

The effect of AI on the wage inequality in the Dutch labour market

An agent-based exploration of skill-biased automation

MSc Engineering and Policy Analysis

Bjarne Timmer

Delft University of Technology

The effect of AI on the wage inequality in the Dutch labour market

An agent-based exploration of skill-biased automation

by

Bjarne Egbertus Jacobus Timmer

Master thesis submitted to Delft University of Technology in partial fulfilment of the requirements for the degree of

Master of Science Engineering and Policy Analysis
Faculty of Technology, Policy and Management
Delft University of Technology

Student name: Bjarne Egbertus Jacobus Timmer
Student number: 5091756
Thesis defence date: 26 August 2026

Graduation committee

Dr. S.T.H. Storm, TU Delft, Economics of Technology and Innovation	Chair & first supervisor
Prof. dr. M.E. Warnier, TU Delft, Systems Engineering	Second supervisor
N.L. Plaisier MSc, Social and Economic Council (SER)	External advisor

This research was carried out during an internship at the Social and Economic Council (SER), The Hague.

Cover image: AI-generated with ChatGPT 5.5 High (OpenAI), June 2026.



Executive summary

Within the OECD, the Dutch labour market occupies a distinctive position: with relatively low unemployment and high labour force participation, it is segmented into a relatively large flexible shell, with roughly forty percent of workers in (temporary) flexible arrangements and the highest flexibility rate of the EU-27, alongside a segment of permanent employees with relatively strong dismissal protection. In the past, the adjustment burden of external shocks has fallen disproportionately on the first group: during the corona crisis the fall in employment sat almost entirely in flexible contracts, with more than one in ten flexible workers losing work in the worst quarters of 2020 against roughly one in fifty permanent employees (CBS, 2025a). It can be expected that generative AI will have significant (distributional) impacts on the Dutch labour market in the next decade, affecting flexible and permanent workers as well as low-skilled and high-skilled workers differently. Importantly, unlike earlier waves of automation, generative AI may lead to the automation of non-routine, cognitive tasks. How fast will Dutch firms deploy the new genAI tools, given that the costs of AI is declining and given the labour market institutions that constrain their choices? Which (routine or non-routine) tasks will be replaced by AI and which (flexible versus permanent) jobs will be affected? How will genAI affect wage growth and wage inequality given the institutional conditions characterising the Dutch labour market?

Existing economic models do not provide an answer: equilibrium models miss the heterogeneity of firms, bounded rationality and the Dutch institutional setting, while existing agent-based models treat the adoption decision as exogenous or as an R&D lottery and do not model the labour-market institutions. The combination of a model analysis allowing for an endogenous, boundedly rational AI adoption within the institutional context of the Dutch labour market is missing.

Addressing this knowledge gap is socially urgent. The Netherlands Scientific Council for Government Policy (“Wetenschappelijke Raad voor het Regeringsbeleid”) compares AI to the steam engine and electricity, the Social and Economic Council (“Sociaal-Economische Raad”) calls for scenario experiments into distributional effects and policy instruments, and the Dutch central bank (“De Nederlandsche Bank”) points to the persistently low productivity growth in a tight labour market.

This research builds an agent-based model in Python (Mesa), building on an existing task-based agent-based model (Upreti & Sridhar, 2024), and extends it with five elements: an explicit investment decision based on net present value (NPV), dismissal protection and the chain regulation (ketenregeling), wage formation with collective-agreement-like rigidity, a declining AI price curve, and a Keynesian feedback from unemployment to demand. Automation is modelled throughout as all-or-nothing substitution, so the model captures AI’s displacement of labour but not its augmenting or task-creating effects; the inequality it reports is therefore an upper bound. The model is a scenario-exploring instrument rather than a point forecast: it traces directions, feedbacks and sensitivities around a calibrated baseline.

Together the eight experiments produce a consistent picture. First, making the adoption decision endogenous tempers automation and inequality considerably relative to the myopic cost-per-unit logic: the expected wage decline through displacement, the Keynesian demand shortfall and a profitability threshold work together as a brake. A decomposition of this gap attributes the large majority of it to the forward-looking investment logic itself and only a small part to the Dutch employment-protection institutions. The most important finding, however, is that the relative price of AI dominates the outcomes, more than any institution: both the level and the trajectory of AI costs determine not just how far automation reaches but which qualitative future unfolds, from near-complete automation under falling costs to a near-standstill under rising costs. Of the institutions studied, only the minimum wage has a clear, but double-edged effect: it compresses the wage gap for those who keep their jobs, but accelerates the displacement of low-skilled workers once it becomes binding. The typically Dutch contract institutions, the transition payment and the chain regulation, act as a small but persistent brake: they slow automation early and modestly lower its long-run level, but raising the transition payment above the statutory level adds nothing, and the contract mix barely matters.

Translated to the real Dutch labour market, the mechanics suggest that AI-driven displacement is unlikely to arrive as waves of dismissals of protected staff: in the model, firms wait for ordinary turnover and automate the vacated task. In reality this would surface as vacancies quietly left unfilled, hiring freezes for starters and non-renewal of flexible contracts, placing the adjustment burden once again on the flexible shell and on labour-market entrants. Dismissal protection buys valuable transition time, the window for retraining and task redesign, but making dismissal more expensive should not be expected to slow automation further. The minimum wage and automation policy turn out to be the same policy in segments where AI is a credible substitute: every increase of the floor strengthens the business case for automating exactly the tasks it protects. And because in reality a large share of generative-AI deployment augments rather than replaces workers, a channel this model deliberately excludes, the reported inequality effects are upper bounds, and the distributional question may shift from who loses their job to who captures the productivity surplus.

The decisive variable, however, is also the most uncertain one. This research added a forward-looking adoption decision, but that decision is made under certainty, and the AI-cost path it responds to is precisely the quantity nobody can forecast: inference prices are falling, yet the firm-side cost of integration and supervision is not falling at the same pace, and early real-world deployments already show AI sometimes costing more than the workers it replaced. Policy should therefore be robust across price futures rather than optimised for the cheap-AI scenario: price-oriented instruments such as a tax on AI use are in principle the most powerful lever the model identifies, but their design must distinguish substitution from augmentation to avoid taxing the beneficial channel. All conclusions are exploratory and hold in the vicinity of the calibrated baseline; their value lies in showing which mechanisms matter and which uncertainties dominate, above all the price path of AI.

Acknowledgements

This thesis marks the end of my MSc in Engineering and Policy Analysis at the Faculty of Technology, Policy and Management of Delft University of Technology.

The choice of subject was, in many ways, an easy one. My fascination with artificial intelligence (AI) goes back years, through a minor in Artificial Intelligence during my bachelor and the countless hours I have spent following its latest developments and tinkering with new models and tools, as my friends and family can readily confirm. My interest in it pulls in two directions at once. I find it genuinely remarkable what these systems can already do, while the same relentless pace leaves me increasingly uneasy about the impact they may have on society. It was precisely that tension that drew me to the subject of this thesis: what AI-driven automation means for work and wages. The elective space in the second year of my master's then let me act on a long-held wish to understand economics more deeply, through the economics and finance track. Bringing these two interests together felt entirely natural, and this thesis is where they meet.

Looking back, what stays with me most is not the model or the results, but the people who made the work both possible and enjoyable. I owe my deepest thanks to my first supervisor and chair of my graduation committee, Dr. Servaas Storm. The many conversations we had about economic theory shaped this research more than anything else, and his willingness to think along, challenge my reasoning and spar with me on ideas pushed the work forward at every stage. I am especially grateful for his availability and dedication, down to answering my emails on weekends, which kept the project moving even when I got stuck. I am equally grateful to my second supervisor, Prof. dr. Martijn Warnier, for his guidance, his clear feedback and the perspective he brought to the work when I risked losing the bigger picture.

I carried out this research during my internship at the Social and Economic Council (SER), and I would like to thank the SER for an environment that was as welcoming and warm as it was instructive. In particular I want to thank Nora Plaisier MSc, my supervisor at the SER, for her support and encouragement throughout. I am also grateful to the many colleagues at the SER who took the time to discuss my work with me. It meant a great deal that I was able to contribute something of my own in return, which made the experience all the more valuable to me. My thanks also go to the Netherlands Bureau for Economic Policy Analysis (CPB) for an interview that gave me valuable modelling advice and helped set the direction of this study.

Finally, I want to thank my family, my friends and my girlfriend. Their patience, encouragement and the welcome distractions they offered carried me through the long days, and I could not have completed this without them. I owe them a special apology, and my thanks, for endlessly listening to AI news, facts and trivia they almost certainly never asked for.

Bjarne Timmer
The Hague, August 2026

Contents

Executive summary	i
Acknowledgements	iii
List of Figures	viii
List of Tables	xi
List of Abbreviations	xiii
1 Introduction	1
1.1 Research goal and questions	3
2 Literature and theoretical framework	4
2.1 From skills to tasks	4
2.2 AI as specific type of automation	4
2.3 Agent-based modelling and AI-adopting ABMs	5
2.4 The Dutch institutional context	6
2.4.1 Permanent vs. flexible contracts	6
2.4.2 Employment protection and severance costs (WAB)	6
2.4.3 Chain regulation (Ketenregeling) (WWZ)	6
2.4.4 CLAs and minimum wage	7
2.5 Answering SQ1: the decisive Dutch mechanisms	7
2.6 Knowledge gap	7
3 Conceptual model	8
3.1 Justification of the ABM approach	8
3.2 Model architecture	9
3.3 Agent interactions within one step	10
3.4 Step order	11
3.5 Components overview	12
3.6 C1 - Production	12
3.6.1 Task structure	12
3.6.2 Leontief production function	13
3.6.3 Task production	13
3.6.4 Human productivity	13
3.6.5 AI productivity	14
3.6.6 Unit labour costs	14
3.6.7 Average cost expansion criterion	15
3.7 C2 - Labour market	15
3.7.1 Exogenous separation	15
3.7.2 Matching	15
3.7.3 Employment and unemployment	16
3.8 C3 - Wage-setting	16
3.8.1 Wage curve	16
3.8.2 Partial wage adjustments	16
3.9 C4 - AI adoption	17
3.9.1 ULC mode (base)	17
3.9.2 NPV-modes	17
3.9.3 Bounded rationality	19
3.9.4 Reactive vs. Proactive adoption	19
3.9.5 Irreversibility	19

3.10	C5 - AI learning curve	20
3.11	C6 - Demand	20
3.11.1	Demand curve	20
3.11.2	Keynesian feedback loop	20
3.12	C7 - Employment protection	21
3.12.1	Contract types and tenure years	21
3.12.2	Severance costs	21
3.12.3	Chain regulation	21
3.13	M1 - Macro-feedback coupling	22
3.14	Answering SQ2: adapting and calibrating the model	22
3.15	Modelling and validation	23
4	Experimental design	24
4.1	Scenario-exploratory demarcation	24
4.2	Baseline configuration	24
4.3	Comparison with Upreti & Sridhar	25
4.4	Scenario axes	25
4.4.1	AI Cost trajectory	25
4.4.2	Comparative advantage flip	26
4.4.3	Institutional levers	26
4.5	Outcome variables	27
4.6	Replication and statistical approaches	28
4.7	Time horizon	28
4.8	Verification and validation	28
4.9	Experimental program	29
5	Results	31
5.1	Experiment 1 - Baseline (mean-field NPV) versus Upreti & Sridhar (ULC)	31
5.2	Experiment 2 - Decision logic: forward-looking NPV versus myopic ULC	33
5.3	Experiment 3 - Institutions: employment protection on/off	34
5.4	Experiment 4 - AI costs trajectory	35
5.5	Experiment 5 - Comparative advantage	36
5.6	Experiment 6 - Minimum wage	37
5.7	Experiment 7 - Severance payment	38
5.8	Experiment 8 - Transition probability	39
5.9	Across the experiments	40
6	Discussion	41
6.1	Summary of the main findings	41
6.2	Interpretation of the results	42
6.2.1	The relative price of AI shapes the outcomes	43
6.2.2	Automation is a level effect, not a growth effect	43
6.2.3	Comparative advantage sets how far automation reaches	44
6.2.4	The minimum wage as the sole binding institution	45
6.2.5	Why the contract institutions hardly bite	46
6.3	Reflection on the agent-based approach	46
6.4	The role of uncertainty	48
6.5	Limitations	49
6.5.1	Pattern calibration and baseline dependence	49
6.5.2	The inequality effects are an upper bound	49
6.5.3	Are jobs just a bundle of tasks?	50
6.5.4	The skill-task split is too coarse	51
6.5.5	Structural and functional simplifications	51
6.5.6	Limitations of the sensitivity analysis	52
6.5.7	What these limitations mean for reading the results	52
7	Conclusions	53
7.1	Answering the research questions	53

7.2	Implications for the Dutch labour market	54
7.2.1	How displacement would surface	54
7.2.2	Augmentation is the largest open question	55
7.2.3	What employment protection does and does not buy	56
7.2.4	The minimum wage: not binding yet	56
7.2.5	Taxing AI: the strongest instrument and the hardest to build	56
7.3	Academic implications	57
7.4	Recommendations for further research	57
7.5	Conclusion	58
	Statement on the use of AI	60
	References	61
A	Implementation	65
A.1	Software stack	65
A.2	Mapping of components to code	65
A.3	The simulation loop and the order of sub-steps	67
A.4	Warm-start: initialising a fully staffed economy	67
A.4.1	What the warm-start does	67
A.4.2	Why the warm-start fills every job	67
A.5	Calibration and default values	68
A.6	Streamlit-dashboard	69
B	Mean-field decision and complete pairwise comparisons	70
B.1	Comparison of the three NPV modes	71
B.2	Experiment 1	74
B.3	Experiment 2	77
B.4	Experiment 3	80
B.5	The skill-wage premium as an initialisation artefact	85
B.5.1	The full mechanism: why the institutional effect is small, saturating and persistent	86
B.5.2	Diagnostics for the mechanism behind the temporal pattern	88
B.6	Experiment 4	93
B.7	Experiment 5	96
B.8	Experiment 6	99
B.9	Experiment 7	102
B.10	Experiment 8	105
B.11	Illustration of the permutation test (Experiment 4)	109
C	Parameter calibration	113
C.1	Calibration strategy	113
C.2	Three structural decisions	113
C.2.1	Wage scale: euro-per-hour-equivalent units	113
C.2.2	AI cost scale	114
C.2.3	Discount rate	114
C.3	Steady-state	114
C.4	The investment-cost magnitude (I_{base})	115
C.5	The baseline parameter set	115
C.6	Known simplifications and caveats	118
D	Verification	119
D.1	Verification	119
D.2	Test specifications	119
D.3	Block 1 - Unit Tests	120
D.3.1	Test 1 - Leontief Output	120
D.3.2	Test 2 - NPV Calculation (All Four Modes)	120
D.3.3	Test 3 - Severance Pay	120
D.3.4	Test 4 - Wage Curve and Partial Adjustment (Half-Life)	120
D.3.5	Test 5 - AI Learning Curve (Exponential Decay)	120

