR, average-price, CCU-enabled scenario, where the simplex analysis shows two distinct basins of attraction: one leading to a CCU-dominated strategy distribution and another leading to a storage-oriented outcome. The ETS sensitivity analysis also shows that the dominant ESS changes with the carbon-price incentive. At lower ETS prices, CCU-including ESS become more prominent; at higher ETS prices, the CCS-dominated ESS becomes more robust across the tested initial conditions.

For the fourth sub-question, concerning the comparison between ABM-EGT-CCNO outcomes and the centralised benchmark, the results show that centralised feasibility and decentralised emergence are related but distinct questions. The CCNO benchmark identifies what can be selected when emitter connection, network design and storage use are optimised together. The ABM-EGT-CCNO model evaluates whether a comparable outcome can arise when emitters revise strategies individually according to private payoffs and evolving tariff signals. In several cases, decentralised adoption reaches storage volumes close to the centralised benchmark,

while relying on a different set of emitters and a different realised network. In other cases, a cost-efficient centralised pathway does not appear in the decentralised process.

Combining these four answers, the central research question can be answered as follows. In the model, a CO₂ value chain emerges when policy incentives, route selection, tariff allocation, regional infrastructure conditions and downstream demand allow decentralised emitters to cross the relevant infrastructure threshold or to access sufficiently attractive utilisation outlets. For CCS, this threshold is shaped by captured volume, spatial position and the cost allocation of shared T&S infrastructure. For CCU, adoption depends on the availability, location and value of utilisation demand. Hence, industrial carbon-management diffusion is best understood as an emergent process in which individual adoption decisions and collective system conditions co-evolve.

6.3 Contributions and implications

This thesis contributes to the modelling of industrial carbon management by linking decentralised adoption dynamics with shared CO₂ infrastructure. Centralised optimisation models are well suited to identifying cost-efficient source–sink connections and infrastructure layouts, whereas ABM and EGT approaches are well suited to representing heterogeneous actors and adaptive strategy revision. The model developed here connects these perspectives by coupling individual strategy revision with T&S costs generated by the realised CCS coalition. Its main added insight is therefore not only the comparison with a coordinated least-cost benchmark, but the identification of stable strategy-distribution attractors and their basins of attraction. This type of stability mapping is not provided by the standalone optimisation model, which selects a coordinated least-cost configuration, nor by a static equilibrium formulation that does not trace the revision path.

The main analytical contribution is to show that critical mass is not a fixed percentage of CCS adopters, but depends on the volume and location of the emitters that enter the coalition. In the simulations, CCS diffusion depends strongly on which emitters enter the coalition. Large emitters in favourable parts of the network can reduce T&S costs enough to support later adoption. Smaller or poorly located coalitions may fail to shift network conditions, even when several emitters temporarily adopt CCS. Early deployment support may be most effective for emitters whose captured volumes and network positions can help the CCS coalition reach the critical-mass threshold.

The results also have a clear implication for T&S tariff design: among the two tariff logics tested, flowshare pricing performs more favourably because it preserves the link between network use, location, and emitter-facing costs. By contrast, uniform average pricing may appear non-discriminatory, but it can weaken the spatial cost signal needed for efficient coalition formation. This is particularly relevant when early CCS adoption is fragile: by limiting the cost burden imposed by more remote users, flowshare pricing preserves the advantage of efficient early adopters and brings the decentralised outcome closer to the centralised social-optimum reference. From

a policy-design perspective, this suggests that location-sensitive T&S tariffs may be preferable to uniform average pricing for efficient early CCS network formation.

CCU becomes most relevant when finite utilisation demand can be matched with emitters at a sufficiently attractive net value, and when CCS has not reached a mature low-cost T&S configuration. Its adoption depends on finite demand, spatial matching, delivery costs, and carbon-accounting treatment. This suggests that CCU can support industrial carbon management in settings where storage is delayed, costly or less accessible, or where local utilisation demand provides a sufficiently attractive outlet for captured CO₂.

The broader implication for policy would be that industrial carbon management requires a value-chain perspective. Carbon prices, CDR credits, T&S tariffs and CCU accounting rules affect private payoffs, but they also influence which emitters enter early, which volumes are added to the system, and whether shared infrastructure moves into a lower-cost regime. In the region studied here, the most favourable conditions for CCS arise when policy incentives, high-volume emitters, favourable locations, and tariff allocation jointly support early network utilisation. For CCU, the key requirement is sufficiently valuable and accessible demand. CCS and CCU deployment therefore depend on the alignment between firm-level incentives and the collective conditions required to build a functioning CO₂ value chain.

6.4 Future research directions

This thesis opens several directions for future research. A first extension concerns CO₂ infrastructure governance and tariff design. The present model translates the T&S network costs returned by the CCNO model into emitter-facing tariffs under a harmonised regional framework. Future work could represent infrastructure operators, regulated cost recovery, and cross-border access rules more explicitly. This would make it possible to study tariff design as part of a broader governance problem for early CO₂ network formation.

A second direction concerns the CCU side of the model and the empirical calibration of downstream market assumptions. In this thesis, utilisation is represented through finite spatial demand nodes, an exogenous CO₂ commodity value and a reduced-form delivery-cost function. Future work could distinguish more detailed utilisation pathways, including permanent and non-permanent products, CO₂ quality requirements, downstream willingness to pay and competition in product markets.

A third direction is to broaden the scenario space and to apply the framework to other regional contexts. Future work could examine carbon-price trajectories, different CCU-demand levels and alternative T&S infrastructure and tariff assumptions. Applying the model to regions with weaker storage access, different industrial structures or stronger utilisation markets would help identify which mechanisms are general and which depend on the trilateral setting. Notwithstanding, this thesis offers one step towards modelling CCS and CCU diffusion as a value-chain process whose outcome depends on decentralised decisions, shared infrastructure, and downstream market conditions.

Appendices

Appendix A

ODD design concepts

This appendix complements the ODD-based model description in Section 3.2 by applying the full design-concept checklist of the ODD protocol to the ABM-EGT model developed in this thesis. The goal is not to introduce new model mechanisms, but to make explicit how the main behavioural and system-level design choices relate to the standard ODD design concepts. The questions used below are those proposed by Railsback and Grimm [53].

A.1 Basic principles

Questions

What general concepts, theories, hypotheses, or modelling approaches underlie the model's design? How is the model related to previous thinking about the problem it addresses? How were these principles incorporated in the model's design [53]?

Answer

The model is constructed on two primary methodological foundations: Evolutionary Game Theory (EGT) and Agent-Based Modelling (ABM). EGT establishes the logic for strategy competition within an evolving population, whereas ABM offers the computational structure for representing individual industrial emitters as heterogeneous, decision-making agents.

The model is also shaped by literature on industrial carbon management value chains, shared CO₂ transport and storage infrastructure, coordination failures, path dependence, carbon accounting, and policy-driven diffusion. These principles are directly incorporated into the model structure. Carbon Capture and Storage (CCS) depends on common transport and storage conditions, while Carbon Capture and Utilisation (CCU) is influenced by spatially distributed demand for CO₂. Policy assumptions are integrated through carbon prices, credits for Carbon Dioxide Removal (CDR), and accounting rules.

The model does not assume the presence of a central planner or operator determining which emitter adopts CCS, CCU, or No CC. Instead, it focuses on decen-

tralised industrial decision-making and examines whether system-level deployment patterns can emerge from these individual decisions.

A.2 Emergence

Questions

What are the model's important results and outputs? Which of them emerge from the mechanistic representation of individual adaptive behaviour, and which are imposed by model rules [53]?

Answer

The principal outcomes of the model are emergent. The model does not pre-determine the final adoption shares of No CC, CCS, or CCU. It also does not enforce the formation of a CCS coalition, the total amount of CO₂ captured, the allocation of captured CO₂ to CCU demanders, or even the existence of tipping-like transitions.

These results arise from the repeated emitter-level game with strategy revision. Each emitter evaluates its feasible strategies based on its private payoff in the current system state. These individual decisions aggregate to generate system-level patterns, including CCS diffusion, CCU uptake, persistent low-adoption outcomes, mixed equilibria, critical-mass effects, and path dependence.

Some elements are specified by scenario design and function as model inputs. For example, CDR credit assumptions, carbon price, CCU total demand, and tariff mode are exogenous inputs. Nevertheless, the final strategy distribution that emerges from these assumptions is not imposed but results from the model's internal dynamics.

A.3 Adaptation

Questions

What adaptive behaviours do agents have? In what ways can they respond to changes in their environment and themselves? How are these behaviours modelled [53]?

Answer

The adaptive agents in the model are large-scale industrial emitters. Their adaptive behaviour involves revising their carbon management strategy, which occurs when they receive a revision opportunity with a specified probability. During revision, each emitter compares the feasible strategies available: No CC, CCS, and CCU.

This behaviour is modelled as direct objective-seeking. Emitters do not replicate observed behavioural patterns. Instead, they evaluate each possible strategy based on a private payoff calculation. In the deterministic version, the emitter selects the strategy with the highest payoff. In the stochastic logit version, strategies with

higher payoffs are more likely to be chosen, but strategies with lower payoffs may still be selected with lower probability.

Emitters can respond to changes in transport and storage (T&S) tariffs, CCU market conditions, and learning effects. However, they do not make direct infrastructure investment decisions, nor do they coordinate collectively with other emitters.

A.4 Objectives

Questions

For adaptive behaviours modelled as direct objective-seeking, what measure of agent objectives is used to rate decision alternatives? How does the objective measure represent processes that link adaptive behaviour to important variables of the agents and their environment [53]?

Answer

Each emitter seeks to maximise its private annual payoff, which is equivalent to minimising the annual cost of its carbon management strategy. This objective reflects the private economic attractiveness of No CC, CCS, and CCU from the perspective of the industrial site.

The payoff calculation incorporates route-specific production and capture costs, remaining or avoided carbon costs, CDR-related revenues (where applicable), T&S tariffs for CCS, CCU revenues, CCU delivery costs, fallback storage costs for unsold CCU volumes, and the impact of learning on capture-specific CAPEX. Consequently, emitter decisions depend on both their own characteristics and the overall system state.

This objective is private rather than social. Emitters do not maximise total system welfare, total CO₂ abatement, or the cost efficiency of the entire CO₂ network.

A.5 Learning

Questions

Do individuals change how they make adaptive decisions over time as a consequence of their experience? If so, how [53]?

Answer

The model incorporates technological learning at the level of the individual emitter, but not behavioural learning. Emitters do not learn from their own past decisions, update beliefs based on experience, or alter their revision rule during the simulation. They apply the same decision logic throughout each run.

Learning is modelled at the technology level. As the cumulative adoption of capture-based strategies increases, the capture-specific CAPEX component decreases up to a lower bound. This is learning-by-doing in carbon capture technology. The

learning effect, therefore, changes the carbon capture payoffs that emitters are facing, but it does not alter the behavioural rule that emitters are using to choose between strategies.

A.6 Prediction

Questions

How do agents predict future conditions in their adaptive behaviour? How does simulated prediction make use of mechanisms such as memory, learning or environmental cues? Or is prediction only tacitly implied in simple adaptive rules [53]?

Answer

The emitters do not have perfect vision of the future system they are about to encounter. During revision, they evaluate their feasible strategies based on the latest available system state. Thus, their decision is “myopic”. They do not predict a sequence of future adoption decisions, future CCNO network updates or even future CCU market-clearing outcomes.

For CCS, emitters use the latest available T&S tariff informations. This information is extracted from the most recent computed CCNO network.

For CCU, emitters evaluate the current CCU market conditions. This includes available demand, delivery costs, CO₂ commodity revenues and fallback storage quantities. Therefore, prediction is not modelled as explicit forecasting. It is represented as the evaluation of strategies in the current perceived state of the system.

A.7 Sensing

Questions

What variables of their environment and themselves are agents assumed to sense and therefore consider in their behaviour? What sensing mechanisms are modelled explicitly, and which sensed variables are simply assumed to be known? With what accuracy are agents assumed to know these variables [53]?

Answer

It is assumed that emitters have access to the information required to evaluate their feasible strategies. This includes their production route costs, capture feasibility, captured CO₂ volume, residual emissions, fossil and biogenic CO₂ shares, and plant location. Emitters also have access to economic and policy parameters relevant to the scenario, such as the carbon price, CDR credit assumptions, CO₂ commodity price assumptions, and accounting rules.

Emitters also have knowledge of system-level “variables” that are needed to compute their payoffs. For CCS, this includes the relevant T&S tariff information based on the current T&S network state. For CCU, this includes the current CO₂

demand market, delivery-cost assumptions and fallback quantities. Emitters also have access to the learning factor, as it has a direct impact on the CAPEX of their capture-based strategies.

No explicit sensing mechanism is modelled in this thesis. When emitters update their strategy, the relevant variables are assumed to be known. However, this does not imply that emitters possess perfect foresight. They are aware of current payoff-relevant information but do not know how other emitters will revise their strategies in future iterations.

A.8 Interaction

Questions

How do the model's agents interact? Do they interact directly with each other, or is interaction mediated through the environment or competition for a resource? With which other agents does an agent interact [53]?

Answer

Interactions between emitters are primarily indirect. Emitters do not negotiate with each other or directly alter each other's state variables. Instead, interaction occurs through shared infrastructure, market-clearing constraints, and aggregate adoption states.

For CCS, interaction between emitters occurs through the CCS coalitions. The emitters choosing CCS as a strategy define the CO₂ volumes that undergo the shared T&S network. This influences the network evaluated by the CCNO model and therefore impacts the T&S tariffs faced by emitters in later revisions. Hence, CCS adoption creates a coordination mechanism: one emitter's strategy can affect the attractiveness of CCS for others.

For CCU, interaction occurs through competition for finite and spatially heterogeneous CO₂ demand. If multiple emitters select CCU, they may compete for the same downstream CO₂ demand. Unsold CCU volumes are redirected to fallback storage, which also influences the payoff of the CCU strategy. Consequently, CCU adoption depends not only on the emitter's own characteristics but also on the number, location, and captured volumes of other CCU adopters.

A.9 Stochasticity

Questions

How are stochastic processes used in the model and why? Are they used for initialisation, variable processes, or empirically determined probabilities [53]?

Answer

Stochasticity is incorporated during initialisation and to represent decentralised decision timing and imperfect strategy revision among emitters. During model initialisation, initial adoption shares are assigned to emitters randomly. Revision opportunities are also stochastic; in each iteration, only a subset of emitters is permitted to re-evaluate and reconsider their strategy, while others retain their current strategy. When the logit revision rule is applied, strategy choice becomes probabilistic by assigning probabilities to payoffs. Strategies with higher payoffs are more likely to be chosen, but those with lower payoffs may still be selected. This approach represents bounded rational decision-making and avoids assuming that all emitters always select the exact best response deterministically.

Reproducibility of individual simulation runs is ensured by using random seeds. To assess the robustness of the model's outcomes, the model can be executed across different random seeds.

A.10 Collectives

Questions

Are collectives, defined as aggregations of agents that affect the state or behaviour of member agents and are affected by their members, represented in the model? If so, how are they represented [53]?

Answer

The model does not represent collectives as agents making independent decisions. There is no dedicated agent representing an emitter consortium, cluster authority, or T&S infrastructure operator making strategic decisions within the ABM. Although there is indeed an agent representing the Carbon Network Operator, this agent is a passive type of agent in ABM.

However, collective effects are central to the model. For example, CCS adopters form a coalition that, while not an autonomous agent, determines the CO₂ volumes passed to the CCNO model. This, in turn, affects the state of the T&S network and the tariffs faced by emitters. Industrial clusters are also relevant, as emitters in close geographical proximity may share infrastructure and experience similar cost-allocation effects, depending on the T&S tariff mode.

Collectives are thus represented as emergent or derived system states rather than as agents with independent behavioural rules.

A.11 Observation

Questions

What outputs from the model are needed to observe its internal dynamics and system-level behaviour? What outputs and analyses are needed to solve the problem the model was designed for [53]?

Answer

The model records two types of outputs: emitter-level and aggregate outputs. Aggregate outputs include adoption shares of No CC, CCS, and CCU, total captured CO₂, CCS and CCU flows, fallback quantities, residual emissions, T&S tariff indicators, learning indicators, and convergence diagnostics. Aggregate outputs enable observation of system evolution across model iterations and assessment of adoption stabilisation.

Emitter-level outputs represent final strategies, payoff components, and related CO₂ flow variables. These outputs are necessary to understand the motivations for adoption or rejection of CCS and to relate aggregate diffusion patterns to heterogeneous industrial conditions.

These outputs are selected to address the research questions of the thesis. They enable comparison of CCS-dominated, CCU-including, mixed, and low-adoption outcomes, and facilitate identification of critical mass and path-dependent behaviour. Additionally, these results support analysis of tariff design and policy assumptions, and allow comparison of decentralised ABM-EGT outcomes with the centralised CCNO benchmark.

Appendix B

Input data and parameter values

This appendix summarises the main numerical inputs used in the reported ABM-EGT simulations. It complements the model description in Chapter 3 and the construction of the CCU-demand database in Appendix C. Unless stated otherwise, the values refer to the final logit runs reported in Chapter 4.