D.4	Block 2 - Conservation Checks	121
D.4.1	Test 6 - Worker Conservation	121
D.4.2	Test 7 - Task Conservation	121
D.4.3	Test 8 - Cost Identity	121
D.5	Block 3 - Edge Cases	121
D.5.1	Test 9 - Trivially Cheap AI => Adoption to 100%	121
D.5.2	Test 10 - Prohibitively Expensive AI => Adoption to 0%	121
D.5.3	Test 11 - <code>employment_protection = False</code> => Institutional Effect = 0	121
D.5.4	Test 12 - <code>init_share_vast = 1</code> => Maximum Severance Component	121
E	Validation	122
E.1	Why pattern-oriented validation	122
E.2	Testable patterns	122
E.2.1	Pattern 1 - Profitability gate	123
E.2.2	Pattern 2 - Slower adoption rate for a more expensive AI	123
E.2.3	Pattern 3 - High severance payment slows automation down	123
E.2.4	Pattern 4 - Demand declines due to mass unemployment	123
E.2.5	Pattern 5 - Induced technical change for a higher minimum wage	123
E.2.6	Pattern 6 - Polarisation dependent on AIs comparative advantage	123
E.3	Experimental design validation	123
E.3.1	Time horizon	123
E.3.2	Monte Carlo replications	124
E.3.3	Scenario	124
E.4	Results per pattern	124
E.4.1	Pattern 1 - Profitability gate	124
E.4.2	Pattern 2 - Slower adoption rate for a more expensive AI	124
E.4.3	Pattern 3 - High severance payment slows automation down	125
E.4.4	Pattern 4 - Keynesian demand decline on mass unemployment	125
E.4.5	Pattern 5 - Induced technical change for a higher minimum wage	125
E.4.6	Pattern 6 - Polarisation dependent on AIs comparative advantage	127
E.5	Synthesis	127
F	Parameters	130
G	Sensitivity analysis	132
G.1	Setup and method	132
G.2	Gini-index	133
G.3	Skill-wage premium (<code>w_h/w_l</code>)	133
G.4	Employment rate for high skilled	133
G.5	Employment rate for low skilled	135
G.6	Limitations of the analysis	135
G.7	Complete parameter overview	135
H	Model assumptions	137
H.1	Type labels	137
H.2	A0 - Model-wide and scope assumptions	137
H.3	C1 - Production	138
H.4	C2 - Labour market	139
H.5	C3 - Wage setting	140
H.6	C4 - AI adoption	140
H.7	C5 - Learning curve	141
H.8	C6 - Demand	141
H.9	C7 - Employment protection	142
H.10	M1 - Macro feedback	143
H.11	Technical assumptions	143
H.12	Most influential assumptions	143

List of Figures

1	UML-class diagram of the model entities. Every box is a separate class with its state variables in the top box and its methods in the bottom box. Coloured labels link the methods to the seven model components (C1-C7), the class LabourMarketModel has the macro feedback loop (M1). The diamonds show the aggregation and composition and the multiplicities are at the ends of the links.	10
2	Agent-interactions within one step. The vertical lines are the entities of Figure 1, the LabourMarketModel orchestrates the sub-step order of step(). The continuous arrows are synchronous calls or information flows, dashed arrows are returns, and self-calls are internal calculations of the LabourMarketModel. The fourteen sub-steps are grouped in seven phases (A-G) and labelled per component (C1-C7). The panel below summarises the macro feedback loop (M1) with one-tick delay.	11
3	Gini-index, Skill-wage premium and Unemployment rate over time of the ULC configuration and Mean-field NPV. Average over 50 seeds with a ± 1 standard deviation band.	32
4	The market price over time of the ULC configuration and Mean-field NPV. Average over 50 seeds with a ± 1 standard deviation band.	32
5	Gini-index, skill-wage premium and unemployment rate over time of the ULC rule and the mean-field NPV with employment protection off (MF-off), both without institutions. Average over 50 seeds with a ± 1 standard-deviation band.	33
6	Gini-index, skill-wage premium and unemployment rate over time with employment protection on (baseline) and off, both in the mean-field NPV mode. Average over 50 seeds with a ± 1 standard-deviation band.	35
7	Gini-index, Skill-wage premium and Unemployment rate over time of the three cost trajectories (decreasing/baseline, flat, rising). Average over 50 seeds with a ± 1 standard deviation band.	35
8	Gini-index, Skill-wage premium and Unemployment rate over time of the four configurations (low, baseline, medium, high). Average over 50 seeds with a ± 1 standard deviation band.	37
9	Employment rate over time per skill group for the four configurations (low, baseline, medium, high). Average over 50 seeds with a ± 1 standard deviation band.	37
10	Gini-index, Skill-wage premium and Unemployment rate over time of the four configurations (base, no minimum wage, low minimum wage, high minimum wage). Average over 50 seeds with a ± 1 standard deviation band.	38
11	Gini-index, Skill-wage premium and Unemployment rate over time of the three configurations (base, no severance, high severance). Average over 50 seeds with a ± 1 standard deviation band.	39
12	Gini-index, skill-wage premium and unemployment rate over time of the five transition-probability configurations (base, low, no, high and complete conversion). Average over 50 seeds with a ± 1 standard deviation band.	39
13	Gini index over the 960-step horizon for the three NPV modes (naïve, adaptive, mean-field). Mean over 50 seeds with a ± 1 standard-deviation band.	72
14	Skill-wage premium over the 960-step horizon for the three NPV modes (naïve, adaptive, mean-field). Mean over 50 seeds with a ± 1 standard-deviation band.	73
15	Unemployment rate over the 960-step horizon for the three NPV modes (naïve, adaptive, mean-field). Mean over 50 seeds with a ± 1 standard-deviation band.	74
16	Gini index over the horizon, Experiment 1 (mean-field NPV baseline versus ULC). Mean over 50 seeds with a ± 1 standard-deviation band.	75

17	Skill-wage premium over the horizon, Experiment 1 (mean-field NPV baseline versus ULC). Mean over 50 seeds with a ± 1 standard-deviation band.	76
18	Unemployment rate over the horizon, Experiment 1 (mean-field NPV baseline versus ULC). Mean over 50 seeds with a ± 1 standard-deviation band.	77
19	Gini index over the horizon, decision-logic step (mean-field protection-off versus ULC). Mean over 50 seeds with a ± 1 standard-deviation band.	78
20	Skill-wage premium over the horizon, decision-logic step (mean-field protection-off versus ULC). Mean over 50 seeds with a ± 1 standard-deviation band.	79
21	Unemployment rate over the horizon, decision-logic step (mean-field protection-off versus ULC). Mean over 50 seeds with a ± 1 standard-deviation band.	80
22	Gini index over the horizon, Experiment 3 (employment protection on versus off). Mean over 50 seeds with a ± 1 standard-deviation band.	81
23	Skill-wage premium over the horizon, Experiment 3 (employment protection on versus off). Mean over 50 seeds with a ± 1 standard-deviation band.	82
24	Unemployment rate over the horizon, Experiment 3 (employment protection on versus off). Mean over 50 seeds with a ± 1 standard-deviation band.	83
25	Paired permutation test for the Gini index, Experiment 3 (employment protection off versus on). The upper panel shows the two configurations (mean ± 1 sd); the lower panel shows the paired difference (A-B) with its 95% simultaneous band, with the significant steps shaded. The significant mass lies at the start and the end of the horizon.	84
26	Paired permutation test for the skill-wage premium, Experiment 3 (employment protection off versus on). Panels as in Figure 25.	85
27	Average severance owed per task over the horizon, Experiment 3 (only defined when employment protection is on). It peaks early on the high-tenure warm-start cohort and then declines before flattening, so the severance brake fades rather than re-strengthening. Mean over 50 seeds with a ± 1 standard-deviation band.	89
28	Average tenure of permanent workers over the horizon, Experiment 3 (only defined when employment protection is on). Tenure stabilises around nine-and-a-half years because exogenous separations cap it, rather than accumulating without bound. Mean over 50 seeds with a ± 1 standard-deviation band.	90
29	Keynesian expansion-gate blocks per step, Experiment 3 (employment protection on versus off). The count rises steeply in the protection-off run while staying flat with protection, the signature of the demand-feedback loop closing the expansion gate. Mean over 50 seeds with a ± 1 standard-deviation band.	91
30	Demand level over the horizon, Experiment 3 (employment protection on versus off). The protection-off run diverges downward late as its lower wage bill feeds the demand intercept. Mean over 50 seeds with a ± 1 standard-deviation band.	92
31	Low-skill employment rate over the horizon, Experiment 3 (employment protection on versus off). The late divergence concentrates here, around 73% without protection against 79% with it, because the closed expansion gate suppresses re-hiring of the cheapest, most automatable workers. Mean over 50 seeds with a ± 1 standard-deviation band.	93
32	Gini index over the horizon, Experiment 4 (decreasing, flat and rising AI-cost trajectories). Mean over 50 seeds with a ± 1 standard-deviation band.	94
33	Skill-wage premium over the horizon, Experiment 4 (decreasing, flat and rising AI-cost trajectories). Mean over 50 seeds with a ± 1 standard-deviation band.	95
34	Unemployment rate over the horizon, Experiment 4 (decreasing, flat and rising AI-cost trajectories). Mean over 50 seeds with a ± 1 standard-deviation band.	96
35	Gini index over the horizon, Experiment 5 (non-routine multiplier 0.50, 0.85, 1.15, 1.50). Mean over 50 seeds with a ± 1 standard-deviation band.	97
36	Skill-wage premium over the horizon, Experiment 5 (non-routine multiplier 0.50, 0.85, 1.15, 1.50). Mean over 50 seeds with a ± 1 standard-deviation band.	98
37	Unemployment rate over the horizon, Experiment 5 (non-routine multiplier 0.50, 0.85, 1.15, 1.50). Mean over 50 seeds with a ± 1 standard-deviation band.	99
38	Gini index over the horizon, Experiment 6 (minimum wage 0, 10, 14.71, 20). Mean over 50 seeds with a ± 1 standard-deviation band.	100

39	Skill-wage premium over the horizon, Experiment 6 (minimum wage 0, 10, 14.71, 20). Mean over 50 seeds with a ± 1 standard-deviation band.	101
40	Unemployment rate over the horizon, Experiment 6 (minimum wage 0, 10, 14.71, 20). Mean over 50 seeds with a ± 1 standard-deviation band.	102
41	Gini index over the horizon, Experiment 7 (severance rate 0, 1/3, 1.0). Mean over 50 seeds with a ± 1 standard-deviation band.	103
42	Skill-wage premium over the horizon, Experiment 7 (severance rate 0, 1/3, 1.0). Mean over 50 seeds with a ± 1 standard-deviation band.	104
43	Unemployment rate over the horizon, Experiment 7 (severance rate 0, 1/3, 1.0). Mean over 50 seeds with a ± 1 standard-deviation band.	105
44	Gini index over the horizon, Experiment 8 (transition probability 0, 0.25, 0.5, 0.75, 1.0). Mean over 50 seeds with a ± 1 standard-deviation band.	107
45	Skill-wage premium over the horizon, Experiment 8 (transition probability 0, 0.25, 0.5, 0.75, 1.0). Mean over 50 seeds with a ± 1 standard-deviation band.	108
46	Unemployment rate over the horizon, Experiment 8 (transition probability 0, 0.25, 0.5, 0.75, 1.0). Mean over 50 seeds with a ± 1 standard-deviation band.	109
47	Permutation test, decreasing versus flat AI-cost trajectory for unemployment rate. $p < 0.001$, significant on 99% of the horizon, $d_z = 3.47$	110
48	Permutation test, decreasing versus flat AI-cost trajectory for the skill-wage premium. $p < 0.001$, significant on 94% of the horizon, $d_z = 3.67$	111
49	Permutation test, decreasing versus rising AI-cost trajectory for the skill-wage premium. $p < 0.001$, significant on 78% of the horizon, $d_z = 2.02$	112
50	Pattern 1 (profitability gate). After an exogenous profitability shock at step 240 raises wages and AI rental costs, average cost exceeds the goods-market price and proactive automation halts. Dutch institutional scenario, 12 Monte Carlo replications.	124
51	Pattern 2 (AI cost vs. adoption). Final AI adoption falls as the AI cost floor rises (15, 20, 25) and rises as the cost-decay rate increases (0.005, 0.020, 0.080). Dutch institutional scenario, 12 replications.	125
52	Pattern 3 (severance vs. adoption). Final AI adoption against the severance rate (0, 1/3, 1, 5, 20); higher severance lowers adoption. Dutch institutional scenario with a short chain limit and high conversion probability, 12 replications.	126
53	Pattern 4 (Keynesian demand). A 35% exogenous employment shock at step 240 lowers aggregate demand and the price level. Dutch institutional scenario, 12 replications.	126
54	Pattern 5 (induced technical change). Final AI adoption against the minimum wage (0, 14.71, 25, 28, 30, 32); adoption rises with the wage floor (Spearman $\rho = +0.97$). Dutch institutional scenario, 12 replications.	127
55	Pattern 6 (polarisation). Skill-wage premium and routine share of automated tasks for non-routine multipliers 0.1, 0.85 and 3.0. Dutch institutional scenario with $w_{\min} = 0$, 12 replications.	128
56	OFAT drivers of the Gini index of worker income, ranked by maximum absolute deviation from baseline (%).	133
57	OFAT drivers of the skill-wage premium (w_h/w_l), ranked by maximum absolute deviation from baseline (%).	134
58	OFAT drivers of the high-skill employment rate, ranked by maximum absolute deviation from baseline (%).	134
59	OFAT drivers of the low-skill employment rate, ranked by maximum absolute deviation from baseline (%).	135

List of Tables

1	AI-cost trajectory configurations. All three share the baseline-decay value.	25
2	Comparative advantage configurations.	26
3	Institutional configurations.	27
4	The complete experimental program across the eight experiments, with the parameter override relative to the baseline and the adoption mode. The mean-field protection-off run (MF-off) appears in both Experiment 2 (versus ULC) and Experiment 3 (versus the baseline).	29
5	Paired comparison of the ULC with the Mean-field approach. Difference = A-B, positive means a higher value for the ULC configuration. Complete matrix in Appendix B.	33
6	Paired comparison of the decision-logic step: ULC versus the mean-field protection-off run (MF-off), both without institutions. Difference = ULC – MF-off; positive means a higher value under ULC. Complete matrix in Appendix B.	33
7	Paired comparison of the institutional step: employment protection off (MF-off) versus the baseline (protection on), both mean-field. Difference = off – baseline; positive means a higher value with protection off. Complete matrix and the permutation band plots in Appendix B.	35
8	Paired comparison of the baseline (A = Decreasing) with the flat and rising trajectory. Difference = A-B, positive means a higher value for the decreasing trajectory. Complete matrix in Appendix B.	36
9	Paired comparison of the baseline (A = 0.85) with the low, medium and high comparative advantage. Difference = A-B, positive means a higher value for the baseline. Complete matrix in Appendix B.	37
10	Paired comparison of the baseline (A = 14.71) with the three alternative minimum wage levels. Difference = A-B, positive means a higher value for the baseline. Complete matrix in Appendix B.	38
11	Paired comparison of the baseline (A = 1/3) with the two alternative severance rates. Difference = A-B, positive means a higher value for the baseline. Complete matrix in Appendix B.	39
12	Paired comparison of the baseline (A = 0.5) with the two alternative transition probabilities. Difference = A-B, positive means a higher value for the baseline. Complete matrix in Appendix B.	40
13	Software dependencies and the exact versions used for the thesis results.	65
14	Repository structure: the location in the code of each part of the implementation.	66
15	Mapping of the conceptual components (C1–C7, M1) to the code.	66
16	Complete pairwise comparison for the three NPV-modes.	71
17	Complete pairwise comparison for experiment 1.	74
18	Decision-logic step of the Experiment 1 decomposition: mean-field NPV with employment protection off versus ULC, both without institutions.	77
19	Pairwise comparison for experiment 3 (employment protection on versus off, both mean-field).	80
20	All pairwise comparisons for experiment 4. Configurations: decreasing costs (baseline; initial AI cost 40, floor 20), flat (initial cost and floor both 30) and rising costs (initial cost 20, floor 40).	93
21	All pairwise comparisons for experiment 5. Non-routine multiplier by label: low = 0.50, base = 0.85, medium = 1.15, high = 1.50.	96