TABLE B.1: Main input data and parameter values used in the reported ABM-EGT simulations

Parameter	Value	Unit	Source	Comment
<i>Reported scenario and behavioural parameters</i>				
Reference ETS price	150	EUR/tCO ₂	Scenario design	Base scenarios
ETS sensitivity values	75; 150; 300	EUR/tCO ₂	Scenario design	Sensitivity analysis
Carbon-accounting variants	CDR; No CDR	–	Scenario design	Route-cost variants
T&S tariff variants	flowshare; endogenous average price	–	Scenario design	Base-scenario dimension
CCU availability	off; on	–	Scenario design	Base-scenario dimension
Base initial strategy shares	100% No CC; 0% CCS; 0% CCU	–	Scenario design	Eight base scenarios
Main random seed	2	–	Final-run configuration	For main figures in Chapter 4
Strategy revision probability	$p_{\text{revision}} = 0.05$	per emitter per iteration	Final-run configuration	Reported logit runs
Strategy revision rule	logit	–	Section 3.2.4.5; [44, 6]	Reported runs
Logit sensitivity parameter	$\beta = 5$	–	Final-run configuration	Payoff sensitivity
Payoff scaling parameter	$\Lambda = 10^6$	EUR/yr	Section 3.2.4.5	Logit normalisation
Mutation probability	$\mu = 0.001$	per revision event	Final-run configuration	Reported logit runs
<i>CCS and T&S parameters</i>				
Bootstrap T&S tariff	35	EUR/tCO ₂	Model assumption	Early CCS tariff
Minimum CCS coalition for network update	2	emitters	Model implementation	Network-evaluation trigger
Capture-CAPEX learning rate	5	%	Model assumption; [55, 54]	Per doubling of capture adopters
Reference deployment level	$N_{\text{ref}} = 1$	project	Section 3.2.4.1	Learning normalisation
Minimum learning multiplier	$f_{\text{min}} = 0.75$	–	Model assumption	CAPEX lower bound
Entrant flow floor	0.01	MtCO ₂ /yr	Model implementation	Entrant-tariff safeguard
Tap-in cost for first CCS entrant in a cluster	3.0	MEUR/yr	Model assumption	New-cluster access proxy
Tap-in cost in an active CCS cluster	0.25	MEUR/yr	Model assumption	Existing-cluster access proxy
Maturity-weighting coefficient	$\eta = 0.5$	–	Model assumption	Entrant-tariff weighting
Mature-cluster threshold	2	MtCO ₂ /yr	Section 3.2.4.3	Flow threshold

Parameter	Value	Unit	Source	Comment
Mature-cluster threshold	3	CCS emitters	Section 3.2.4.3	Emitter-count threshold
<i>CCU parameters</i>				
CO ₂ commodity selling price	109.09	EUR/tCO ₂	Section 3.2.4.4	Biomass-substitution value
Dry-biomass reference price	200	EUR/t biomass	dry [3]	Rounded central assumption
Biomass carbon fraction	0.5	tC/t biomass	dry [35]	Generic dry-biomass fraction
Fixed CCU delivery cost	20	EUR/tCO ₂	Appendix F; [57, 5]	
Distance-dependent CCU delivery cost	0.15	EUR/tCO ₂ /km	Appendix F; [57, 5]	
CCU total demand	28.853	MtCO ₂ /yr	Appendix C	Total aggregated demand
Zero-sale grace period	5	iterations	Model implementation	Before reclassification eligibility
Zero-sale reclassification probability	0.05	per iteration after eligibility	Model implementation	Reported logit runs

Appendix C

Construction of the CO₂ Demand Database

This appendix outlines the construction of the CO₂ commodity demand database utilised in the model. The database is not intended to provide a highly detailed forecast of future CO₂ demand, as such projections are difficult to justify for 2050 due to significant uncertainties regarding future CO₂ markets, industrial transformation, and policy conditions. Instead, the database offers a spatially distributed representation of potential CO₂ demand within the study region. Its primary purpose is to enable the model to assess the impact of a CO₂ commodity demand on decentralised adoption decisions among No CC, CCS, and CCU. The focus is on providing a rough estimate of CO₂ commodity demand to facilitate the inclusion of a CO₂ demand market for evaluating CCU within the model.

For modelling purposes, the demand database is divided into three blocks: greenhouse horticulture, conventional food and beverage merchants, and industrial CCU demand. This decomposition does not represent a comprehensive forecast of all potential CO₂ uses in 2050. Rather, it identifies categories for which relevant quantitative evidence is available to illustrate spatial adoption mechanisms considered in this thesis. Greenhouse horticulture is included due to Dutch sources indicating a significant future requirement for externally supplied CO₂ as the sector transitions away from natural gas combustion [60, 25, 38]. Food and beverage demand is included as this sector is identified as one of the dominant parts of the merchant CO₂ market by current European CO₂ market data [18]. Finally, the industrial CCU block is included as European CCU scenarios and roadmaps identify fuels, chemicals, materials and mineralisation as major future utilisation pathways for captured CO₂ [36, 13, 43].

The three blocks are maintained separately because they correspond to distinct types of CO₂ utilisation. Greenhouse horticulture and food and beverage demand primarily represent direct or conventional merchant uses of CO₂, while the industrial CCU block reflects future transformation-based applications of CO₂ as a carbon

feedstock. This separation also reduces the risk of double-counting, since industrial CCU sources are primarily linked to utilisation routes such as e-fuels, chemicals, and materials, rather than conventional food-sector applications.

C.1 Definition of total sectoral demand

The total CO₂ demand, over the trilateral region, used in the model is defined as the sum of three demand blocks:

$$D_{\text{total}} = D_{\text{greenhouse}} + D_{\text{food \& beverage}} + D_{\text{industrial CCU}}. \quad (\text{C.1})$$

The central value for greenhouse horticulture is 2.5 MtCO₂/yr. This is in line with the expected external supply of CO₂ in Dutch greenhouse horticulture as the sector moves away from combustion of natural gas and thus loses part of its internally generated CO₂ supply [60, 25, 38].

To estimate food and beverage demand, the database begins with reported European CO₂ consumption and isolates the share attributable to the merchant food and beverage market. This provides an estimate of the conventional European food and beverage CO₂ demand, which is then scaled and regionalised to the study area. The way regionalisation is applied is explained in Section C.2. The total value used in the model after regionalisation is 1.095 MtCO₂/yr [18, 42].

The industrial CCU demand estimate is the largest and most uncertain block. This block does not reflect the current merchant CO₂ demand. It is an estimate for potential future CO₂ use in industrial applications such as synthetic fuels, chemicals, materials, and mineralisation. The central value is based on a European 2050 CCU roadmap estimate of the use of CO₂ in the EU. This European value is then scaled down to the study region, leading to a regional industrial CCU demand estimate of 25.258M tCO₂/yr [13].

The resulting central demand values are summarised in Table C.1.

TABLE C.1: Central 2050 CO₂ demand used in the ABM-EGT model.

Demand block	Demand [MtCO ₂ /yr]
Greenhouse horticulture	2.500
Food and beverage	1.095
Industrial CCU demand estimate	25.258
Total	28.853

These values should be treated as modelling inputs rather than exact predictions. The aim is to give a plausible order of magnitude for future demand of CO₂ in the region studied in this thesis, while keeping the model transparent and reproducible.

C.2 Regionalisation to the study area

Some demand estimates used to construct the database are available only at the European level, rather than for the specific geographical scope of this thesis. As explained in the ODD protocol, the area of study consists of Belgium, the Netherlands, and the North Rhine-Westphalia part of Germany. In the absence of a strong public dataset for 2050 CO₂ demand in this specific region, European-level demand values are downscaled with a GDP-based allocation factor.

The regional allocation factor is defined as:

$$\lambda_{\text{ROI}} = \frac{GDP_{\text{BE}} + GDP_{\text{NL}} + GDP_{\text{NRW}}}{GDP_{\text{EU}}}. \quad (\text{C.2})$$

Using 2024 GDP values, this gives:

$$\lambda_{\text{ROI}} \approx 0.146. \quad (\text{C.3})$$

The GDP values for Belgium, the Netherlands and the European Union are taken from World Bank data, while the value for North Rhine-Westphalia is taken from NRW.Global Business [64, 46]. This factor is used as a first-order regionalisation technique. It assumes that the share of future demand for CO₂ utilisation in the study region is proportional to the current economic weight of the region in the European Union.

This approximation is necessary. Future CO₂ demand will also be influenced by factors such as industrial structure, port activity, hydrogen availability, chemical clusters, infrastructure development, and policy incentives. However, for the purposes of this thesis, it is assumed that the GDP-based factor offers a transparent and reproducible method for translating European-level demand estimates into a regional input for the model.

C.3 Spatial allocation of demand nodes

The agent-based evolutionary model requires a spatial representation of CO₂ demand. A single aggregate regional demand value is insufficient, as the distance between emitters and potential CCU demand points affects transport costs and, consequently, influences incentives for decentralised adoption. Therefore, total demand is distributed across a set of spatial demand nodes.

The allocation is intentionally approximate. The objective is not to predict the exact 2050 CO₂ demand of individual firms, cities, or industrial sites. Rather, the goal is to develop a plausible spatial structure that enables the model to assess whether local CO₂ demand can enhance the attractiveness of CCU for certain emitters and influence adoption dynamics at the system level.