22	All pairwise comparisons for experiment 6. Minimum wage by label: no minimum wage = 0, low = 10, base = 14.71, high = 20.	99
23	All pairwise comparisons for experiment 7. Severance rate by label: no severance = 0, base = 1/3, high = 1.0.	102
24	All pairwise comparisons for experiment 8. Conversion probability by label: no conversion = 0, low = 0.25, base = 0.50, high = 0.75, complete = 1.0.	105
25	Baseline parameter set: value, calibration tier and basis for each parameter. The component identifiers (C1–C7) refer to chapter 3.	115
26	Summary of verification tests	119
27	Overview of the six validation patterns	122
28	Validation results: pass/fail and supporting evidence per pattern	128
29	Model parameters overview	130
30	Parameters, baseline values and tested ranges. Each numeric parameter is evaluated at five points: the baseline and $\pm 25\%$ and $\pm 50\%$ around it, clipped to the parameter's admissible domain (probabilities to $[0, 1]$; counts to integers ≥ 1). * nr_multiplier uses the scenario range of subsection 4.4.2. † alpha_min uses a fixed grid $\{0; 0.125; 0.25; 0.375; 0.50\}$. The switches employment_protection and ai_irreversible are toggled off/on. 135	
31	Assumptions that affect the whole model and set the boundaries of the system. They do not belong to a single component.	137
32	Leontief task technology and productivity: how tasks, workers and AI generate output. . .	138
33	Hiring, matching, separation and vacancy handling. The rules of the labour market. . .	139
34	A wage-supply curve per skill group with partial adjustments.	140
35	The investment decision: unit costs and NPV in four adoption modes.	141
36	The path of the AI cost over time.	141
37	The goods market and the two-class consumption function that links unemployment to demand.	142
38	Contracts, severance and the chain regulation. The Dutch institutional rules of the model. 142	
39	The connection layer between all the different components.	143
40	Technical assumptions	143

List of Abbreviations

The model is described as seven components (C1–C7) plus one macro-feedback coupling (M1); these and the other abbreviations used in the thesis are listed below. Wherever an abbreviation appears for the first time in a chapter it is a clickable link that jumps to its entry here.

Model components

C1	Production	section 3.6
C2	Labour market	section 3.7
C3	Wage-setting	section 3.8
C4	AI adoption	section 3.9
C5	AI learning curve	section 3.10
C6	Demand	section 3.11
C7	Employment protection	section 3.12
M1	Macro-feedback coupling	section 3.13

Abbreviations

ABM	agent-based model
AC	average cost
AI	artificial intelligence
API	application programming interface
CAO	collectieve arbeidsovereenkomst (Dutch collective labour agreement; see CLA)
CBS	Centraal Bureau voor de Statistiek (Statistics Netherlands)
CGE	computable general equilibrium
CI	confidence interval
CLA	collective labour agreement
CPB	Centraal Planbureau (Netherlands Bureau for Economic Policy Analysis)
DNB	De Nederlandsche Bank (Dutch central bank)
DSGE	dynamic stochastic general equilibrium
EMA	exponential moving average
EU	European Union
EUR	euro (currency unit)
GDP	gross domestic product
GPT	generative pre-trained transformer
HBO	hoger beroepsonderwijs (Dutch higher professional education)
ILO	International Labour Organization
ISCO	International Standard Classification of Occupations
LLM	large language model

Abbreviations

MBO	middelbaar beroepsonderwijs (Dutch intermediate vocational education)
MF	mean-field (NPV expectation mode; “MF-off” = mean-field with employment protection disabled)
MPC	marginal propensity to consume
NPV	net present value
OECD	Organisation for Economic Co-operation and Development
OFAT	one-factor-at-a-time (sensitivity analysis)
PIAAC	Programme for the International Assessment of Adult Competencies
SER	Sociaal-Economische Raad (Social and Economic Council)
SML	suitability for machine learning
UC	unit cost
ULC	unit labour cost
UML	unified modeling language
UWV	Uitvoeringsinstituut Werknemersverzekeringen (Dutch Employee Insurance Agency)
WAB	Wet Arbeidsmarkt in Balans (Dutch Balanced Labour Market Act)
WACC	weighted average cost of capital
WO	wetenschappelijk onderwijs (Dutch academic / university education)
WRR	Wetenschappelijke Raad voor het Regeringsbeleid (Netherlands Scientific Council for Government Policy)
WWZ	Wet Werk en Zekerheid (Dutch Work and Security Act)

1

Introduction

The Dutch labour market is institutionally different from labour markets in other members of the OECD. The Netherlands combines relatively strict employment protection for permanent workers with an unusually large flexible layer of temporary and part-time workers. According to the CBS in 2024 around forty percent of the workers had a flexible employment relationship (CBS, 2024c). On the EU-index of flexible employment the Netherlands ranks first in the EU-27, at 21.8 percent, well above the EU-27 average of 11.0 percent (CBS, 2024a). At the same time the collective labour agreements (CLA), the foundations of the Dutch polder model¹, are beginning to erode. In 2010, 75.8 percent of Dutch workers had a CLA compared to 71.8 percent in 2022 (De Beer & De Sauvage Nolting, 2024). The OECD highlights the Netherlands as exhibiting labour market dualism, with relatively strict employment protection for permanent contracts and more flexible regulation for temporary and non-standard forms of work (OECD, 2019).

In this institutional context the rise of generative AI has begun. Where earlier automation waves mainly affected routine physical and cognitive tasks, large language models (LLMs) seem to target non-routine cognitive tasks (Brynjolfsson et al., 2018; Webb, 2019). Prices of AI tools are declining rapidly. According to the AI Index Report 2025 the inference costs for a model on GPT-3.5-level between November 2022 and October 2024 decreased by a factor of 280 (Maslej et al., 2025). Epoch AI calculated that the compute-efficiency of pre-training grows by a factor of 2.4 per year (Cottier et al., 2024). However, these declining unit costs do not necessarily imply that AI adoption will become cheaper in absolute terms. The long-term pricing strategies of AI companies remain uncertain, while many leading firms currently appear to rely on high capital expenditures and have not yet demonstrated stable profitability. Moreover, gains in model capability have so far tended to require more compute, both in training and inference, so a falling cost per unit of compute need not make AI cheaper to deploy in absolute terms. The economic feasibility of adoption on task-level however depends on the relation between the AI costs and the labour costs, like Svanberg et al. (2024) showed for computer vision: only 23 percent of the wages paid for AI-exposed vision tasks would be cost-effective to automate. The relevant question therefore concerns the institutional conditions under which AI pays off, rather than AI's raw technical capability.

The WRR (2021) compared AI to the steam engine and electricity as a new general purpose technology. The SER (2025) advises to do scenario experiments to anticipate on possible distributional effects and analyse the effects of policy instruments. DNB (2024) points towards the inertia of the Dutch labour productivity growth, while the labour market is very tight. On top of that AI is currently only adopted in the most productive firms (CBS, 2025d).

Existing economic models fail to capture this setting. Equilibrium models like Acemoglu and Restrepo (2018, 2019) give a clear framework, automation moves labour through a displacement effect and gives

¹The Dutch "polder model" refers to the Netherlands' consensus-oriented system of socio-economic governance, in which employers' organisations, trade unions and the state negotiate labour-market and wage policies through consultation. Collective labour agreements are a central instrument of this model, because they translate negotiated compromises between social partners into binding sectoral or company-level employment conditions.

labour back through a reinstatement effect, but fail to capture the heterogeneity of firms, bounded-rational investment decisions, and the specific institutions that define the Dutch case. Agent-based models (ABMs) are theoretically better equipped to implement these (Fagiolo & Roventini, 2017; Tefatsion, 2006). However, state-of-the-art ABMs like Upreti & Sridhar (2024), Dosi et al. (2010, 2019, 2022) and Dawid et al. (2023; 2014), still model the adoption decision largely as exogenous and do not model labour market institutions. The combination of endogenous, bounded-rational AI-adoption with a Dutch institutional context is missing in the literature.

This thesis aims to fill this gap with an agent-based model built in Python using the Mesa library² and building upon the model of Upreti & Sridhar (2024). The model of Upreti & Sridhar (2024) is expanded upon in five different dimensions. First, firms choose explicitly between human labour and AI through an investment decision based on the net present value (NPV); second, the model of this thesis includes employment protection and the chain regulation (ketenregeling); third, wage setting is based on a wage curve with partial adjustment that resembles the rigidity of Dutch CLA negotiations indirectly; fourth, the AI price dynamics is modelled using an exogenous exponential cost-decay curve; and finally, the Keynesian feedback loop couples unemployment back to the aggregate demand. By these additions the effects of the permanent/flex share, severance costs, the chain regulation, the wage rigidity of CLA-like bargaining institutions and the minimum wage can be analysed on the Gini-index, on wage inequality and on unemployment.

The primary methodological contrast is between this forward-looking NPV decision and a myopic unit-labour-cost rule; the three NPV expectation modes (naive, adaptive and mean-field) form a secondary axis whose effect turns out to be of second order (Appendix B), so the mean-field mode is used as the benchmark throughout. The scope is deliberately bounded: the model represents employees only (self-employment, which makes up a large part of the Dutch flexible shell, is outside it), a single aggregate sector rather than the sectoral Dutch dual economy, and two fixed skill groups without schooling or mobility between them. Crucially, automation is modelled as all-or-nothing substitution: the model captures the displacement effect of AI but not the augmentation (human-AI complementarity) or reinstatement (new-task) channels through which AI can also raise labour demand, so the inequality it reports should be read as an upper bound. These boundaries, and their consequences, are set out in full in Appendix H.

The hypothesis follows from Acemoglu's (2002; 2011) induced technical change and the capital-deepening results of Vergeer and Kleinknecht (2010): rigidities that keep labour costs high and predictable steer investment towards labour-saving technologies. In the Dutch dual model, two opposing forces are expected to govern the pace of adoption. On the one hand, institutions that raise or stabilise labour costs, above all a binding minimum wage, make human labour more expensive relative to AI and are therefore expected to accelerate adoption and deepen its displacement effect. On the other hand, the forward-looking NPV decision, being more complete than a myopic unit-cost rule and combined with the cost of employment protection, is expected to temper adoption at stabilisation: firms that anticipate displacement-driven wage softening and that carry severance obligations internalise more of the cost of automating than a short-sighted rule would.

These dynamics are expected to act on the two outcomes differently. The net distributional effect, wage polarisation or compression, is expected to depend on which skill group is hit harder, which in turn follows from the relative advantage of AI on routine versus non-routine tasks: following Upreti and Sridhar (2024), automating predominantly non-routine tasks displaces high-skilled workers and is expected to compress the skill-wage premium, whereas automating predominantly routine tasks displaces low-skilled workers and widens it. Unemployment, by contrast, is expected to rise under automation regardless of which group is displaced. Here the protection institutions, severance and the chain regulation, act through two opposing channels: they raise the cost of labour, which over the long run pushes firms towards automation, but they also shield incumbent workers from immediate dismissal, which is expected to dampen the rise in unemployment in the short run. The minimum wage, by contrast, is expected to push unemployment up by accelerating the displacement of low-skilled labour, even as it compresses the wage gap. Whether these institutional brakes are in fact strong enough to redirect adoption, relative to the falling price of AI, is an empirical question that the experiments address.

²The complete source code of the model, the experiments and the analysis scripts is openly available at https://github.com/Bedjarn/LabourMarket_ABMTU; the software implementation is documented in Appendix A.

1.1. Research goal and questions

This thesis researches how endogenous AI-adoption, inside the Dutch institutional context, affects wage inequality and unemployment, and what instruments can mitigate these effects. The research question is as follows:

To what extent does endogenous AI adoption, within the context of the Dutch labour market, affect wage inequality and unemployment, and what policy instruments can effectively mitigate these effects?

The main question is split in four sub-questions:

1. What mechanisms of task-based AI-automation (specifically which regulations) in the Dutch labour market are decisive for the interaction between technology-adoption and wage formation?
2. How can the agent-based model of Upreti & Sridhar (2024) methodologically be changed and calibrated to internalise investment decisions of firms (based on NPV) and labour market frictions?
3. How does endogenizing the investment decision of AI adoption affect the simulated causal relationship between technology use and both wage inequality and unemployment, compared to exogenous scenarios? Specifically: which feedback loops come into play between wage levels, labour scarcity, and adoption speed?
4. What is the expected effectiveness of different policy instruments in mitigating wage inequality and unemployment, and how do the Dutch labour-market institutions reinforce or weaken these feedback loops?

The first sub-question is answered in chapter 2 (literature and theoretical framework), after the relevant academic literature is discussed. The second is answered across the conceptual model (chapter 3) and its implementation (Appendix A), with the explicit answer stated in section 3.14. The third and fourth sub-questions are addressed by the experiments in chapter 5 (results), interpreted in the discussion (chapter 6) and answered in the conclusions (chapter 7).

2

Literature and theoretical framework

This chapter builds the theoretical framework in four steps. First the change from a skill-based to a task-based vision on automation (section 2.1). Next the specific position of AI in that task-based approach (section 2.2). Afterwards the methodological choice for agent-based modelling and the current ABM-literature for automation and labour markets (section 2.3). Finally, the institutional characteristics of the Dutch labour market and their relevance for AI adoption (section 2.4), with which SQ1 is answered (section 2.5) and the knowledge gap is formulated (section 2.6).

2.1. From skills to tasks

The theoretical base of this research lies in the change from skill-biased technical change to the task-based framework. The work of Autor, Levy and Murnane (2003) reformulated technological change as a displacement in the demand of tasks instead of skills. They proposed that computers replace workers in routine cognitive and manual tasks, and complement workers in non-routine problem solving and communication-intensive tasks.

Acemoglu and Autor (2011) formalised this approach in a model wherein tasks are units that produce outputs and skills are capacities that allow workers to perform the task. Technology moves the relative demand between skill levels and also redistributes tasks between labour and capital. Goos et al. (2014) showed empirical evidence that job polarisation (simultaneous growth of high- and low-paying jobs with a shrinking middle) in sixteen West-European countries in 1993-2010 occurred and is explained largely by routine-biased technological change and offshoring. Balsmeier & Woerter (2019) show using Swiss firm data that digitalisation increases the employment opportunities for highly educated and decreases for low-educated people, especially for machine-based technologies.

Recent theoretical additions come from Acemoglu and Restrepo (2018, 2019). In their task-based framework capital automates tasks that previously required labour (the displacement effect, whereby the labour share of value added falls and labour demand is suppressed even as productivity rises). Simultaneously new tasks can be created wherein labour has a comparative advantage (the reinstatement effect) that raises labour demand. The result of both effects determines whether technological improvements are labour saving or labour enhancing.

2.2. AI as specific type of automation

AI is harder to position in the classical task framework than earlier automation waves. Brynjolfsson et al. (2018) introduced the suitability for machine learning (SML) measure and found that AI hits different occupations than earlier automation waves: most occupations have some SML-tasks, but only a few occupations are fully automatable. Webb (2019) analyses AI-patents and shows that AI, unlike earlier robots and software, largely exposes high-skilled tasks. This upward reach does not, however, displace the routine dimension as the primary axis of automatability. Frey and Osborne (2017) estimate that around 47% of US employment sits in a high-risk category, precisely because routine, codifiable tasks are the easiest to specify and hand to a machine, whereas non-routine perception, creative and

social tasks form persistent engineering bottlenecks. AI thus broadens exposure into higher-skilled, non-routine work without overturning this ordering: routine tasks remain the most readily automated, while non-routine tasks are exposed only partially and resist automation longest. Svanberg et al. (2024) introduce the concept of minimum viable scale, that bridges the gap between technological exposure and economical adoption: AI is only adopted when the costs of the transformed task are below the labour costs.

Data wise Gmyrek, Berg and Bescond (2023) from the ILO have made a score of generative AI-exposure on a ISCO-4-digit level. Georgieff and Hye (2022) have country specific cross-country evidence on the relationship between AI-exposure and employment development. Hernandez and Brown (2020) show a 44x reduction in training compute. The Stanford AI Index (Maslej et al., 2025) shows that the cost of inference decreases by more than two orders of magnitude in 18 months. However, the long-term pricing strategies of AI companies remain unclear. Although the cost of computing power has historically declined, many leading AI firms currently appear to operate with high capital expenditures and uncertain profitability. Moreover, improvements in model performance often remain closely linked to increases in compute requirements, both during training and inference. This means that falling unit costs of compute do not necessarily translate into cheaper AI adoption for firms. More capable models may still become more expensive in absolute terms if they require substantially more computational resources.

Two recent contributions sharpen this picture and motivate the comparative-advantage axis of this thesis. Acemoglu and Restrepo (2022) show that the relative wage declines of worker groups specialised in routine tasks in industries undergoing rapid automation account for a large share of the changes in the US wage structure since 1980. Their results provide an empirical magnitude for the task-displacement channel isolated in this model. For generative AI specifically, Eloundou et al. (2024) estimate potential occupational exposure to large language models and find that, unlike earlier waves of automation, exposure generally increases with wages and educational requirements, extending well into higher-skilled cognitive work. Their exposure measure does not, however, distinguish between labour displacement and labour augmentation. Acemoglu (2025) further cautions that the aggregate productivity gains associated with current AI advances, and the resulting wage increases, may remain modest over the coming decade. Taken together, these contributions do not establish a unique mapping between AI exposure and the routine or non-routine content of tasks. This thesis therefore treats AI's task orientation as an explicit scenario axis rather than fixing it beforehand (subsection 4.4.2). The baseline uses a routine-biased calibration that remains close to the historical pattern of automation, while alternative scenarios allow AI's comparative advantage to shift toward non-routine tasks.

2.3. Agent-based modelling and AI-adopting ABMs

The study of dynamics, heterogeneity of firms, unequal workers, bounded-rationality, institutional friction and endogenous decision-making does not fit right in representative-agent equilibrium models. Tesfatsion (2006) describes agent-based computational economics as the computational study of economic processes modelled as dynamic systems of interacting agents. Fagiolo & Roventini (2017) build upon this and argue that a feasible approach sees the economy as a complex evolving system wherein heterogeneous agents through far-from-equilibrium interactions continuously changing the structure of the system. That description fits the research question.

Within the agent-based macro models there are two types of frameworks: the Keynes+Schumpeter models (K+S) of Dosi et al. (2010) and the EURACE@Unibi-model of Dawid & Neugart (2023). Dosi et al. (2010) introduced the K+S model that combines Schumpeterian endogenous growth through R&D in the capital goods sector with Keynesian aggregate demand dynamics through consumption and investments. Dosi, Pereira, Roventini and Virgillito (2019) expand on this with a labour augmented version that models the interaction between firms and workers and lets both functional inequality and personal inequality arise from the model. Dosi et al. (2022) presents a multi-sector variant with endogenous technological paradigms. EURACE@Unibi documented in Dawid et al. (2014) is built as a policy analysis tool with heterogeneous consumers, capital good producers, regions and a financial sector. Dawid and Neugart (2023) builds on this with a task-based expansion that shows how automation redistributes income.

Upreti and Sridhar (2024) are the direct starting point for this thesis. They are chosen over the more established K+S (Dosi et al.) and EURACE (Dawid et al.) frameworks because those model technological change as an R&D lottery or in abstract task terms and cannot directly address the routine/non-routine composition central to the AI question, whereas Upreti and Sridhar embed the task-based frame explicitly and stay compact enough to extend with Dutch institutions. Their agent-based model simulates firms and workers as rational agents, where firms combine tasks to a homogenous service and stochastically adopt automation of feasible tasks. Their most important finding is that automation of routine-tasks enlarges the wage inequality while automation of non-routine tasks shrinks the wage inequality. The power of their approach is that they use the task-based frame explicitly in the ABM. The limitation is on three different points: First, automation adoption is not explicitly driven by a cost-benefit investment decision: the trade-off between investment costs and future savings is not modelled. Second, they miss institutional heterogeneity: employment protection, contract types and minimum wage are not included in the model specification. Third, AI costs remain static over time. Dosi, Pereira, Roventini and Virgillito (2019) addresses some institutional features, such as labour market regulations, but their K+S framework looks at technological change as an R&D lottery. AI-adoption is not modelled and also not treated as a forward-looking investment decision. Dawid and Neugart (2023) models tasks and automation explicitly, but their institutional context is relatively abstract, limiting its use for contexts similar to the Dutch labour market.

2.4. The Dutch institutional context

The Dutch labour market combines a relatively large share of flexible employment with strong protection for the permanent workers. Four institutional elements are decisive.

2.4.1. Permanent vs. flexible contracts

The sharp distinction between permanent and flexible contracts is one of the most distinctive characteristics of the Dutch labour market. In 2024 forty percent of the workers had a flexible employment relationship, including 2.7 million workers with a temporary or variable contract and 1.2 million self-employed (CBS, 2024c). Within the EU-27 the Netherlands is ranked first on the share of flexible employment (21.8 percent), while the EU-27 average is 11.0 percent (CBS, 2024a). In sectors such as the hospitality industry ("horeca") (64.8 percent) and culture, sport and recreation (31.8 percent) flex employees are the majority (under categories permanent, self-employed and flex) (CBS, 2025b). The distinction is sharp both statistically and legally: the contract type determines the degree of protection and the severance costs. For this research it is crucial, because a firm's adoption decision directly depends on the mix of permanent/flex workers in an affected task.

2.4.2. Employment protection and severance costs (WAB)

The "Wet arbeidsmarkt in balans" [Law labourmarket in balance] entitles all workers to a severance payment (called "transitievergoeding" in Dutch) from the first day of employment upon involuntary dismissal, calculated as 1/3 month salary per tenure year ("Wet Arbeidsmarkt in Balans", 2021). While this applies to both permanent and flex contracts in principle, effective termination costs differ substantially. Permanent workers require an active dismissal procedure via UWV³ or the cantonal court, flex contracts are typically allowed to expire. This asymmetry in effective termination costs is directly relevant to the firm's adoption decision: automating a task occupied by permanent workers carries substantially higher severance costs than the same task performed by flex workers.

2.4.3. Chain regulation (Ketenregeling) (WWZ)

The chain regulation is covered in the "Wet werk en zekerheid" [Law work and security] where the maximum time of a temporary contract is 36 months ("Wet Werk En Zekerheid", 2020). If this is exceeded the contract will become a permanent contract, unless the employer breaks the chain. For employers this is a periodical pressure: the contracts of flex workers who approach the chain limit have to be con-

³UWV stands for *Uitvoeringsinstituut Werknemersverzekeringen*. Upon dismissal of an employee with a permanent contract the employer should in some cases ask for permission of the UWV, mostly for dismissal because of business reasons or after long-term absence. UWV tests whether the employer has enough reason to dismiss and if the employer met procedural conditions. This makes dismissal of permanent workers often more formal, time consuming and more expensive than not extending a flex or temporary contract.

verted to permanent ones (where the firm accepts a future severance obligation) or the contract has to be terminated. The chain regulation is the second institutional mechanism that steers the contract mix in firms, besides the severance payment.

2.4.4. CLAs and minimum wage

The CLAs in the Netherlands are the cornerstone of the Dutch polder model, but are under fire. For tens of years the coverage was around 80 percent, in 2022 only 71.8 percent is covered by a CLA (De Beer & De Sauvage Nolting, 2024). The universally binding declaration ensures that CLA agreements also count for non-joined employers within the same sector (“Wet op het algemeen verbindend en het onverbindend verklaren van bepalingen van collectieve arbeidsovereenkomsten”, 2023). For this model the behaviour is relevant: CLA negotiations produce wage stickiness. Wages react slowly to shocks due to agreements usually taking some years to conclude, renegotiations taking time and information-asymmetry (Caloia et al., 2023). The minimum wage (in 2026 €14.71) is a hard bottom line for wages (“Wet minimumloon en minimumvakantiebijslag”, 2026).

2.5. Answering SQ1: the decisive Dutch mechanisms

The prioritisation of these institutions is, besides the literature, based on orienting conversations with labour-market experts at the Dutch Social and Economic Council (SER) and the Netherlands Bureau for Economic Policy Analysis (CPB). From these discussions it became apparent that the permanent/flex split, the severance costs, the chain regulation and the CLA agreements are, in their view, especially characteristic of the Dutch labour market and therefore particularly relevant to this question, which is precisely why they are worth being modelled and tested separately in a Dutch setting rather than assumed away. The experts regarded these institutions as decisive for the labour market at the macro level. They were not formal interviewees, the conversations were orienting and produced no explicit qualitative data.

Based on this SQ1 is answered as follows. The mechanisms of task-based AI adoption that are decisive in the Dutch context for the interaction between technology-adoption and wage setting are:

- The legal minimum wage as an absolute wage floor.
- The transition payment of 1/3 monthly salary per tenure year
- The chain regulation of 36 months that converts flex workers to permanent workers
- The indirect wage-stickiness due to CLA coverage
- The economies of scale (declining average costs) on the AI side
- The task heterogeneity between firms through a firm-specific routine/non-routine distribution

2.6. Knowledge gap

The knowledge gap has two dimensions. The first dimension is technological: in the current ABM literature AI adoption is modelled either exogenously or through R&D lotteries on a productivity scale. The bridge between technical exposure, economic feasibility, and firm-level investment decisions with NPV, expectations is not taken into account in the current frameworks.

The second dimension is institutional. The Dutch labour market is an outlier in the OECD for the permanent/flex distribution. None of these ABMs takes this into account, neither does it model severance payments or chain regulations. The model in this thesis fills both gaps by making the decision making process of AI adoption endogenously through four different NPV-modes with all different expectation calculations, and by explicitly implementing the institutions in the code.

3

Conceptual model

This chapter describes the conceptual model of the agent-based model (ABM) that models the interaction of AI-adoption, job availability and wage setting of the Dutch labour market (or similar labour markets). The model builds upon the task-based approach of Autor et al. (2003) and the productivity model of Upreti & Sridhar (2024), and expands upon their model with an adoption decision based on the net present value (NPV). This chapter gives the equations and structure, the software realisation is documented in Appendix A, and the complete source code is openly available at https://github.com/Bedjarn/LabourMarket_ABM_TU.

The central methodological contribution is to replace the myopic, stochastic adoption rule of Upreti & Sridhar (2024) with an explicit, forward-looking investment decision. The primary contrast in this thesis is therefore between a short-sighted unit-labour-cost (ULC) rule and a forward-looking net-present-value (NPV) rule. Within the NPV family, three expectation modes (naive, adaptive and mean-field) are compared as a secondary axis that probes how sensitive the outcomes are to the way firms form wage expectations; this expectation choice turns out to be a second-order effect (Appendix B), and the mean-field mode, in which firms anticipate the economy-wide wage-softening that ongoing automation causes, is used as the benchmark throughout. The core question is how the adoption decision, and the expectations behind it, shape the resulting wage inequality and unemployment.

The model is described below as seven components (C1-C7) plus one cross-cutting macro-feedback coupling (M1). This component framing is the reference vocabulary used in the following chapters. The components are: C1 Production, C2 Labour market, C3 Wage-setting, C4 AI adoption, C5 AI learning curve, C6 Demand, and C7 Employment protection.

3.1. Justification of the ABM approach

The question that is central in this thesis has three characteristics that do not fit in representative-agent equilibrium models. First, there is heterogeneity on multiple dimensions: firms differ in their task package and technological access, workers in skill type and contract, and decisions are made under bounded-rationality. Second, there is a path dependence: a firm that invests today in AI has other cost levels tomorrow and will take other expansion and adoption decisions. Third, there are institutional frictions: employment protection works differently for permanent and flex-workers, and the chain regulation links history to the contract type in a way that is hard to aggregate.

Tesfatsion (2006) and Fagiolo and Roventini (2017) offer the conceptual justification: the economic system as an ecology of heterogeneous agents that through far-from-equilibrium interaction change the structure of the system. Grimm et al. (2005) give the methodological base: pattern-oriented modelling, wherein a model is tested on the capability to create multiple empirical patterns simultaneously. Alternatives such as DSGE or CGE models require equilibrium or suppress the dynamics that are researched based on the setup.

3.2. Model architecture

The model is implemented as an agent-based model in Mesa and consists of four classes, shown in Figure 1: the environment class *LabourMarketModel* (the market) and three agent- and data classes, *Producer*, *Worker* and *Task*. The market is both the environment and the scheduler. It manages the population of the producers and workers and keeps track of the aggregated economy. *Producers* are the only entities taking decisions. Each firm has a set package of tasks and decides how much it produces, who it hires and which tasks are automated. A *Task* is the entity of both production and automation: it can be done by human *Workers* or AI-units, they register the investment costs, investment step and the NPV at the moment of investment. *Workers* are passive on purpose: they have a skill level, wage, tenure and contract type, but do not take any decisions themselves. By focussing all the decisions on a firm level the decision-making process is manageable and isolated; every change to the state of a *Worker* is made by the *LabourMarketModel*.

The associations in Figure 1 show how the entities interact with each other. The *LabourMarketModel* manages the populations of *Producers* and *Workers*, while each *Producer* has 1 or more *Tasks* (in this paper 20): *Tasks* do not exist without a firm that owns them. A *Task* can have zero or more *Workers* assigned to it, and a *Worker* is at any given moment employed by at most one *Producer*. The coloured labels link every method to one of the seven model components, from production (C1) to employment protection (C7), so the diagram can also be used as a map between the conceptual mechanisms and their implementation (Appendix A). The logic of a component is inside the class itself, while the step() method of the *LabourMarketModel* creates the macro feedback loop (M1) that connects the component to one dynamic system.




Figure 1: UML-class diagram of the model entities. Every box is a separate class with its state variables in the top box and its methods in the bottom box. Coloured labels link the methods to the seven model components (C1-C7), the class *LabourMarketModel* has the macro feedback loop (M1). The diamonds show the aggregation and composition and the multiplicities are at the ends of the links.

3.3. Agent interactions within one step

Within one tick the labour market follows a set order of steps, that Figure 2 shows from top to bottom besides the three entities. The fourteen sub-steps are of *step()* and are grouped into seven phases. In phase A the *LabourMarketModel* sets the clock forward, it updates the AI-costs from the cost curve and sets the delayed wage bill and profit income that is used for the demand in this step. In phase B every *Producer* produces, after which the *LabourMarketModel* sets the price through the inverse demand curve. Phase C applies the chain regulation and converts flex workers to permanent workers or lays them off. Phase D handles exogenous separations, every separation activates a recall for that firm that can automate the freed up task, the reactive channel of the model. Phase E updates the wages following a curve with partial adjustments and adaptive expectations. In Phase F firms place hiring requests, shielded by a profitability condition that also handles the proactive automation channel, and links the unemployed to the remaining vacancies. Phase G ends the tick with raising the tenure, the NPV update, passing on the new wage levels to the *Worker* objects, a repricing at the end for consistency and the collection of data.