The total demand of each demand block is distributed over representative spatial nodes:

$$D_i = w_i D_b, \tag{C.4}$$

where D_i is the demand allocated to node i , D_b is the total demand of demand block b , and w_i is the allocation weight of node i within that block. In the absence of reliable 2050 demand forecasts at the node level, equal weights are assigned within each demand category. This approach prevents the dataset from conveying an unwarranted level of precision.

C.3.1 Greenhouse horticulture

Demand for greenhouse horticulture is assigned to four Dutch greenhouse cluster locations. These locations represent areas where greenhouse horticulture is spatially concentrated, such as Westland, Aalsmeer, Venlo and Boskoop. The selected nodes should be considered as cluster centroids rather than single greenhouse companies. The location choice is based on the spatial organisation of Dutch greenhouse regions and Greenports [26, 28].

Demand in the central greenhouse is 2.5 MtCO₂/yr. It is evenly split among the four greenhouse nodes with 0.625 MtCO₂/yr per node. This spatialisation reflects the fact that greenhouse CO₂ demand is expected to be mainly concentrated in the Netherlands, while avoiding an unrealistically detailed representation of individual greenhouse sites.

C.3.2 Food and beverage demand

Food and beverage demand is allocated to existing beverage and brewery production sites in Belgium, the Netherlands, and the German portion of the study region. This allocation reflects typical merchant demand for CO₂, such as beverage carbonation and other food sector applications.

The selected nodes are based on large and publicly identifiable production locations. Examples of such places include breweries, and soft-drink production sites. These are sites of companies like AB InBev, Heineken and Coca-Cola Europacific Partners [1, 31, 14]. The selected locations are chosen because they provide realistic spatial anchors for the standard food and beverage CO₂ demand.

The central food and beverage demand is 1.095 MtCO₂/yr. This is evenly distributed across 13 food and beverage nodes, averaging ~ 0.084 MtCO₂/yr per node. These values should not be considered site-specific demand forecasts. They are modelling allocations employed to represent the spatial presence of conventional merchant CO₂ demand.

C.3.3 Industrial CCU demand estimate

The largest component of the database is the industrial CCU demand estimate, which aligns with the anticipated future demand for CO₂ in industrial CCU applications such as synthetic fuels, chemicals, materials, and mineralisation. Due to

C. CONSTRUCTION OF THE CO₂ DEMAND DATABASE

the lack of reliable forecasts for the distribution of this demand across individual industrial sites in 2050, demand is allocated to existing industrial, chemical, and port clusters.

The selected industrial nodes are not assumed to be future CO₂ customers. They represent the locations of cluster-centers where one might plausibly locate future CCU demand. This constructed database includes eight industrial and chemical clusters:

1. Antwerp-Bruges
2. North Sea Port/Ghent
3. Rotterdam
4. Moerdijk
5. Chemelot/Geleen
6. CHEMPARK Leverkusen-Dormagen-Krefeld
7. Marl Chemical Park
8. Duisburg/Ruhr

These clusters are supported by public information on chemical, port and industrial activity in the region [47, 8, 12, 24, 49, 50, 48, 19, 9, 11, 15, 10, 22, 56].

The industrial CCU demand estimate is 25.258 MtCO₂/yr. This is divided evenly over the eight industrial cluster nodes, or about 3.157M tCO₂/yr per node. This even distribution avoids over-concentration of future industrial demand to a few large hubs and the illusion of facility-level accuracy.

C.4 Final structure of the demand database

The final demand database contains 25 spatial demand nodes. Table C.2 summarises the spatial structure of the database.

TABLE C.2: Spatial structure of the central CO₂ demand database.

Demand block	Total demand [MtCO ₂ /yr]	Nodes	Demand per node [MtCO ₂ /yr]
Greenhouse horticulture	2.500	4	0.625
Food and beverage	1.095	13	0.084
Industrial CCU demand estimate	25.258	8	3.157
Total	28.853	25	—

This database is therefore a compromise between geographical realism and methodological caution. The use of an aggregated demand value would ignore the spatial

dimension entirely, which is relevant in terms of incentives for CCU adoption. On the other hand, a very detailed facility-level database would imply a level of accuracy that is not found in the evidence on demand in 2050. The selected approach is to use real geographical anchors and broad sectoral demand estimates, and the interpretation of the node-level values is explicitly approximate.

In this model, these demand nodes are used to represent the spatial availability of the CO₂ utilisation opportunities. Their role is to influence the relative attractiveness of CCU for emitters close to potential demand points. Thus, the database is intended to support the analysis of adoption dynamics, coordination effects and CCS-CCU competition, not to predict the precise future size of individual CO₂ markets.

Appendix D

Convergence test

This appendix describes the stopping criterion used in the ABM simulations. The model is not required to reach a state in which no emitter ever changes strategy again. This is especially relevant for the logit revision rule. In this revision rule, the choice of strategies is probabilistic. Logit dynamics are typically studied as stochastic best response processes, where the long run behaviour is characterised by convergence to, or fluctuations around, a stationary distribution, rather than convergence to a deterministic absorbing state [6]. The analysis of agent-based model outputs also often requires treating the simulation results as iterative series and testing the stability of the aggregate indicators [39]. Hence convergence here is defined on the aggregate level: the model is considered converged when the adoption shares of No CC, CCS and CCU are approximately stationary.

At each iteration t , the model records the vector of adoption shares:

$$x_t = (x_t^{NO}, x_t^{CCS}, x_t^{CCU}), \quad (\text{D.1})$$

where x_t^{NO} , x_t^{CCS} and x_t^{CCU} are the shares of emitters adopting No CC, CCS and CCU, respectively. In CCS-only simulations, the CCU share remains equal to zero, but the same logic can still be applied.

The convergence test is based on the evolution of these adoption shares over moving windows. The idea is simple: the model compares recent average adoption shares to average adoption shares in the previous period. If these averages are very close, then the system is no longer strongly growing toward one strategy or the other. For a window of length W , the average share of strategy s over the most recent window is:

$$\bar{x}_{t,W}^s = \frac{1}{W} \sum_{k=t-W+1}^t x_k^s. \quad (\text{D.2})$$

The average over the previous window of the same length is:

$$\bar{x}_{t-W,W}^s = \frac{1}{W} \sum_{k=t-2W+1}^{t-W} x_k^s. \quad (\text{D.3})$$

The short-term drift is then defined as:

$$d_t = \max_{s \in \{NO, CCS, CCU\}} |\bar{x}_{t,W}^s - \bar{x}_{t-W,W}^s|. \quad (D.4)$$

This indicator measures whether the average adoption shares are still moving over the short term. If d_t is small, this means that the newest window will look very similar to the previous one.

But a short-term test in isolation can be misleading. The system may seem stable for a short time, but could still slowly drift over a longer period. For this reason, the logit convergence test also uses a long-term drift computed in the same way but with a longer window W_L :

$$d_t^L = \max_{s \in \{NO, CCS, CCU\}} |\bar{x}_{t,W_L}^s - \bar{x}_{t-W_L,W_L}^s|. \quad (D.5)$$

The short-term drift thus checks for the local stability of the system, and the long-term drift checks for the presence of a slower structural movement.

The third applied convergence criterion is the "volatility" of adoption shares in the most recent long window. This is basically the maximum standard deviation over this window:

$$\sigma_t = \max_{s \in \{NO, CCS, CCU\}} \text{std}(x_{t-W_L+1}^s, \dots, x_t^s). \quad (D.6)$$

This criterion is needed because two windows can have similar average values even if the adoption shares still fluctuate strongly within the window. The volatility test, hence, ensures that the system is not only stable on average, but also sufficiently calm around that average.

In the main logit configuration, the model is considered stationary when the following three conditions are satisfied at the same time:

$$d_t < \varepsilon_{\text{mean}}, \quad d_t^L < \varepsilon_{\text{long}}, \quad \sigma_t < \varepsilon_{\text{vol}}. \quad (D.7)$$

The first condition limits short-term drift, the second limits long-term drift, and the third limits persistent stochastic fluctuations. The model does not stop as soon as these conditions are satisfied once. They must remain satisfied for several consecutive iterations. This avoids stopping the simulation during a temporary calm period. The values of the parameters used in the simulation are shown in the Table [D.1](#).

Because the long-term drift compares two long windows, the full logit convergence test can only be evaluated after at least $2W_L$ iterations. With $W_L = 100$, this means that the complete stationarity test can only start from iteration 200. This also ensures that the model has enough history before declaring convergence.

D. CONVERGENCE TEST

TABLE D.1: Convergence parameters used in the main logit simulations

Parameter	Value	Role
Short window W	15 iterations	Short-term drift window
Long window W_L	100 iterations	Long-term drift and volatility window
$\varepsilon_{\text{mean}}$	7×10^{-3}	Maximum accepted short-term drift
$\varepsilon_{\text{long}}$	1.3×10^{-2}	Maximum accepted long-term drift
ε_{vol}	2.3×10^{-2}	Maximum accepted volatility
Confirmation windows	5	Consecutive successful checks required
Maximum iterations	5000	Hard stopping limit

The code also includes a separate stopping rule for best-response simulations. Since best-response dynamics are less stochastic than logit dynamics, this rule can be stricter: in addition to drift and volatility, it checks whether the average number of switching emitters and new capture adopters has become sufficiently small. However, the main simulations in this thesis rely on the logit stationarity criterion described above.