Figure 2: Agent-interactions within one step. The vertical lines are the entities of Figure 1, the LabourMarketModel orchestrates the sub-step order of step(). The continuous arrows are synchronous calls or information flows, dashed arrows are returns, and self-calls are internal calculations of the LabourMarketModel. The fourteen sub-steps are grouped in seven phases (A-G) and labelled per component (C1-C7). The panel below summarises the macro feedback loop (M1) with one-tick delay.

3.4. Step order

Every time period the model goes through the following steps:

1. Update AI-learning curve and raise model step counter.
2. Calculate the lagged wage bill.
3. Let all firms produce and trim AI capacity.
4. Calculate the market price.
5. Chain regulation (if employment protection is turned on).
6. Exogenous separation of workers.
7. Update wages.
8. Update the adaptive expectations (NPV).

9. Match and hire.
10. Raise tenure of working workers (if employment protection is turned on).
11. Update the smoothed displacement expectations (used in the mean-field NPV).
12. Update the new wages to workers.
13. Recalculate production and price for data collection.
14. Data collection.

At $t = 0$ there is a warm-start where all tasks are filled with workers (preferred to the cheapest labour), wages are set on target (so no partial adjustments), and at $t = 0$ the data is collected before the first time step.

3.5. Components overview

Each of the seven components is presented below in the same structure: the purpose, the equations, the key assumptions, and the interfaces with other components. The component order is $C1 \rightarrow C7$, followed by the macro-feedback coupling M1 in section 3.13.

3.6. C1 - Production

Component C1 turns the inputs assigned to each task into firm-level output and prices each task's input choice. It bundles the task structure, the Leontief production function, the productivity functions, and the unit-cost comparison.

3.6.1. Task structure

Each producer j has N tasks, divided into two types based on the firm-specific parameter $\alpha_j \in [0, 1]$, which represents the share of routine tasks. The number of routine tasks is $N_r = [\alpha_j \cdot N]$, and the number of non-routine tasks is $N_{nr} = N - N_r$. Within each type, tasks are ordered by a complexity index x , where $x = 1$ denotes the simplest task.

The routine/non-routine distinction follows Autor, Levy and Murnane (2003), which provides the core framework for task-based automation. Heterogeneity in α_j captures differences in task composition across firms: firms with a high α_j (for example those resembling clerical or bookkeeping occupations) are more exposed to AI-driven automation than firms with a low α_j (for example those resembling hairdressers or carpenters).

The distribution of α_j can be generated in two modes. In the “uniform” mode, firms draw α_j from $U(\alpha_{\min}, \alpha_{\max})$, where $\alpha_{\min}$ and $\alpha_{\max}$ are configurable parameters. In the “data” mode, α_j is drawn from the occupational exposure scores reported by Gmyrek, Berg and Bescond (2023), aggregated to the ISCO-08 4-digit level by taking the mean across task-level scores. All experiments in this thesis use the uniform mode with $\alpha_{\min} = 0.25$ and $\alpha_{\max} = 0.75$; the data mode is an implemented but unused option, retained so that empirically grounded exposure profiles can be substituted in future work.

Gmyrek, Berg and Bescond (2023) was chosen as the exposure source for the data mode because it is the closest conceptual match to the present model. The ILO measure developed by Gmyrek et al. (2023) starts from task-level assessments of generative-AI exposure and aggregates them to detailed occupations, which can be translated naturally into a single continuous parameter $\alpha_j \in [0, 1]$. In the ABM this provides a proxy for the share of tasks within a firm that is potentially susceptible to GenAI-based automation, while remaining straightforward to implement.

Felten et al.'s (2021) AI Occupational Exposure measure is not used, although it is an important benchmark in the literature, because it is less closely aligned with the purpose of this model. Felten et al. (2021) construct an occupation-level exposure measure by mapping AI application domains onto human abilities in O*NET. This makes it an ability-based and US-centred indicator of general AI exposure, rather than a task-level measure of generative-AI automatability. For a task-based ABM in which heterogeneity enters through the share of potentially automatable tasks, the ILO data therefore offers a better match.

Georgieff and Hyee (2022) are also not used as the main calibration source. Their work is valuable for analysing the relationship between AI exposure and labour-market outcomes across countries, but it is less suitable here because it adapts the Felten framework to PIAAC data and works at the ISCO-08

2-digit level. That level of aggregation is too high for calibrating firm-level heterogeneity in an ABM that assigns exposure at a more disaggregated occupational level.

At the same time, this importance of the choice should not be overstated. In the present model, the exposure measure does not by itself generate the main results, it only shapes the initial distribution of α_j across firms. The substantive model dynamics are generated endogenously through wages, AI costs, severance costs, demand feedbacks, and boundedly rational investment decisions. As a result, the qualitative conclusions are expected to depend more on the relative ordering and dispersion of exposure across firms than on the exact score assigned to any individual occupation. In that sense, the choice of exposure metric matters mainly for the calibration of heterogeneity, not for the core causal mechanisms of the model.

In reality, the share of routine tasks is not randomly distributed, but sector- and firm-specific. A fully empirical distribution, for example weighted with detailed Dutch occupational microdata, would be more realistic, but such data are not publicly available from CBS at the required level of detail. As a consequence, firms in the current model partly function as stylised representations of occupational task profiles rather than fully diversified real-world firms. A richer extension would be to model tasks along two dimensions, routine/non-routine and cognitive/manual, but this lies beyond the scope of this thesis.

3.6.2. Leontief production function

Production follows a Leontief production function on task level:

$$Q_j = \min_x (u_x / c_x) \quad (1)$$

where Q_j is the output of producer j , u_x is the produced amount of task x and c_x the input coefficient. In this model $c_x = 1$ for all tasks (but implemented for future research simplification), so output is equal to the lowest task contribution: $Q_j = \min_x (u_x)$. This is the bottleneck principle: output is decided by the least productive task.

The Leontief production has been used in task-based automation literature before (Upreti & Sridhar, 2024). The strict complementarity enables the idea that all the tasks are necessary for the end product. A limitation of this framework is that there is no substitution between tasks. In practice some firms could partly skip tasks or combine them.

3.6.3. Task production

The contribution of task x depends on who performs the task:

$$u_x = n_{AI} \cdot \pi_{AI}(x) \text{ if automated} \quad (2a)$$

$$u_x = \sum_i \pi_h(x, s_i) \text{ if human staffed} \quad (2b)$$

where n_{AI} is the number of AI units, $\pi_{AI}(x)$ the productivity function of one AI, and $\pi_h(x, s)$ the productivity of a human worker with skill level $s \in \{high, low\}$; in equation (2b) the sum runs over the workers i staffing the task, each with skill s_i . Upon automation n_{AI} is chosen such that the AI output is at least as high as the output of the separated workers, such that the Leontief output does not collapse.

3.6.4. Human productivity

The productivity functions include a comparative advantage per skill level.

Non-routine tasks

$$\pi_h(x, high) = \exp(a \cdot x) \quad (3a)$$

$$\pi_h(x, low) = \exp(\xi \cdot a \cdot x) \quad (3b)$$

Routine tasks

$$\pi_h(x, low) = \exp(a \cdot x) \quad (3c)$$

$$\pi_h(x, high) = \exp(a \cdot x) \quad (3d)$$

The parameter $a > 0$ determines the scale of the productivity increase for complexity, and $\xi \in [0, 1]$ dampens the productivity upon a skill-task-mismatch.

Following Autor, Levy, and Murnane (2003), the assumption is that highly skilled workers have a comparative advantage in non-routine relative to routine tasks. It is captured by allowing productivity in non-routine tasks to depend on skill, while keeping routine-task productivity identical across skill groups. The parameter $a > 0$ determines how productivity rises with task complexity, and $\xi \in [0, 1]$ captures whether low-skilled workers can effectively perform non-routine tasks. This avoids imposing the stronger and less well-supported assumption that low-skilled workers have a comparative advantage in routine tasks.

3.6.5. AI productivity

AI productivity decreases with task complexity and is scaled through a non-routine multiplier:

$$\pi_{AI}(x) = \phi_0 \cdot \exp(-\phi_1 \cdot x) \cdot \mu(\text{type}) \quad (4)$$

With:

$$\mu(\text{type}) = \begin{cases} 1 & \text{if routine task} \\ \mu_{nr} & \text{if non-routine task} \end{cases}$$

The parameter μ_{nr} determines the relative suitability of AI for non-routine tasks:

- $\mu_{nr} < 1$: AI is less productive for non-routine tasks (the standard assumption).
- $\mu_{nr} = 1$: AI is as productive for both routine and non-routine tasks.
- $\mu_{nr} > 1$: AI is more productive for non-routine tasks.

The decrease of productivity with complexity (ϕ_0 = base productivity, ϕ_1 = complexity penalty and x = complexity level) resembles that AI excels at standardised tasks and struggles with more complex tasks (Brynjolfsson et al., 2018; Webb, 2019). The parameter μ_{nr} allows for calibration on both traditional automation (industrial revolution) and generative AI (LLM's that take over cognitive tasks). This distinction is a core parameter for the scenario analysis.

Both AI and human productivity are very simplified and generalised. Productivity remains the same over longer periods of time. AI units do not learn over time, and neither do humans. As a result, adoption here is driven by the falling relative price of AI rather than by improving AI capability: the model studies the affordability dynamics of automation, not its capability dynamics. Capability enters only as the comparative-advantage scenario axis (the non-routine multiplier), held fixed within each run. A freshly hired worker is as productive as a veteran worker who spent 40 years on the job as long as they have the same skill level.

3.6.6. Unit labour costs

For each task the firm compares the costs per unit task output.

$$UC_h(x, s) = w_s / \pi_h(x, s) \quad (5)$$

$$UC_{AI}(x) = k_{AI} / \pi_{AI}(x) \quad (6)$$

where UC_h and UC_{AI} are the costs per unit of task output for human and AI input, w_s is the wage of skill level s (set in C3), and k_{AI} the rental cost of one AI unit (set in C5).

The preferability is: $\min(UC_h(x, high), UC_h(x, low), UC_{AI}(x))$

3.6.7. Average cost expansion criterion

Because all 20 tasks are complementary inputs in a Leontief technology, a single bottleneck factor's unit price is in the wrong unit to compare directly with the price of one unit of one output. When using the firm's current average total cost as the expansion gate this is mitigated. Average total cost is the sum of all variable costs (labour wages plus AI rental costs) divided by current output:

$$AC_j = \frac{\sum_{w \in p_j} w_{s(w)} + \sum_{x \in X_j: t_x = \text{automated}} n_{AI,x} \cdot k_{AI}}{Q_j} \quad (7)$$

where p_j is firm j 's set of employed workers (with wage $w_{s(w)}$ set by each worker's skill $s(w)$), X_j its set of tasks, t_x the state of task x (automated or human), $n_{AI,x}$ the number of AI units on task x , and Q_j the firm's output.

The firm invests in extra capacity only when current price covers average costs: expand if $AC_j < p$, where p is the current market price (set in C6).

C1 reads wages from C3, AI rental costs from C5, and severance-induced cost increments from C7 via the NPV in C4; it writes output to C6 and triggers expansion-driven adoption requests to C4.

3.7. C2 - Labour market

Component C2 allocates workers to tasks and processes separations, hires and dismissals. It bundles the exogenous separation process and the two-round matching procedure.

3.7.1. Exogenous separation

$$P(\text{separation}|\text{working}) = \delta \quad (8)$$

where δ is the per-step job-separation probability of an employed worker. With L workers employed this generates $\delta \cdot L$ separations per period on average. This is the channel that drives the reactive adoption channel. Upon separation a firm receives information on the productivity of the separated worker; this is used to set n_{AI} if the firm decides to automate the task.

The generalisation of identical separation rates is a simplification; in reality this differs per sector. There is no endogenous separation. Firms can only proactively automate through the bottleneck channel, not by complete reorganisations. This restricts the automation speed.

3.7.2. Matching

The matching procedure has two rounds:

Round 1, preference: Requests are matched to unemployed workers of the requested skill level.

Round 2, cross-skill substitution: Unmatched requests try to match with an alternative pool if:

$$UC_h(x, fallback) \leq \theta \cdot UC_h(x, preference) \quad (9)$$

where $\theta \geq 1$ is the cross-skill cost-tolerance and *fallback* and *preference* are the alternative and preferred skill levels. At $\theta = 1$ there is no cross-skill substitution. At a large θ : almost everyone is acceptable. The order of requests is randomized, so step order does not bias matching. Workers accept any offer for a job.

The model is demand driven: workers have no agency, there is no search friction, and there is a generalised wage. This is a strong simplification.

3.7.3. Employment and unemployment

A worker is employed when assigned to a task and unemployed otherwise; there is no search or participation margin, so the labour force is the full, fixed worker population. Per skill level $s \in \{h, l\}$ the employment rate is the share of that group currently assigned to a task,

$$e_s = \frac{L_s}{N_s} \quad (10)$$

with L_s the number of employed workers of skill s and N_s the (constant) size of that skill group. The aggregate unemployment rate reported throughout this thesis is the population-weighted complement of these two rates,

$$u = 1 - \frac{N_h \cdot e_h + N_l \cdot e_l}{N_h + N_l} = 1 - \frac{L_h + L_l}{N_h + N_l} \quad (11)$$

that is, the share of all workers not currently matched to a task. At $t = 0$ the warm start assigns one worker to every task, so with 450 workers against 400 tasks in the baseline the run begins at $u \approx 11\%$. This is a starting value rather than a lower bound: a task can hold more than one worker, and firms request an additional worker whenever a task is a production bottleneck and human labour still undercuts AI on that task, so total employment can exceed the number of tasks and u can fall below its initial level. Movements in either direction are therefore endogenous, driven by automation, exogenous separations and demand contraction on the one side and by bottleneck-driven hiring on the other.

C2 supplies labour stock to C1, triggers reactive-channel adoption events in C4 on every separation, and writes employment counts to C3 and to the wage-bill aggregate used by C6.

3.8. C3 - Wage-setting

Component C3 translates labour-market tightness into wages per skill level, with the minimum wage and CLA stickiness as institutional rigidities.

3.8.1. Wage curve

Wages are set by an upward-sloping wage curve per skill level (Blanchflower & Oswald, 1995):

$$w_s^* = \max(w_{\min}, a_s + b_s \cdot L_s) \quad (12)$$

where w_s^* is the target (goal) wage for skill level s , L_s the number of employed workers of that skill, and $w_{\min}$ an explicit minimum wage. When $w_{\min} = 0$ the wage floor is the intercept a_s (reservation wage). The parameter b_s determines the wage sensitivity for employment: high L_s (tighter market) $\rightarrow$ higher wages.

3.8.2. Partial wage adjustments

$$w_{s,t} = \max(w_{\min}, \lambda \cdot w_s^* + (1 - \lambda) \cdot w_{s,t-1}) \quad (13)$$

with wage adjustment speed of $\lambda \in (0, 1)$. Half-life after a shock: $t^{1/2} = \ln(2) / (-\ln(1 - \lambda))$ periods. The minimum wage acts like a double floor: both the goal and the adjusted wage.