The simulation can stop in two different situations. First, it stops if the convergence conditions, as explained before, are met for the required number of consecutive checks. In that situation, the final state is considered an approximately stationary regime of adoption. Secondly, the simulation stops if the maximum number of iterations is reached. In that case, the simulation has not met the convergence criterion within the allowed computational horizon. In that situation, the final state must be interpreted with that limitation in mind.

Appendix E

Learning-by-doing for capture CAPEX




FIGURE E.1: Learning-by-doing factor applied to the capture-specific CAPEX.

Appendix F

Calibration of the reduced-form CCU delivery-cost function

This appendix explains the reduced-form CO₂ delivery-cost function used in the CCU submodel. In the ABM, CCU demanders are represented as passive demand nodes with a location and an annual demand. The model thus does not include a detailed logistics optimisation calculation. However, each potential CCU transaction still requires a per-tonne delivery cost between the emitter and the demander. The function computing that cost must hence be simple enough to be computed for every emitter-demander pair, while remaining consistent with the order of magnitude reported in the literature.

We choose here to represent the full physical delivery-chain cost as a fixed component and a distance-dependent component:

$$c_{i,d}^{chain} = c_{truck}^{fix} + c_{truck}^{dist} d_{i,d}, \quad (\text{F.1})$$

where $d_{i,d}$ is the distance between emitter i and demander d in km. This structure is in line with the techno-economic descriptions of liquefied CO₂ trucking chains that separate distance-insensitive logistics steps, such as liquefaction, buffer storage, conditioning and loading, from the transport operation itself, whose cost increases with distance [5]. In this reduced-form function, the fixed component is the first group of costs, and the distance-dependent component is the trucking part of the delivery chain.

The numerical calibration is based on the cost range reported by Sneyd et al. for CO₂ transport in e-fuel supply chains. They suggest that CO₂ transport costs might be in the range of 20 – 50 EUR/tCO₂ for distances up to about 200 km [57]. The values are treated as levelised per-tonne delivery costs for representative CO₂ flows, not as fixed project costs or costs re-calculated endogenously per transaction volume. The lower value, 20 EUR/tCO₂, is used as a short-distance cost floor, while the upper value, 50 EUR/tCO₂, is imposed at 200 km:

$$c^{chain}(0) = 20, \quad c^{chain}(200) = 50. \quad (\text{F.2})$$

With the assumption of a linear function $c^{chain}(d) = a + bd$, this gives $a = 20$ and:

$$b = \frac{50 - 20}{200} = 0.15 \text{ EUR/tCO}_2/\text{km}. \quad (\text{F.3})$$

The resulting full delivery-chain function is therefore:

$$c_{i,d}^{chain} = 20 + 0.15d_{i,d}. \quad (\text{F.4})$$

This formulation should be viewed as a reduced form approximation rather than a full engineering model of liquefied CO₂ trucking. In particular, the model does not explicitly link the cost of delivery to the volume transported in each transaction. This is a modelling simplification: calibration values are already expressed in EUR/tCO₂ and are used as representative levelised costs. Thus, the function captures the first-order effect of distance on delivery cost, but not economies of scale, truck utilisation rates, storage sizing, loading constraints or route-specific logistics.

The second modelling decision concerns the sharing of costs. Verbist et al. assume that the CO₂ transport costs are shared between both sides of a transaction between an upstream CO₂ provider and a downstream CO₂ user [61]. This thesis extends that logic to the cost of the entire delivery-chain. The physical supply-chain cost is still the cost in Equation F.4, but only half of that cost is included in the emitter's CCU payoff. The other half is implicitly attributed to the downstream demander.

So the cost of delivery on the emitter side is:

$$c_{i,d}^{del,E} = \frac{1}{2}c_{i,d}^{chain} = \frac{1}{2}(20 + 0.15d_{i,d}) = 10 + 0.075d_{i,d}. \quad (\text{F.5})$$

As explained in the ODD protocol, demanders are passive agents in the ABM and have no own payoff function. Hence, their share of the delivery-chain cost is not separately modelled. So the net revenue for the emitter from selling one tonne of CO₂ to demander d is:

$$NR_{i,d} = p^{CO_2} - c_{i,d}^{del,E}. \quad (\text{F.6})$$

Appendix G

Reading simplex plots

Several figures in Chapter 4 summarise adoption dynamics using simplex plots. Instead of displaying absolute numbers of emitters, these figures represent the proportion of emitters adopting each available strategy. This approach enables a concise comparison of initial conditions, final outcomes, and transition directions across multiple simulations. This appendix provides a brief guide to interpreting the one-dimensional simplex for two-strategy cases and the triangular simplex for three-strategy cases.

One-dimensional simplex plots

When only two strategies are available, the simplex reduces to a line segment. In CCS-only scenarios, this line illustrates the trade-off between No CC and CCS. Movement to the right indicates a higher CCS share, while movement to the left indicates a higher No CC share. Because only two strategies are present, the shares are complementary; for example, a CCS share of 0.7 corresponds to a No CC share of 0.3.

In these plots, empty circles denote initial conditions and red dots denote final conditions. Grey arrows indicate the transition from each initial condition to its final state. If many arrows converge in the same region, the system exhibits convergence toward a similar outcome from different starting points. Conversely, if arrows originating from different regions terminate in different final states, the outcome is path-dependent.




FIGURE G.1: Example of a one-dimensional simplex plot for a two-strategy case.

Triangular simplex plots

When three strategies are available, the simplex is depicted as a triangle. In this thesis, the three vertices correspond to pure strategy distributions: 100% No CC, 100% CCS, or 100% CCU. Points within the triangle represent mixed strategy distributions. For instance, a point near the CCS vertex indicates a high CCS share, while a point near the centre represents a more balanced mixture of the three strategies.

The value of each strategy share is determined using the grid lines; lines parallel to one side of the triangle indicate constant values of the opposite component. In the figures presented, the blue axis represents the CCS share, the grey axis represents the No CC share, and the green axis represents the CCU share. The three values obtained from the plot always sum to one.

As in the one-dimensional case, empty circles denote initial conditions and red dots denote final conditions. Grey arrows summarise the movement from initial to final strategy shares. Red trajectories illustrate selected examples of the full transition path during the simulation, facilitating the identification of whether a final outcome is reached directly or through intermediate strategy changes.




FIGURE G.2: Example of a triangular simplex plot for a three-strategy case.


Appendix H

Seed-dependency study

This appendix checks whether the main results depend strongly on the stochastic seed. A seed is the number used to initialise the random parts of the simulation. Changing the seed keeps the same model, input data and scenario parameters, but changes the order of revision opportunities and stochastic strategy choices.

The objective is not to obtain identical trajectories for every seed. Since the model uses stochastic logit revision, different seeds are expected to change the timing of adoption and may slightly affect final outcomes. The figures below compare the adoption-share trajectories across several seeds for the eight studied scenarios from zero-adoption. The horizontal axis is expressed as a percentage of the maximum number of iterations, so that runs with different convergence times can be compared visually. Solid lines show CCS adoption and dashed lines show No CC adoption.

CCS-only scenarios



(A) S1: CDR, flowshare tariff, CCS only

(B) S3: CDR, average-price tariff, CCS only

FIGURE H.1: Adoption-share trajectories in the CCS-only scenarios with CDR credits.

Figure H.1 shows that the CDR cases are robust in terms of final outcome. In both S1 and S3, all seeds converge towards quasi the same CCS-dominated state. The transition does not occur at the same moment for every seed, due to change in

emitter's revision moments. Due to this change in the revision moments, the critical mass can form at complete different moment, but the final outcome type remains the same. This supports the interpretation of Chapter 4 which is: in the CDR cases, stochasticity mainly affects the timing of the transition, not the existence of the high-CCS outcome.




FIGURE H.2: Adoption-share trajectories in the CCS-only scenarios without CDR credits.

Figure H.2 shows that the No CDR cases are also robust in terms of final outcomes. In S5, all seeds move towards quasi the same CCS-dominated outcome. But, on the contrary, the transition timing vary like in the CDR cases. S7 behaves differently. CCS remains marginal across seeds and No CC remains dominant. This confirms that, under average pricing and without CDR credits, the model does not reliably activate the CCS infrastructure feedback from the zero-capture initial condition.

CCU-enabled scenarios

The same seed-dependency check is applied to the CCU-enabled scenarios. These runs include a third strategy, so the seed can affect both the timing of CCS emergence and the role of CCU.

Figure H.3 confirms the robustness of the CDR cases when CCU is available. In S2 and S4, seeds may change the duration and intensity of the transitional CCU phase, but they do not change the broad final outcome. The system still converges towards a CCS-dominated state. This is consistent with the interpretation in Section 4.1.1. When CCS crosses the infrastructure threshold, CCU does not remain a separate dominant final pathway under CDR credits.