During the warm-start ($t = 0$) wages are immediately set on the goal wage so the model starts from a consistent initial state. From $t = 1$ onwards the partial adjustments activate.

As Keynes (Corden, 1978) already argued, money wages tend to be more rigid than real wages, in modern labour economics, sticky wages are understood as wages that adjust only slowly to shocks and changing labour-market conditions. λ is a proxy for CLA rigidity and models sticky wages, a phenomenon that wages react slowly to changes in labour market conditions, because of e.g. contract lengths, negotiation delays and information lags (Caloia et al., 2023). The minimum wage reflects the Dutch minimum wage law ("Wet minimumloon en minimumvakantiebijslag", 2026).

An important assumption is that all workers with the same skill level earn an identical wage, regardless of their employer or task. There is no downward nominal rigidity above $w_{\min}$ (wages can fall freely): the partial-adjustment speed λ slows how fast wages move toward their target in either direction, but it does not impose the asymmetric, downward stickiness that characterises Dutch CLA wage-setting, where negotiated nominal wages rarely fall. There is therefore no wage-bargaining mechanism, and the institution most central to the polder model enters only as a symmetric adjustment lag. This is a strong simplification and a substantial difference from the Dutch labour market with CLAs.

C3 reads employment counts from C2, writes wages back to worker objects (used by C1), and supplies the wage bill input to C6.

3.9. C4 - AI adoption

Component C4 determines, per task and per firm, whether to automate, given current and expected costs, severance obligations, and bounded-rational expectations. It bundles the ULC and the three NPV modes, the bounded-rationality filter, the reactive and proactive automation channels, and the irreversibility switch.

3.9.1. ULC mode (base)

In the unit labour cost (ULC) mode the adoption decision is short-sighted and immediately reversible:

$$\text{Automate task } x \text{ if } UC_{AI}(x) < \min_s UC_h(x, s) \quad (14)$$

This resembles the implementation of Upreti & Sridhar (2024) and serves as the baseline to isolate the effects of the NPV modes. In Upreti & Sridhar's model AI adoption is irreversible, but the implementation costs also do not play a role in the decision-making process.

3.9.2. NPV-modes

In the NPV-modes the automation is an investment. The NPV is:

$$NPV(x) = -(I(x) + S(x)) + \sum_{t=1}^T \left[\Delta C^t(x) / (1+r)^t \right] \quad (15)$$

where r is the per-step discount rate and T the planning horizon (number of steps); the upfront investment cost $I(x)$, severance cost $S(x)$ and per-period saving $\Delta C^t(x)$ are defined below.

With complexity dependent investment costs:

$$I(x) = I_0 \cdot (1 + \kappa \cdot x) \quad (16)$$

where I_0 is the base investment cost and κ the complexity-scaling of that cost. The severance cost $S(x)$ enters as an extra upfront cost. $S(x) = 0$ when employment protection is turned off or when there are no permanent workers on the tasks (see section 3.12 for a more extensive explanation).

And periodical savings:

$$\Delta C^t(x) = UC_h^{\text{expected}}(x, t) - UC_{AI}^{\text{expected}}(x, t) \quad (17)$$

where $UC_h^{\text{expected}}(x, t)$ and $UC_{AI}^{\text{expected}}(x, t)$ are the expected human and AI unit costs for task x at future step t ; the expectation modes below differ only in how the wage path, and hence UC_h^{expected} , is formed.

The adoption decision is:

$$\text{Automate task } x \text{ if } NPV(x) > \eta \cdot I(x) \quad (18)$$

where $\eta > 0$ is the hurdle rate. The model saves the calculated NPV on the moment the firm invests (*ai_investment_npv* in the model) for possible analysis. All firms are completely rational on the AI-learning curve. The difference is completely in the expectations of future wages.

The hurdle rate reflects that firms demand higher returns than $NPV > 0$ because of uncertainty and organisational inertia (Meier & Tarhan, 2007). The complexity scaling of κ resembles that more complex tasks demand higher implementation efforts (e.g. fine-tuning, model development, integration, testing).

Naive expectations (NPV-mode)

Firms expect that current wages remain unchanged:

$$w_s^{\text{expected}}(t) = w_s^{\text{now}} \quad \forall t \in 1, \dots, T \quad (19)$$

Naive firms underestimate the wage decreasing effects of automation and overestimate the future savings.

Adaptive expectations (NPV-mode)

Firms extrapolate the recent wage trends through an exponential moving average (EMA):

$$\Delta \overline{w}_{s,t} = \omega \cdot (w_{s,t} - w_{s,t-1}) + (1 - \omega) \cdot \Delta \overline{w}_{s,t-1} \quad (20)$$

Every time step the wages change. But because one observation is noisy, firms use a smoothed trend using EMA. $(w_{s,t} - w_{s,t-1})$ is the current wage difference between this period and the previous one, e.g. if the wage is now 34.8 and was 35.0 it is -0.2. $\Delta \overline{w}_{s,t-1}$ is the smoothed trend, the 'memory' of the firm. ω is the weight of the newest observation. E.g. $\omega = 0.3$ influences the expected wage for 30% and the memory 70%.

The expected wage is:

$$w_{\text{floor},s} = \max(w_{\min}, a_s)$$

$$w_s^{\text{expected}}(t) = \max(w_{\text{floor},s}, w_s^{\text{now}} + t \cdot \Delta \overline{w}_s) \quad (21)$$

With a_s and $s = \text{high}, \text{low}$ as reservation wage and $w_{\min}$ as minimum wage. Adaptive firms see the decreasing wage trends and discount these. This slows down automation as opposed to naive firms but is always behind reality.

Adaptive learning is more commonly used for bounded rationality in macroeconomic models (Evans & Honkapohja, 2001). The EMA-implementation is computationally simple and has only parameter ω .

Mean-field expectations (NPV-mode)

Firms project labour availability through a mean-field approach:

$$L_s^{\text{expected}}(t) = \max(0, L_s^{\text{now}} - t \cdot \overline{A} \cdot \sigma_s) \quad (22)$$

where $\overline{A}$ is the EMA-smoothed number of new workers displaced by automation per period. It follows the same smoothing as equation (20) with weight ω : $\overline{A}_t = \omega \cdot D_t + (1 - \omega) \cdot \overline{A}_{t-1}$, with D_t the number of workers displaced by automation in step t . σ_s is the share of skill level s in the current employment. The expected wages follow from the wage curve:

$$w_s^{\text{expected}}(t) = \max(w_{\min}, a_s + b_s \cdot L_s^{\text{expected}}(t)) \quad (23)$$

This approach is a cheap and clean way of approximating the recent average displacement. Firms here form expectations over *economy-wide* automation: they anticipate that ongoing automation across all

firms raises unemployment, which presses wages down through the wage curve and lowers the savings expected from automating. This is a mean-field belief about the aggregate, not a strategic calculation of a single firm's own marginal contribution to unemployment. It nonetheless creates a self-restraining feedback at the system level: the more automation firms collectively expect, the less attractive further automation looks.

It is still a bounded rationality approach, not a fully rational expectation. Strategic moves are not captured (first mover advantages). Mean-field is the baseline NPV mode of this thesis (see section 4.2).

3.9.3. Bounded rationality

Not every firm evaluates automation on every timestep:

$$\text{Evaluate if } U(0, 1) < p_{\text{eval}} \quad (24)$$

where $U(0, 1)$ is a uniform random draw and p_{eval} the per-step probability that a firm evaluates a task for automation.

This restriction of constant evaluation tries to resemble organisational inertia and limited information processing capacity (Hannan & Freeman, 1984; Simon, 1955). It also prevents the unrealistic scenario of immediate, simultaneous adoption by all firms.

3.9.4. Reactive vs. Proactive adoption

Reactive channel

Upon exogenous separation (a worker leaves the firm) a firm has to decide what to do with the now empty (or understaffed) task. Based on equation (24) it is decided if the task is reevaluated, if automation is the preferred option the task will be automated. The number of AI units is chosen such that $n_{\text{AI}} \cdot \pi_{\text{AI}}(x) \geq \pi_h(x, s_e)$ (productivity of the worker that left), so the output does not collapse.

If the task is not evaluated or automated the cheapest (UC) human replaces the old worker to prevent output collapse.

Proactive channel

Upon bottleneck expansion a firm evaluates (one evaluation per firm per period decided by equation (24)) if it is cheaper to expand capacity using AI or with labour. With proactive automation existing workers are fired and replaced by AI units that at least produce the same output.

The difference resembles two relevant decision moments: the substitution decision (reactive) and the expansion decision (proactive). Concretely the model uses two automation moments. Reactive evaluations are drawn per separation event: each exogenously departing worker triggers an independent p_{eval} draw at the firm. Proactive evaluations are drawn once per firm per step: a single p_{eval} draw determines whether all of that firm's bottleneck tasks are evaluated for proactive automation in that step. The asymmetry reflects that worker exits are local, while bottleneck-driven expansion is an organisational decision that bundles all of a firm's open positions into one deliberation.

This asymmetry is reinforced by the severance costs. Severance is paid for every permanent worker who is dismissed by an active firm decision. In the reactive channel that means: the worker whose exogenous separation triggered the evaluation does not generate severance (they have left of their own accord), but if the firm subsequently decides to automate the task and dismiss any remaining permanent co-workers, severance is paid for those co-workers. In a fully proactive automation severance is paid for every permanent worker on the task. Flex workers never generate severance, regardless of channel. Reactive automation is therefore on average cheaper than proactive automation, but not zero-severance, the gap closes as the share of multi-worker tasks rises.

3.9.5. Irreversibility

The model has a certain AI reversibility option that determines whether automation is irreversible during a planned time horizon (T).

- $ai_{irreversible} = True$: an automated task can be reevaluated after T steps. Afterwards it is again compared, if labour is cheaper the task can return to human labour. Upon renewal the clock is reset without investment costs.
- $ai_{irreversible} = False$: free de-automation on every step. This makes it possible to isolate the effect of investment irreversibility.

Now two different dimensions can be separately varied: 1. How far can a firm look forward (and how) (ULC vs. NPV) and 2. Is the investment irreversible. This tries to mimic theoretically that the irreversibility creates a real-options-effect that rewards waiting (Dixit & Pindyck, 1994).

C4 reads wages from C3, AI costs from C5, severance rules from C7, and the goods-market price from C6 (for the profitability gate). It writes adoption decisions to C1 and triggers dismissals to C2.

3.10. C5 - AI learning curve

Component C5 drives the exogenous cost trajectory of AI inputs.

The AI-costs decrease over time using an exogenous learning curve:

$$k_{AI}(t) = k_{\min} + (k_{AI_0} - k_{\min}) \cdot \exp(-\lambda_k \cdot t) \quad (25)$$

With a half-life $t^{1/2} = \ln(2) / \lambda_k$. All firms can see this curve. k_{AI_0} is the initial AI rental cost, $k_{\min}$ is the AI cost floor but can also be used as a ceiling to experiment with rising costs, and λ_k is the decay speed.

The functional form is a reduced-form exponential decay toward a floor. It is deliberately not an experience (learning) curve in the cumulative-output sense of Wright (1936), in which unit cost falls as a power of cumulative production; here the decline is driven by time rather than by accumulated adoption, and is exogenous: the pace of the curve does not depend on adoption or investment levels. In the NPV-calculation all firms anticipate the same future AI costs. An endogenous curve could be an extra feedback loop but introduces even more parameters and is outside the scope of this thesis.

Because $k_{\min}$ can act as a ceiling and λ_k can be set negative, the same functional form represents declining, flat and rising AI-cost trajectories, this is the basis for the AI-cost scenario axis in section 4.4.

C5 writes $k_{AI}(t)$ to C1 (unit labour cost comparison) and to C4 (NPV calculation).

3.11. C6 - Demand

Component C6 endogenizes aggregate demand so that automation-driven changes in wages and profits feed back into the goods market.

3.11.1. Demand curve

$$p = \max(0.01, A - B \cdot Q) \quad (26)$$

where p is the market price, $Q = \sum_j Q_j$ is the total output, A is the exogenous demand and B is the price sensitivity for the amount of goods produced. The 0.01 is built in so under certain conditions the price cannot become negative.

3.11.2. Keynesian feedback loop

$$A^t = A_0 + \gamma \cdot Y_{t-1} + \gamma_\pi \cdot \Pi_{t-1} \quad (27)$$

where A^t is the demand intercept fed into equation (26), A_0 the autonomous (baseline) demand, Y_{t-1} the total wage bill of the previous time step, and Π_{t-1} the profit of the producers from the previous time step. The delay prevents simultaneous calculations. Equation (27) is the two-class consumption function: Here γ and γ_π play the role of the propensities to consume out of wage and profit income respectively, with $\gamma > \gamma_\pi$ reflecting that workers spend a larger share of their income than firms (Lavoie & Stockhammer, 2013). They are not literal dimensionless MPCs: each maps an income level (the wage

bill or profit) onto the demand intercept, and their values are pattern-calibrated rather than estimated (Appendix C).

The feedback mechanism forms the following chain of reactions:

Automation $\rightarrow \downarrow L \rightarrow \downarrow Y/\Pi \rightarrow \downarrow A \rightarrow \downarrow p \rightarrow \downarrow \text{profit} \rightarrow \downarrow \text{incentive to expand}$

The connection to the rest of the model happens in equation (7), where firm expansion is gated using $AC_j < p$.

The Keynesian feedback loop is the idea of destroying demand by automation: if AI suppresses labour, the wage bill decreases and thus demand will decrease (Storm, 2017). This kind of mechanism is written in one of Keynes' essays (Hagemann, 2019) and is one of the mechanisms of the Keynesian view on macroeconomics. In the model wealth effects, interest rates, savings and governmental redistributions are not included.

C6 reads the lagged wage bill and profit from C3/C1, and writes the price p back to C1's profitability gate and to firm revenue.

3.12. C7 - Employment protection

Component C7 encodes the two core elements of Dutch employment-protection legislation: the severance payment (WAB) and the chain regulation (WWZ). The Dutch labour market has employment protection through the WAB and the WWZ ("Wet Arbeidsmarkt in Balans", 2021; "Wet Werk En Zekerheid", 2020). The model translates two core elements: the severance payment and the chain regulation (ketenregeling). These are activated by turning *employment_protection* on in the model.

3.12.1. Contract types and tenure years

Each worker has a contract type $c \in \{flex, permanent\}$ and tenure years τ . Each model step is one month (set using *steps_per_year* = 12). Upon the warm-start a part (*init_share_vast*) of the workers gets a permanent contract with a randomly sampled number of tenure years, drawn uniformly between the chain limit and initial tenure max years $\times$ steps per year. Flex workers draw a tenure uniformly between 0 and the chain limit minus one. Tenure is raised by one at the end of each model step for all working workers.

3.12.2. Severance costs

$S(x)$ is taken as an upfront cost in the NPV-equation (15) and is summed over the permanent workers that the firm would proactively dismiss to enable automation. In the reactive partial-vacancy case this includes the remaining co-workers. In the reactive full-vacancy case $S(x) = 0$ because the only worker on the task has already separated.

$$S(x) = \sum_{w \in \text{permanent workers on task } x} \text{severance rate} \cdot w_s \cdot \frac{\text{tenure years}_w}{\text{steps per year}} \quad (28)$$

This follows from the WAB-formula (severance rate = 1/3, so one third of a monthly salary per tenure year).