Figure H.4 also confirms the robustness of the No CDR cases. In S6, the broad outcome remains a lower-volume CCS outcome with only residual CCU. Regarding the transition, it is a bit more sensitive to the stochasticity than in the previous scenarios. Finally, the S8 scenario also demonstrates a high robustness across different seeds.

H. SEED-DEPENDENCY STUDY




FIGURE H.3: Seed-dependency of adoption-share trajectories in the CCU-enabled scenarios with CDR credits.




FIGURE H.4: Seed-dependency of adoption-share trajectories in the CCU-enabled scenarios without CDR credits.


Overall, the seed-dependency study supports the main interpretation of Chapter 4. The stochastic seed affects the timing of tipping points, the smoothness of adoption trajectories, and sometimes the really exact final adoption level. However, the final outcomes remains quasi the same, hence the scenario classification don't change with different seeds.

Appendix I

CCU delivery-flow maps

This appendix shows the final CCU delivery-flow maps for two runs in which CCU reaches a big final adoption share. The first map represents scenario S8 at the reference ETS price of 150 EUR/tCO₂. The second map shows S2 at an ETS price of 75 EUR/tCO₂.

The maps show the spatial matching between CCU participants and CO₂ demand nodes. Green points represent emitters participating in CCU, red points represent sites adopting CCS or No CC, turquoise circles indicate CO₂ demand points, and turquoise lines indicate positive emitter–demand delivery connections. These figures complement the discussion in Sections 4.1.5.



(A) Final CCU emitter–demand delivery connections in scenario S8 at the reference ETS price of 150 EUR/tCO₂.

(B) Final CCU emitter–demand delivery connections in scenario S2 at an ETS price of 75 EUR/tCO₂.

FIGURE I.1: CCU delivery-flow maps

Appendix J

Additional CCS network maps

This appendix reports larger versions of the CCS network maps used to compare the centralised CCNO reference with the decentralised ABM-EGT-CCNO outcomes.




FIGURE J.1: Supplementary CCS network maps.




FIGURE J.2: Supplementary CCS network map for the decentralised ABM-EGT-CCNO outcome S5 under the No-CDR setting.

Bibliography

- [1] AB InBev. Belgium locations. AB InBev Europe Careers, 2026. Accessed: 2026-05-07.
- [2] M. Akhurst, P. Brownsort, N. Burnside, R. S. Haszeldine, G. Johnson, C. Littlecott, P. Parmiter, V. Scott, G. Sim, and J. Stewart. SCCS Recommendations and Conference 2013 Executive Summary: Unlocking North Sea CO₂ Storage for Europe: Practical Actions for the Next Five Years. Technical report, Scottish Carbon Capture and Storage, 2013.
- [3] Argus Media. Argus biomass markets. Technical report, Argus Media, 2025.
- [4] W. B. Arthur. Competing technologies, increasing returns, and lock-in by historical events. *The Economic Journal*, 99(394):116–131, 1989.
- [5] M. Ashkavand, M. Scheffler, W. Heineken, M. D. Solomon, and T. Birth-Reichert. Techno-economic assessment of liquefied CO₂ transport via trucking. *International Journal of Greenhouse Gas Control*, 147:104491, 2025.
- [6] V. Auletta, D. Ferraioli, F. Pasquale, P. Penna, and G. Persiano. Convergence to equilibrium of logit dynamics for strategic games. *Algorithmica*, 76(1):110–142, 2016.
- [7] L. E. Blume. The statistical mechanics of strategic interaction. *Games and Economic Behavior*, 5(3):387–424, 1993.
- [8] Business in Antwerp. Europe’s largest integrated chemical cluster: A hub of innovation and sustainability. Business in Antwerp, 2026. Accessed: 2026-05-07.
- [9] Chemelot. Chemelot climate neutral by 2050. Chemelot, 2020. Accessed: 2026-05-07.
- [10] Chemiepark Marl. The marl chemical park. Chemiepark Marl, 2026. Accessed: 2026-05-07.
- [11] CHEMPARK. Facts. CHEMPARK, 2026. Accessed: 2026-05-07.
- [12] City of Ghent and North Sea Port. Towards a new carbon-neutral economy in the ghent area of north sea port: Exploratory study for a ccu hub. Technical report, City of Ghent, 2019. Accessed: 2026-05-07.

- [13] CO2 Value Europe. The contribution of carbon capture and utilisation (ccu) towards climate neutrality in europe: An eu roadmap for ccu by 2050. Technical report, CO2 Value Europe, 2024. Accessed: 2026-05-07.
- [14] Coca-Cola Europacific Partners. 2024 cdp corporate questionnaire. Technical report, Coca-Cola Europacific Partners, 2024. Accessed: 2026-05-07.
- [15] CURRENTA. Company. CURRENTA, 2026. Accessed: 2026-05-07.
- [16] M. Cătuți and D. Vangenechten. The european commission’s industrial carbon management strategy provides technocratic answers to an unfinished political debate.
- [17] De Tijd. Zonder basf komt ook CO₂-afvang van andere chemiebedrijven in gedrang.
- [18] European Biogas Association. Biogenic co2 from biogases. Technical report, European Biogas Association, 2025. Includes European CO2 market data based on S&P Global Commodity Insights (2024). Accessed: 2026-05-07.
- [19] European Commission. Just transition platform project fiche: Chemelot circular hub. Technical report, European Commission, 2023. Accessed: 2026-05-07.
- [20] European Commission. Industrial Carbon Management Strategy. Technical Report Comm. 52024DC0062, European Commission, 2024.
- [21] European Commission, Joint Research Centre. AIDRES - Advancing industrial decarbonization by assessing the future use of renewable energies in industrial processes. JRC Data Catalogue, 2026. Accessed: 2026-05-11.
- [22] Evonik. Marl. Evonik Worldwide Locations, 2026. Accessed: 2026-05-07.
- [23] G. Faber, A. W. Ruttinger, T. Strunge, T. Langhorst, A. Zimmermann, M. van der Hulst, F. Bensebaa, S. M. Moni, L. Tao, G. Guillén-Gosálbez, S. Suh, S. McCord, P. Styring, and V. Sick. Adapting technology learning curves for prospective techno-economic and life cycle assessments of emerging carbon capture and utilization pathways. *Frontiers in Climate*, 4:820261, 2022.
- [24] Fluxys. Fluxys belgium, north sea port and arcelormittal belgium receive eu funding for ghent carbon hub study. Fluxys Press Release, 2022. Accessed: 2026-05-07.
- [25] Glastuinbouw Nederland. Co2-voorziening. Glastuinbouw Nederland, 2026. Accessed: 2026-05-07.
- [26] Glastuinbouw Nederland. Greenports. Glastuinbouw Nederland, 2026. Accessed: 2026-05-07.

- [27] R. Golombek, M. Greaker, S. Kverndokk, and L. Ma. The transition to carbon capture and storage technologies. Technical Report Working Paper 9047, CESifo, 2021.
- [28] Greenports Nederland. Regionale greenports. Greenports Nederland, 2026. Accessed: 2026-05-07.
- [29] V. Grimm, U. Berger, F. Bastiansen, S. Eliassen, V. Ginot, J. Giske, J. Goss-Custard, T. Grand, S. K. Heinz, G. Huse, A. Huth, J. U. Jepsen, C. Jørgensen, W. M. Mooij, B. Müller, G. Pe’er, C. Piou, S. F. Railsback, A. M. Robbins, M. M. Robbins, E. Rossmanith, N. Rüger, E. Strand, S. Souissi, R. A. Stillman, R. Vabø, U. Visser, and D. L. DeAngelis. A standard protocol for describing individual-based and agent-based models. *Ecological Modelling*, 198(1–2):115–126, 2006.
- [30] V. Grimm, U. Berger, D. L. DeAngelis, J. G. Polhill, J. Giske, and S. F. Railsback. The odd protocol: A review and first update. *Ecological Modelling*, 221(23):2760–2768, 2010.
- [31] Heineken Nederland. Onze brouwerijen. Heineken Nederland, 2026. Accessed: 2026-05-07.
- [32] S.-H. Hwang and L. Rey-Bellet. Positive feedback in coordination games: Stochastic evolutionary dynamics and the logit choice rule. *Games and Economic Behavior*, 126:355–373, 2021.
- [33] International Association of Oil & Gas Producers (IOGP). The potential for CCS and CCU in europe. Technical report, IOGP, 2019. Madrid Forum report.
- [34] IPCC. Good practice guidance for land use, land-use change and forestry. Technical report, Intergovernmental Panel on Climate Change, 2003.
- [35] IPCC. 2006 ipcc guidelines for national greenhouse gas inventories: Volume 2, chapter 6, reference approach. Technical report, Intergovernmental Panel on Climate Change, 2006.
- [36] A. Itul, A. Diaz Rincon, O. D. Eulaerts, A. Georgakaki, M. Grabowska, Z. Kapetaki, E. Ince, S. Letout, A. Kuokkanen, A. Mountraki, D. Shtjefni, and M. Jaxa-Rozen. Clean energy technology observatory: Carbon capture utilisation and storage in the european union – 2023 status report on technology development, trends, value chains and markets. Technical report, Publications Office of the European Union, Luxembourg, 2023. JRC134999. Accessed: 2026-05-07.
- [37] L. R. Izquierdo, S. S. Izquierdo, and W. H. Sandholm. *Agent-Based Evolutionary Game Dynamics: A Guide*. 2024. Book.
- [38] Kas als Energiebron. Co2-voorziening. Kas als Energiebron, 2026. Accessed: 2026-05-07.

- [39] J.-S. Lee, T. Filatova, A. Ligmann-Zielinska, B. Hassani-Mahmooui, F. Stonedahl, I. Lorscheid, A. Voinov, J. G. Polhill, Z. Sun, and D. C. Parker. The complexities of agent-based modeling output analysis. *Journal of Artificial Societies and Social Simulation*, 18(4):4, 2015.
- [40] K. Lin, N. Wei, D. Jiang, and Y. Zhang. Progress and perspectives in source-sink matching for CCUS: A critical review. *Renewable and Sustainable Energy Reviews*, 226:116313, 2026.
- [41] A. Mahler and E. M. Rogers. The diffusion of interactive communication innovations and the critical mass: the adoption of telecommunications services by german banks. *Telecommunications Policy*, 23(10–11):719–740, 1999.
- [42] MarketDataForecast. Europe carbon dioxide market size, share, trends and growth forecast report, 2025–2033. MarketDataForecast, 2025. Commercial market report used only as a short-to-medium-term growth proxy. Accessed: 2026-05-07.
- [43] C. G. Martinez et al. Clean energy technology observatory: Carbon capture, utilisation and storage in the european union – 2024 status report on technology development, trends, value chains and markets. Technical report, Publications Office of the European Union, Luxembourg, 2024. JRC139285. Accessed: 2026-05-07.
- [44] R. D. McKelvey and T. R. Palfrey. Quantal response equilibria for normal form games. *Games and Economic Behavior*, 10(1):6–38, 1995.
- [45] S. Nie, G. Cai, Y. Li, Y. Chen, R. Bai, L. Gao, and X. Chen. To adopt CCU technology or not? an evolutionary game between local governments and coal-fired power plants. *Sustainability*, 14(8):4768, 2022.
- [46] NRW.Global Business. Business location nrw: Europe’s heartbeat. NRW.Global Business, 2026. Accessed: 2026-05-07.
- [47] Port of Antwerp-Bruges. Chemicals. Port of Antwerp-Bruges, 2026. Accessed: 2026-05-07.
- [48] Port of Moerdijk. Industrial park moerdijk. Port of Moerdijk, 2026. Accessed: 2026-05-07.
- [49] Port of Rotterdam. Storage and use of co2. Port of Rotterdam, 2026. Accessed: 2026-05-07.
- [50] Porthos. Co2 reduction through storage under the north sea. Porthos, 2026. Accessed: 2026-05-07.
- [51] N. Poulsen, S. Holloway, F. Neele, N. A. Smith, and K. Kirk. *CO2StoP Final Report. Assessment of CO2 Storage Potential in Europe. European Commission Contract No ENER/C1/154-2011-SI2.611598*, volume 2014 of *Danmarks og Grønlands Geologiske Undersøgelse Rapport*. GEUS, Copenhagen, 2014.

- [52] M. A. Rahmanta, B. M. Sopha, and A. M. Sri Asih. Modeling and optimization of carbon capture, utilization, and storage networks: A review of mathematical approaches. *Energy Reports*, 15:109162, 2026.
- [53] S. F. Railsback and V. Grimm. *Agent-Based and Individual-Based Modeling: A Practical Introduction*. Princeton University Press, 2 edition, 2019.
- [54] E. S. Rubin, I. M. L. Azevedo, P. Jaramillo, and S. Yeh. A review of learning rates for electricity supply technologies. *Energy Policy*, 86:198–218, 2015.
- [55] E. S. Rubin, S. Yeh, M. Antes, M. Berkenpas, and J. Davison. Use of experience curves to estimate the future cost of power plants with CO₂ capture. *International Journal of Greenhouse Gas Control*, 1(2):188–197, 2007.
- [56] D. Schüwer, K. Arnold, K. Biege, et al. Co₂ reuse nrw: Evaluating gas sources, demand and utilization for co₂ and h₂ within the north rhine-westphalia area with respect to gas qualities. Technical report, Wuppertal Institute for Climate, Environment and Energy, 2015. Accessed: 2026-05-07.
- [57] A. Sneyd, S. Chan, A. Mitchell, J. Howes, and S. Baltac. CO₂ Transport for e-Fuels Production in Europe. Technical report, Transport & Environment, 2025.
- [58] K. D. P. Storrs, I. Lyhne, and R. Drustrup. A comprehensive framework for feasibility of CCUS deployment. *International Journal of Greenhouse Gas Control*, 125:103878, 2023.
- [59] G. van der Laan and A. F. Tieman. Evolutionary game theory and the modeling of economic behavior. *De Economist*, 146(1):59–89, 1998.
- [60] N. J. A. van der Velden and P. X. Smit. Co₂-behoefte glastuinbouw 2030. Technical report, Wageningen Economic Research, 2019. Accessed: 2026-05-07.
- [61] F. Verbist, J. Meus, J. A. Moncada, P. Valkering, and E. Delarue. Implications of the EU ETS on the level-playing field between CCS & CCU. *International Journal of Greenhouse Gas Control*, 136:104165, 2024.
- [62] F. Verbist, J. Meus, P. Valkering, and E. Delarue. Unravelling co₂ value chain participation under negative emission pricing and eu industry relocation, 2024. Supplementary Information (AIDRES, Table S4).
- [63] J. W. Weibull. An introduction to evolutionary game theory. Technical Report Working Paper No. 347, The Industrial Institute for Economic and Social Research, 1992.
- [64] World Bank. GDP (current US\$). World Bank Open Data, 2025. Accessed: 2026-05-07.
- [65] B. Yu, J. Fu, and Y. Dai. Multi-agent simulation of policies driving CCS technology in the cement industry. *Energy Policy*, 199:114527, 2025.

- [66] Zero Emissions Platform (ZEP). Building a CO₂ transport infrastructure for europe. Technical report, Zero Emissions Platform, 2011.
- [67] T. Zhao and Z. Liu. A novel analysis of CCS technology adoption: An evolutionary game model. *Energy*, 189:116352, 2019.

Code of conduct and transparency statement on the use of GenAI for KU Leuven-students (academic year 2025–2026)

Generative AI (GenAI) assistance tools can be used to generate various types of content, including text, images, code, video, music, or combinations thereof. Common examples of such tools include ChatGPT, Google Gemini, Microsoft Copilot, Midjourney, Claude.ai, Perplexity.ai, and DALL·E, among others.

This code of conduct is a tool that helps students to be transparent about the use of GenAI and fits within the university's principles on academic integrity.

Important guidelines and remarks

- **Sensitive or personal data:** Some GenAI tools protect your input and sensitive or personal data better than others. There is often no transparency on what the owners of the AI applications do with the data entered. Therefore, **do not enter sensitive or personal data in free GenAI tools**. More info about what sensitive and personal data are, is described in the [three confidentiality levels of the KU Leuven data classification model](#) (non-confidential, confidential, strictly confidential). For confidential data you should preferably use **M365 Copilot Chat with your KU Leuven account**. If you use another GenAI tool, you must be **absolutely certain** it does **not store or reuse the data you enter**. **Try to avoid entering these data. Do so only if strictly necessary, and only to the extent required. For personal data, work with anonymous or pseudonymized data.** Additionally, in case of strictly confidential data this should first be discussed with the teaching staff of the course or your thesis supervisor.
- **Copyrighted materials:** For lawfully obtained copyrighted material you should preferably use **M365 Copilot Chat with your KU Leuven**

account. If you use another GenAI tool, you must be **absolutely certain** it does **not store or reuse the data you enter**.

- GenAI assistance may not be used for data or topics covered by a **Non-Disclosure Agreement (NDA)**. Even in the absence of an NDA, certain information may still need to be treated as confidential—for example, due to regulatory requirements or the risk of significant harm to the university if disclosed. In case of doubt, check with your teaching staff or supervisor.
- If your master’s thesis is under **embargo**, you should first discuss with your supervisor whether the use of GenAI is permitted.
- Before using a GenAI tool, always consider whether its use is responsible, including from a **sustainability perspective** (e.g. when using GenAI as a search engine, language assistant, ...).
- **Take a scientific and critical attitude** when interacting with GenAI assistance and interpreting its output, that may not always be correct.
- As a student you are responsible for complying with Article 84 of the Regulations on Education and Examinations: your report or thesis should reflect your own knowledge, understanding and skills. Be aware that plagiarism rules also apply to (work that is the result of) the use of GenAI assistance tools.