3.12.3. Chain regulation

A flex worker whose tenure years reaches the chain-limit, has to be converted to a permanent contract or is not extended:

- With chance p_{convert} the worker's contract is converted to permanent: the worker stays on his task and its contract type is changed
- Else it is not extended: the worker is laid off without automation-callback (there is no reactive evaluation), and the firm places a regular hiring-request for a new flex-worker that is handled by match and hire.

- If there are no unemployed workers available as substitution the worker is extended as a permanent worker so the task is not empty.

The choice for no automation-callback at a layoff reflects that the chain regulation is a legally enforced decision, not organisational. It is not a natural moment for an evaluation moment: there is no bottleneck that needs expansion, and no search activity stimulated by a specific problem (Cyert & March, 1963). The task runs operationally, only the contract type may change. Forcing an automation-evaluation on this moment would give the firm an extra, not motivated by reality, moment evaluation possibility.

Above all, this would create double counting: if the firm wants to automate because of expansion, this already happens through the existing channel. The proactive channel captures it when a bottleneck expansion makes AI look cheaper; the reactive channel captures it when a worker leaves of their own accord. The chain regulation does not need to open its own automation channel. The effect of adoption runs indirectly through the wage costs and severance obligations that enter the NPV-equations and the marginal-cost calculation.

The specification reflects the legal reality in the Netherlands (WAB, WWZ), but leaves out several elements: there is no UWV review, no interruption of the chain regulation through breaks between contracts, no distinction between layoffs for economic reasons and layoffs on other grounds, and no additional compensation on top of the severance payment. The severance rate is assumed to be uniform across all workers, whereas in practice the formula becomes different for long employment histories. These simplifications make the mechanism computationally manageable and help isolate the core effect: employment protection makes proactive automation more costly and slows adoption, especially in firms with many long-tenured workers.

C7 reads worker tenure and contract type from C2, writes severance increments into C4's NPV, and triggers contract conversions or layoffs that flow back to C2.

3.13. M1 - Macro-feedback coupling

The seven components are tied together by a single macro-feedback loop. Automation in C4 lowers employment in C2 and wages in C3, which lowers the aggregate wage bill, which lowers next-tick aggregate demand in C6, which lowers the goods-market price, which can close C1's profitability gate, stalling expansion and further adoption pressure in C4. The loop is closed but lagged by one tick to avoid simultaneity. Its strength depends primarily on the difference between the marginal propensities to consume out of wages and out of profits ($\gamma - \gamma_\pi$): if these are equal the loop collapses and demand behaves as if exogenous.

3.14. Answering SQ2: adapting and calibrating the model

The methodological changes introduced in the model of Upreti and Sridhar can be classified into five components, matching the five extensions stated in chapter 1. First, adoption becomes an explicit investment decision with an NPV formula, a severance component, a hurdle rate and four behaviour modes (a myopic ULC rule and three forward-looking NPV expectation modes), replacing the stochastic adoption of Upreti and Sridhar; the ULC-versus-NPV contrast is the primary contribution, while the three NPV modes differ only at second order (Appendix B). Second, employment protection and the chain regulation (ketenregeling) are added through contract types, a severance payment and a 36-month conversion limit. Third, wage setting follows a wage curve with partial adjustment that captures the rigidity of Dutch CLA bargaining. Fourth, an exogenous exponential cost-decay curve makes the AI cost dynamic, with a configurable initial value, floor and decay. Fifth, a Keynesian feedback loop links aggregate demand endogenously to wages and thus employment. The task structure itself, a firm-specific routine share α and a per-task complexity index, underlies the production component; the implementation can optionally seed α from Gmyrek et al. AI-exposure data per ISCO-4-digit occupation, although all reported experiments use the uniform draw.

In addition, there is bounded rationality via *p_evaluate* (not every firm evaluates all the time), and a two-channel adoption-trigger-mechanism (reactive upon separation, proactive upon expansion) that shows the heterogeneity in adoption moments.

Calibration happens on two levels. Where directly accessible and fitting, Dutch values are used for

initialisation, e.g. the chain limit (36 months) and the severance rate (1/3). Parameters that are not directly measurable or fitting are calibrated using pattern-matching on stylised facts (see Appendix E and Appendix C). A full list of all the parameters is given in Appendix F.

3.15. Modelling and validation

The model has been verified through unit, conservation and edge-case tests (Appendix D) and subjected to pattern-oriented checks against six behavioural patterns (Appendix E). All six pass, but they differ in evidential weight: several (such as the profitability gate, the severance term and the minimum-wage response) follow fairly directly from the model equations and act as internal-consistency checks close to verification, whereas the genuinely emergent patterns (the Keynesian demand decline and the polarisation response) carry the stronger behavioural-plausibility evidence. Taken together, these tests support the model's internal consistency and the plausibility of its behavioural dynamics, without implying a direct one-to-one empirical fit. The high-level synthesis and an honest discussion of the verification & validation limitations are given in section 6.5.

The translation of this conceptual model into executable code, the software stack, the mapping of the components onto the code, the simulation loop and the warm-start procedure, is documented in Appendix A; the complete, runnable source code, including all experiment and analysis scripts, is openly available at https://github.com/Bedjarn/LabourMarket_ABM_TU. Next, chapter 4 builds directly on this model and specifies the experiments that put it to work.

4

Experimental design

This chapter sets what is experimented, against what, and how the results are measured. In chapter 5 the results are presented, this chapter specifies the configurations, how the outcomes are read, and the experimental design as a whole.

The design is scenario-exploratory rather than predictive: the goal is to explore how the model outcomes vary on different sets of plausible futures and institutional choices, not to predict a point estimate for the Dutch labour market. After the demarcation (section 4.1), the baseline against which every scenario is compared is defined (section 4.2). section 4.3 examines the difference between the NPV approach and ULC approach of Upreti & Sridhar (2024). section 4.4 sets the three scenario axes, the AI-cost curve, the comparative advantage and the Dutch institutional instruments, each with the parameter values. section 4.5 to section 4.7 set the outcome variables, the replication and statistical tests, and the time horizon. section 4.9 summarises the complete experimental design.

4.1. Scenario-exploratory demarcation

The calibration does not pretend to represent the real future of the Dutch labour market. The baseline (section 4.2) is a transparent reference point, defined in such a way that each reported scenario can be read as a deviation from a single, well-documented configuration. The three axes of section 4.4 move the model away from that baseline along dimensions for which the empirical literature offers no conclusive evidence: the trajectory of AI rental costs over the simulated horizon, the comparative advantage assumption inherent in the non-routine multiplier, and the calibration of the Dutch institutional regime. By treating each of these as an explicit axis rather than an assumption, the analysis maps out a space of plausible outcomes instead of fixating on a single point. This approach follows the exploratory scenario-based advice of the SER (2025) on generative AI in the Dutch labour market.

4.2. Baseline configuration

All scenario configurations are evaluated against one baseline. The baseline uses the parameters of Appendix F (with the calibration of the parameters in Appendix C) and the mean-field NPV mode of component C4 (section 3.9). The choice for the NPV-mode is theoretical: of the four adoption rules, the mean-field rule is the only one in which firms internalise the macro-feedback of their own adoption: as automation raises unemployment, the wage curve lowers wages and therefore the savings expected from further automation. Naïve expectations ignore the wage trend and adaptive expectations are always lagging. Mean-field is thus the most defensible boundedly-rational benchmark, and every scenario in this thesis is evaluated against it.

The decision is empirically tested (Appendix B). A pairwise permutation test between the three NPV modes, the same test as in section 4.6, shows that there are significant differences between the modes, but these differences are small. The order aligns with the theory: mean-field has a systematically higher Gini-index and unemployment rate, followed by adaptive and then naïve, exactly the order as expected.

The naïve and adaptive modes therefore serve only to justify the choice of mean-field as the behavioural benchmark; they are not run across the scenario axes of section 4.4. Every reported scenario uses the mean-field mode, and the complete pairwise comparison of the three modes is given in Appendix B.

4.3. Comparison with Upreti & Sridhar

To quantify what the NPV and institutional expansion add to the framework of Upreti & Sridhar (2024), the baseline is compared to a Upreti & Sridhar parameter configuration. ULC is shortsighted, has no hurdle rate, no employment protection and approaches the adoption logic of Upreti & Sridhar (2024) from this model. The Upreti & Sridhar configuration is not meant as a literal replication of their model but as the nearest counterpart in this model, the difference between both runs isolates the effect of the investment logic and other additions.

The comparison reports the difference in the three primary outcome variables of section 4.5, the Gini-coefficient, the skill-wage premium and the unemployment rate, not as a single end value but during the whole run, with a paired test from section 4.6. The difference itself is a result: It shows how much of the model behaviour that shows up in the subsequent scenarios is attributable to the NPV-investment logics or the other components. Both configurations are run with the same seeds.

Because the baseline carries both the forward-looking investment logic and the Dutch institutions while the ULC configuration carries neither, this single contrast bundles the two. To separate them, a third configuration is added: the mean-field NPV with employment protection switched off (MF-off), which keeps the forward-looking logic but drops the institutions. The baseline-to-ULC gap is then decomposed into a decision-logic step (MF-off versus ULC), reported as Experiment 2, and an institutional step (baseline versus MF-off), reported as Experiment 3 (section 5.3).

4.4. Scenario axes

Three scenario axes around the baseline are explored. The AI cost trajectory discussed in subsection 4.4.1 handles the future path of the AI costs as an assumption instead of strict. The comparative advantage flip in subsection 4.4.2 researches the effects of whether AI has a comparative advantage on routine or non-routine tasks. subsection 4.4.3 varies the institutions of the Dutch labour market around their legal value.

4.4.1. AI Cost trajectory

The path of the AI costs is determined by the three parameters of the learning curve in component C5 (section 3.10): the initial costs, the floor, and the decay speed. With the decay speed on its baseline-value the three trajectories below are set by varying the initial costs and the floor, so that it shows both a rising and a descending curve.

In the descending curve, the baseline path that is used elsewhere in the thesis, starts from an initial value and over time decreases to the floor following the exogenous exponential cost-decay curve of component C5. In the flat trajectory the floor is set to the initial costs, thus the costs remain constant over time, this isolates the role of the price dynamics. In the rising trajectory the floor and initial costs are swapped: the cost starts at a lower value and rises following a similar curve to the higher value.

The rising variant is included because the assumption of decreasing costs is not a set fact. Two supply mechanisms can plausibly flip the direction. First, more scarce computational power and energy (due to better models demanding more computational power). Second, suppliers of large language models can move their pricing strategy towards higher prices per call if the context windows have become larger and users are locked-in, even if the raw costs of inference keep decreasing.

The table below reports the parameter values used in the three trajectories. All other parameters take their baseline-values as shown in the parameter calibration appendix (Table 1).

Table 1: AI-cost trajectory configurations. All three share the baseline-decay value.

Trajectory	Start k_{ai}	Floor $k_{ai, floor}$
Decreasing (baseline)	40	20

Trajectory	Start k_ai	Floor k_ai, floor
Flat	30	30
Rising	20	40

4.4.2. Comparative advantage flip

The non-routine multiplier in component C1 (section 3.6) scales AI productivity on non-routine tasks relative to routine tasks: routine tasks are the fixed reference at $\phi(x)$, non-routine tasks are set at $\mu_{nr} \cdot \phi(x)$. A value below one therefore does not make AI better at routine work, it makes AI worse at non-routine work.

The empirical literature is divided on which side the advantage falls. The standard view from earlier automation waves, in Brynjolfsson et al. (2018) and Webb (2019), places the multiplier below one: AI complements non-routine tasks rather than substituting for them. Recent evidence on large language models points the other way, with substantial productivity gains in cognitive and communicative tasks that earlier automation could not address.

Four configurations bracket both sides. At 0.50 AI is markedly weaker on non-routine work, a sharpened version of the pre-generative-AI assumption. The baseline of 0.85, a mild non-routine penalty, is a deliberately conservative routine-biased anchor consistent with section 2.2. At 1.15 AI has a small non-routine advantage and at 1.50 a clear one, matching the claim that generative AI substitutes for cognitive tasks.

Because the multiplier scales one side while holding the other fixed, it moves the average level of AI productivity as well as its distribution across task types: with a routine share drawn between 0.25 and 0.75, economy-wide average AI productivity runs from roughly 0.75ϕ at $\mu_{nr} = 0.50$ to 1.25ϕ at $\mu_{nr} = 1.50$. This axis therefore varies the reach of automation jointly with its direction, and part of the difference between configurations reflects how much work is automatable at all rather than which group performs it.

Because pattern 6 of the validation (Appendix E) showed that the skill-wage premium is not monotonic in this parameter, the low-skilled employment rate in chapter 5 is reported alongside the wage premium, so that the effects remain interpretable (Table 2).

Table 2: Comparative advantage configurations.

Scenarios	Value
Large advantage for routine tasks	0.50
Small advantage for routine tasks (baseline)	0.85
Small advantage for non-routine tasks	1.15
Large advantage for non-routine tasks	1.50

4.4.3. Institutional levers

Four Dutch institutional settings are varied separately around the baseline: the minimum wage (component C3, section 3.8), employment protection as a whole (the on/off switch of component C7, section 3.12), the severance rate (component C7) and the transition probability (p_convert, component C7). When one is varied the others are kept on their baseline value so the effect of each is isolated on the outcome variables.

The minimum hourly wage is varied across four points: no floor, a value lower than the current minimum wage, the statutory Dutch level of EUR 14.71 applicable from January 2026, and one significantly higher level examining the channel of induced technical change of Acemoglu (2002) and Vergeer and Kleinknecht (2010). The severance pay rate is varied across three points covering no protection, the statutory one-third-month-per-year of service under the Balanced Labour Market Act (WAB, 2021), and a higher level (1.0 month per year of service) at which the severance pay term becomes more prominent in the firm's investment decision. Finally the transition probability is changed indicating that more/less flex contract are transitioned to a permanent contract. This isolates the effect of the

share of permanent contracts versus higher severance payments per worker. In addition, employment protection is switched off entirely against the on-baseline. This toggles the whole of component C7 at once, the severance payment, the chain regulation and the build-up of permanent contracts, isolating the aggregate effect of the Dutch protection regime rather than one of its internal margins; it is the same institutional step that decomposes Experiment 1 (section 4.3) and is reported as Experiment 3.

The table below summarises the levers. Bold values are the baseline values (Table 3).

Table 3: Institutional configurations.

Lever	Values
Minimum wage $w_{\min}$ (€)	0 10 14.71 20
Employment protection (on/off)	off on
Severance payment severance_rate (month/tenure year)	0.0 0.33 1.0
Transition probability p_{convert}	0.0 0.25 0.5 0.75 1.0

4.5. Outcome variables

Three primary outcomes are reported for all configurations, and they are best read as a complementary triplet rather than as three independent dimensions. The **unemployment rate** is the cleanest measure of displacement: it captures directly how many workers are substituted by AI. The **skill-wage premium** (w_h/w_l), the ratio of the high-skilled wage to the low-skilled wage at the same time step, is the cleanest measure of the wage structure: it captures the polarisation or compression of wages among the employed while ignoring the unemployed. The **Gini-index** of worker income sits between the two. It is computed over all workers at the end of the simulation with the unemployed counted at zero income, so it combines two channels: wage dispersion among the employed and exclusion through unemployment. The weight of these two channels is not fixed but shifts with the employment level. When employment is high the block of zero incomes is small and the Gini mostly reflects wage dispersion among the employed; as unemployment rises, the growing mass of zero incomes increasingly dominates the index. In the experiments the unemployment channel turns out to be the larger mover, so the Gini co-moves strongly with the unemployment rate, but it still carries a wage-dispersion component that the unemployment rate alone does not. The Gini should therefore be read as a composite employment-and-wage measure, interpreted together with the unemployment rate (the displacement channel) and the skill-wage premium (the wage-structure channel).