Exam Regulations Article 84: *“Every conduct individual students display with which they (partially) inhibit or attempt to inhibit a correct judgement of their own knowledge, understanding and/or skills or those of other students, is considered an irregularity which may result in a suitable penalty. A special type of irregularity is plagiarism, i.e. copying the work (ideas, texts, structures, designs, images, plans, codes , ...) of others or prior personal work in an exact or slightly modified way without adequately acknowledging the sources. Every possession of prohibited resources during an examination (see article 65) is considered an irregularity.”*

- In order to maintain academic integrity and avoid plagiarism, **more information about being transparent on the use of GenAI assistance and about correctly citing and referencing GenAI** can be found on this website for students (**Dutch/English**).
- **Additional reading: KU Leuven guidelines on responsible use of Generative AI tools, and other information** (**Dutch/English**)

A few final words

If you are uncertain whether or not you should declare your use of GenAI tools, we suggest that you discuss this with your instructor or supervisor. It is always safer to declare GenAI use, even when it is not strictly required. However, declaring GenAI use does not entail that its use is

allowed; the right column in the table below provides more detailed instructions in this regard (code of conduct).

Moreover, advanced AI tools are evolving rapidly, and their capabilities have expanded significantly in a short period of time. As a result, we do not yet have all the answers about their responsible use. Finally, it is important to follow-up on the most recent evolutions in AI technologies, to have an open mind but also to be a bit cautious, to communicate with instructors, teaching assistants, supervisors and peers, to be as transparent as we can, and to learn together as we move along.

Student name: Florian Verwilghen

Student number: R1044952

Please indicate with "X" whether it relates to a course assignment or to the Bachelor's or Master's thesis:

☐ This form is related to a **course assignment**.

Course name: **Course number:**

This form is related to my ☐ **Bachelor's** or ☒ **Master's thesis**.

Title Master's thesis: Understanding the Emergence of CO₂ Value Chains: An Agent-Based Evolutionary Game Approach Coupled with CO₂ Network Optimisation

Supervisor: Prof. Erik Delarue and Dr. Jorge Andrés Moncada

Daily supervisor: Flore Verbist, Dr. Jorge Andrés Moncada, Dr. Veis Karbassi

Please indicate with "X":

☐ I did not use any GenAI assistance tool.

☒ I did use GenAI Assistance. In this case **specify which ones** (e.g. ChatGPT, M365 Copilot, ...): Chatgpt, Grammarly, DeepL, Semantic Scholar

GenAI assistance used as/for:	Name of the GenAI tool(s) used If helpful, also describe in which way you were using GenAI related to what is specified as code of conduct.	Code of conduct: For each of the categories below, always take into account the important guidelines and remarks mentioned above (e.g. copyrighted data, sensitive or personal data,...).
As a language assistant for reviewing or improving texts I wrote myself	<i>Grammarly, DeepL, Chatgpt (But almost only Grammarly and DeepL)</i>	This use is similar to using spelling and grammar check tools. In general, you do not have to refer to such kind of GenAI use in the text. However, be careful: When using GenAI tools on texts you did not write yourself to improve the text, you have to refer to the original source or author, otherwise you are committing plagiarism and thus an irregularity.
As a paraphrasing tool	<i>Grammarly</i>	This use is allowed except when it is prohibited by the teacher or the program of study. You may paraphrase your own text or texts by an author other than yourself and take inspiration from what a GenAI tool or another tool suggests (unless it is not allowed). In general, you do not have to refer to such kind of GenAI use or other paraphrasing tools in the text. However, be careful: If it entails text by an author other than yourself, you are not allowed to include that paraphrased text without reference to the original source or author. Without such reference, you would be committing plagiarism and thus an irregularity.

For translation aid to improve texts I wrote myself or to better understand text from others	<i>DeepL, Chat-GPT</i>	This use is allowed except when it is prohibited by the teacher or the program of study. It is similar to using translation tools (Google translate, DeepL, ...). In general, you do not have to refer to such kind of GenAI use in the text. However, be careful: • You are not allowed to include that translated text without reference to the original source or author. Without such reference, you would be committing plagiarism and thus an irregularity. • Always check the translated text for correctness and meaning.
As a search engine to get information on a topic or to search for existing research on the topic	<i>Semantic Scholar, Chatgpt</i>	This use is similar to e.g. a Google search or checking Wikipedia. If you write your own text based on this information, you do not have to refer to the use of GenAI in the text. You only have to refer to the existing research and references you have checked and used (without such references you would be committing plagiarism and thus committing an irregularity). However, be careful: • Be aware that the output of the GenAI tool cannot be guaranteed as a 100% reliable source of information. The output may not be entirely correct and/or be limited due to the databases it uses. Moreover, knowledge evolves and may change over time; therefore, the database of the GenAI tool may not be up to date. Therefore, verify the information and do not just copy-paste it as you should understand and critically process everything you are writing.

For literature search	<i>Semantic Scholar</i>	This use is comparable to e.g. a Google Scholar search. You do not have to refer to such kind of GenAI use; you only have to refer to the literature references you have checked and used (without such references you would be committing plagiarism and thus committing an irregularity). However, be careful: • Be aware that the search output is restricted to the database the GenAI tool is built on. After this initial search, look for scientific sources and conduct your own analysis of the source documents. Interpret, analyse and process the information you obtained; verify it and do not just copy-paste it as you should understand and critically process everything you are writing. • Be aware that some GenAI tools may output no or wrong references. As a student you are responsible for further checking and verifying the absence or correctness of references; do not just copy-paste it.
------------------------------	-------------------------	---

For generating programming code	<i>Chatgpt : Not to generate code directly, but to debug with messages from the terminal, and explain how solve the problem in the concerned files. Also used to help with the creation of some graphs and more precisely the selection of the best-looking adoption trajectories and the colours' Matplotlib codes in the Simplex plots (See Chapter 4).</i>	Use of GenAI for coding is allowed except when it is prohibited by the teacher or the program of study. If used for coding, correctly mention the use of GenAI assistance in accordance with the instructions on the page on being transparent about the use of GenAI.
For generating new ideas (research)	/	This use of GenAI is allowed except when it is prohibited by the teacher or the program of study. Correctly mention the use of GenAI assistance in accordance with the instructions on the page on being transparent about the use of GenAI. Be careful: • Further verify in this case whether the idea is novel or not. It is likely that it is related to existing work. If so, that existing work should be correctly referenced in the text (without such reference you would be committing plagiarism and thus committing an irregularity).

For generating synthetic data	/	Use of GenAI for generating synthetic data is allowed, provided that it is methodologically and ethically justifiable, except when it is prohibited by the teacher or the program of study. Always correctly mention the use of GenAI assistance in accordance with the instructions on the page on being transparent about the use of GenAI. Be careful: • Always carefully evaluate the generated synthetic data for quality and possible bias since the output is highly dependent on the quality of the data on which the models are trained.
For generating blocks of text (other than the allowed use without referencing mentioned above)	Chatgpt : Not to generate directly blocks of text. For some passages, I gave text that I wrote myself and asked to help bringing more structure to it.	According to Article 84 of the Regulations on Education and Examinations your text should allow to correctly and properly assess your own knowledge, understanding and skills. Therefore, inserting blocks of text without quotes and a reference to GenAI assistance in your work is not allowed. Be careful: • If it is really needed to insert a block of text from a GenAI tool, for instance because of the nature of your assignment, mention it as a citation by using quotes and correctly mention the use of GenAI assistance in accordance with the instructions page on being transparent about the use of GenAI. • However, in general, such GenAI use should be kept to an absolute minimum; you should always check the original sources.

For generating visuals, video or audio	/	Use of GenAI for generating visuals, audio or video is allowed except when it is prohibited by the teacher or the program of study. Be careful: • If used, refer to GenAI for visuals according to the usual referencing style, following the instructions on the page on referencing GenAI. • If you work with existing visuals, audio or video, that existing work should be correctly referenced in the text (without such reference you would be committing plagiarism and thus an irregularity). • Explain the usage in the methods section (if there is one) and optionally attach (or link to) the prompts with the full output (or history).
Other use (specify here; this may also include a combination of types of use mentioned above):	<i>Chatgpt: I used chatgpt to help with the LaTeX encoding of all my sources in my bibliography.</i>	To make sure other use of GenAI is allowed within the course or thesis, and if so, the conditions that may apply, contact the teaching staff of the course or the supervisor of the thesis beforehand and explain the intended GenAI use. Also inform the programme director. Motivate how you would comply with Article 84 of the exam regulations. Explain the use and the added value of the AI tool you consider to use and how it is in accordance with the assignment or thesis and the KU Leuven guidelines on responsible use of Generative AI tools. Depending on the kind of GenAI use, it may be needed to properly reference it in the text, in accordance with the instructions page on being transparent about the use of GenAI.