These three outcomes map directly onto the two dimensions the main research question asks about, wage inequality and unemployment, and therefore onto sub-questions 3 and 4. The unemployment rate answers the unemployment dimension directly. Wage inequality is captured at two levels: the skill-wage premium measures inequality in the wage *structure* between the skill groups, the polarisation-or-compression question that follows from the task-based literature, while the Gini index measures total income inequality across all workers, including the exclusion of those displaced into unemployment. Reporting both is deliberate. Sub-question 3, on how endogenous adoption shifts inequality and unemployment relative to exogenous scenarios, requires separating the wage-structure channel (the premium) from the displacement channel (the unemployment rate) that the composite Gini blends together; the Upreti & Sridhar (2024) comparison and the feedback loops behind it (chapter 5) are read off exactly this separation. Sub-question 4, on the effectiveness of policy instruments, is then judged by how each institutional lever moves the same triplet: an instrument that compresses the premium but raises unemployment, and with it the Gini, is assessed differently from one that lowers all three. The triplet therefore lets every experiment in chapter 5 be evaluated against the exact terms of the research question, rather than against a single aggregate that would hide which channel is moving.

Three extra indicators are reported if they help to explain the existing scenario-outcomes: the low-skilled employment rate, the high-skilled employment rate and the goods price. The routine/non-routine share is mainly informative for subsection 4.4.2, where the wage-premium signal of polarisation is not monotonic in the comparative advantage parameter. The price is informative for subsection 4.4.1 and for the Upreti & Sridhar comparison of section 4.3, where the feedback loop can directly mitigate or reinforce the cost effects on adoption.

4.6. Replication and statistical approaches

Each configuration is run with 50 different seeds, counting up from 42. The same seed set is used for the different configurations: the seed determines all the stochasticity in the simulation, and thus one seed gives the same initial pattern, the same exogenous separations and the same matching order. These are common random numbers: it makes the comparison between two configurations paired and removes the partly stochastic component, such that the effect to be analysed is isolated.

Why test on the difference of the trajectory and not the end point. The model is stochastic: each seed produces another path. Comparing two configurations on one value, the last step of the one against the other, throws away the whole trajectory and thus the model behaviour that got it there. Moreover, the goal is to analyse whether two configurations show different dynamics, not whether they differ at a single moment in time. This weighs heavily because the three NPV modes are hardly distinguishable. The path to that point captures the interesting information not the point itself. An end-value comparison with a difference of e.g. 1 standard deviation is weak, it only uses one of the steps, and the hypothesis could be rejected or accepted based on random oscillations of the model on that specific moment in time (e.g. the baseline having a small peak and the scenario a small crash). What is necessary is a test that assesses the whole trajectory and corrects for the fact that the trajectory is inspected on many different timesteps.

The test. For two configurations A and B and a given outcome variable, the seed-by-seed difference forms the matrix $D_s(t) = A_s(t) - B_s(t)$, where s runs over the 50 seeds and t over the steps. At each step, the studentized mean-difference statistic $t(t) = \frac{\text{mean}_s D_s(t)}{\text{standard error}_s D_s(t)}$ is computed. The global test statistic is the supremum ($\sup_t |t(t)|$) over the horizon. The null hypothesis, both configurations produce identical trajectories, makes the A/B label exchangeable within each seed, so flipping the sign of an entire seed-difference vector is a valid permutation. By building the null distribution of ($\sup_t |t(t)|$) overall (2^{50}) sign patterns, in practice, a Monte Carlo sample of 5,000 permutations, a global p-value is obtained that simultaneously corrects for multiple testing over time, using a sup-(t) or max-statistic procedure. The same resampling provides the critical value (c), the $((1 - \alpha))$ -quantile of the null distribution, and thus a simultaneous confidence band (mean difference (t) $\pm c \cdot$ standard error (t)). The steps at which this band excludes zero are exactly the steps where the two curves differ significantly, with family-wise error-rate correction. In this way, the test indicates not only whether the trajectories differ, but also where along the horizon.

What is reported. Two summaries of this test carry the interpretation in chapter 5. The *significant horizon* is the fraction of the 960 steps at which the simultaneous band excludes zero; it measures how persistent a difference is. The effect size d_z is Cohen's paired d_z : per seed the difference is first averaged over the horizon, $\Delta_s = \text{mean}_t D_s(t)$, and then $d_z = \text{mean}_s \Delta_s / \text{sd}_s \Delta_s$, the mean per-seed difference divided by its spread across seeds. Its sign gives the direction of the effect and its magnitude how large and how consistently the two configurations differ; because the common random numbers strongly shrink the across-seed spread, d_z can become large. As a rough guideline $|d_z| < 0.3$ is negligible, 0.3–0.8 small, 0.8–2 medium, and > 2 large. With 50 seeds nearly every p-value lies below 0.001, so the p-value is of little value, the reporting in chapter 5 therefore leans on d_z and the significant horizon.

4.7. Time horizon

Each simulation runs for 960 timesteps, equal to 80 simulated years for 12 steps per year. This horizon is long enough for the AI-curve (subsection 4.4.1) to reach the floor and for the model to stabilise. It is not chosen for any similarity with a specific real period: the scenario-exploratory goal of section 4.1 makes the absolute horizon less important than the trajectories of the configurations.

4.8. Verification and validation

The model has endured a verification test that tests the mechanical correctness of each component, and pattern-oriented checks based on six stylised facts of labour market behaviour. All verification tests and all six patterns pass, though the patterns differ in evidential weight: several act as internal-consistency checks close to verification while the demand-decline and polarisation patterns are genuinely emergent (Appendix E). The complete description of the verification and validation is in Appendix D and

Appendix E. An honest discussion of the limitation of the validation tests is given in section 6.5. Besides the verification and validation a one-factor-at-a-time sensitivity analysis (Appendix G) is done for all model parameters which show which assumptions (Appendix H) influence the outcomes the most.

4.9. Experimental program

The design choices of section 4.2 to section 4.4 specify together a program of eight experiments, summarised in Table 4: the comparison with Upreti and Sridhar (Experiment 1) and its two-step decomposition into the decision logic (Experiment 2) and the institutions (Experiment 3), the cost trajectory (Experiment 4), the comparative advantage (Experiment 5), and three further institutional levers (Experiments 6 to 8).

For each run the complete trajectory of the three primary outcomes is saved, together with the secondary indicators that help explain some configurations. The trajectory tests between configurations are calculated following section 4.6. The results are presented and interpreted in chapter 5, the broader implications for the Dutch labour market follow in chapter 6.

Table 4: The complete experimental program across the eight experiments, with the parameter override relative to the baseline and the adoption mode. The mean-field protection-off run (MF-off) appears in both Experiment 2 (versus ULC) and Experiment 3 (versus the baseline).

Exp.	Configuration	Override as opposed to the baseline	Mode
E1	baseline meanfield	(baseline)	mean-field NPV
E1	upreti sridhar ulc	Same parameters, ULC logic	ULC
E2	meanfield protection off (MF-off)	employment_protection=False; versus ULC	mean-field NPV
E3	employment protection on (baseline)	(baseline)	mean-field NPV
E3	employment protection off (MF-off)	employment_protection=False	mean-field NPV
E4	ai cost decline	(baseline) k_ai=40, floor=20	mean-field NPV
E4	ai cost flat	k_ai=30, k_ai_floor=30	mean-field NPV
E4	ai cost rising	k_ai=20, k_ai_floor=40	mean-field NPV
E5	nr multiplier low	nr_multiplier=0.50	mean-field NPV
E5	nr multiplier base	(baseline) 0.85	mean-field NPV
E5	nr multiplier medium	nr_multiplier=1.15	mean-field NPV
E5	nr multiplier high	nr_multiplier=1.50	mean-field NPV
E6	w min 0	w_min=0	mean-field NPV
E6	w min 10	w_min=10	mean-field NPV
E6	w min 14.71	(baseline) 14.71	mean-field NPV
E6	w min 20	w_min=20	mean-field NPV
E7	severance 0.0	severance_rate=0.0	mean-field NPV

Exp.	Configuration	Override as opposed to the base-line	Mode
E7	severance 0.33	(baseline) 1/3	mean-field NPV
E7	severance 1.0	severance_rate=1.0	mean-field NPV
E8	Transition probability p_convert 0.0	p_convert=0.0	mean-field NPV
E8	Transition probability p_convert 0.25	p_convert=0.25	mean-field NPV
E8	Transition probability p_convert 0.5	p_convert=0.5	mean-field NPV
E8	Transition probability p_convert 0.75	p_convert=0.75	mean-field NPV
E8	Transition probability p_convert 1.0	p_convert=1.0	mean-field NPV

5

Results

This chapter presents the results of the experimental design of chapter 4. The aim is descriptive and exploratory: per experiment this chapter shows how the model outcomes develop over time and how every scenario differs from the baseline. The translation of these patterns to the Dutch labour market and broader policy implications follow in chapter 6 and chapter 7, this chapter remains a description of what the model generates and which internal mechanism drives what without making any conclusions on the real world.

The eight experiments follow the program of section 4.9. Experiment 1 compares the calibrated baseline (mean-field NPV) with the myopic ULC rule of Upreti & Sridhar (2024); Experiments 2 and 3 decompose that comparison into the contribution of the forward-looking decision logic (Experiment 2) and of the Dutch employment-protection institutions (Experiment 3). Experiment 4 varies the trajectory of AI costs and Experiment 5 the comparative advantage of AI across tasks. Experiments 6 to 8 vary three institutional levers in turn: the minimum wage (Experiment 6), the severance rate (Experiment 7), and the flex-to-permanent transition probability (Experiment 8).

Each configuration is run with 50 seeds, starting from 42, with the same seed set for every configuration. The comparison between two configurations is therefore paired: per seed the pair shares the same starting position, exogenous separations and matching order, so that only the effect of the scenario remains. The figures show, per outcome variable, the average over the 50 seeds with a ± 1 standard-deviation band over the complete horizon of 960 steps (80 years at 12 steps per year). Two configurations are compared with the paired permutation test of section 4.6.

The results are read mainly through two summaries of that test, both defined in section 4.6. The effect size d_z is the standardised paired difference: its sign gives the direction of the effect and its magnitude indicates how large and how consistent the difference is across seeds (as a guide, $|d_z| > 2$ counts as large). The significant horizon is the share of the 960 steps at which the simultaneous confidence band excludes zero, that is, how persistent the difference is. Because with 50 seeds almost every p-value lies below 0.001, the p-value itself carries little information here.

For each experiment, the main text shows the trajectory figures and a compact table comparing every scenario variant against the baseline on the three primary outcomes: the Gini index, the skill-wage premium (w_h/w_l), and the unemployment rate. The full pairwise comparisons (all pairs, with p-values and confidence intervals), an example of a pairwise permutation test, and larger per-experiment figures are in Appendix B. Where this chapter reports a set of scenarios around the baseline, Appendix G adds a one-factor-at-a-time sensitivity analysis on all parameters.

5.1. Experiment 1 - Baseline (mean-field NPV) versus Upreti & Sridhar (ULC)

The first experiment isolates what the NPV-investment logic and the institutional expansions add as opposed to the short-sighted ULC-decision making which mimics the approach of Upreti & Sridhar

(2024). These configurations run with the same seeds and parameters.

For all three outcomes (Figure 3) the ULC-configuration is structurally higher than the mean-field baseline, and the difference grows over the horizon instead of disappearing. For the income inequality and the unemployment rate the gap is large and spans the whole horizon ($d_z \approx 9.9$ on both the Gini and unemployment, significant on 100% of the horizon, Table 5). The skill-wage premium separates the two configurations far less ($d_z \approx 3.2$, 97%): both runs compress the wage inequality in the first phase, after which ULC becomes somewhat higher. The difference between the two configurations is mainly in the level of automation and unemployment, which leads into a different wage structure of the workers.

The difference is explainable due to the adoption logic component C4. The ULC-decision making only compares the actual costs per unit in component C1 and is shortsighted. Once the AI-curve (C5) decreases below the wages (C3) there is automation. The mean-field NPV weighs the future savings, with a hurdle rate and, in the baseline, the severance costs in C7 as an extra hurdle. The key point is that the mean-field expectation accounts for economy-wide wage suppression: more automation leads to unemployment (C2), which through the wage curve (C3) presses wages down and therefore lowers the expected savings. On top of that the macro-feedback (M1) works: lower wages lower the wage bill and thus the demand (C6), which presses the price and disincentives the expansion gate (C1). This shows in the goods market price (Figure 4), which falls further and faster under ULC (to ≈ 190 against ≈ 340 for the mean-field), reflecting its higher automation and suppressed wage bill. ULC misses both slowing mechanisms and thus automates to a much higher level. The gap between the configurations is to an extent the addition of the NPV-logic, the institutions and the Keynesian feedback to the base dynamics of the model.




Figure 3: Gini-index, Skill-wage premium and Unemployment rate over time of the ULC configuration and Mean-field NPV. Average over 50 seeds with a ± 1 standard deviation band.




Figure 4: The market price over time of the ULC configuration and Mean-field NPV. Average over 50 seeds with a ± 1 standard deviation band.

Table 5: Paired comparison of the ULC with the Mean-field approach. Difference = A-B, positive means a higher value for the ULC configuration. Complete matrix in Appendix B.


Outcome	Difference (ULC - baseline)	95%-CI	d_z	Sign. horizon
Gini	+0.174	[0.169, 0.179]	9.86	100%
Skill-wage premium	+0.071	[0.064, 0.077]	3.15	97%
Unemployment	+0.181	[0.176, 0.186]	9.79	100%

This single contrast bundles two changes at once: the move from a myopic unit-cost rule to a forward-looking investment rule, and the addition of the Dutch employment-protection institutions. To separate the two, the next two experiments dissect the additions separately: Experiment 2 isolates the decision logic (MF-off versus ULC) and Experiment 3 the institutions (baseline versus MF-off).

5.2. Experiment 2 - Decision logic: forward-looking NPV versus myopic ULC

The second experiment isolates the decision logic by comparing the two institution-free configurations: the forward-looking mean-field NPV with employment protection off (*MF-off*) against the myopic ULC rule. Both carry no severance, no chain regulation and no permanent contracts, so the only thing that differs is how firms decide to automate. This is the first step of the decomposition of Experiment 1. For the Gini and unemployment the ULC rule is structurally higher than MF-off, by a large margin that holds over the entire horizon ($d_z \approx 6.9$, significant on 100% of the horizon); the skill-wage premium differs less ($d_z \approx 3.0$, 99%) (Figure 5, Table 6).

The mechanism is component C4. The ULC rule compares only current unit costs (C1) and automates the moment the falling AI curve (C5) drops below the wage (C3); it is short-sighted and misses two brakes that the mean-field NPV builds in. First, the NPV weighs future savings against a hurdle rate, so a task is automated only when the discounted savings clear a margin, and the investment is locked in over the horizon (irreversibility). Second, the mean-field expectation anticipates economy-wide wage suppression: more automation raises unemployment (C2), which through the wage curve (C3) presses wages down and lowers the expected savings, while the Keynesian feedback (M1) reinforces this through demand and the price. ULC sees none of this and overshoots to a much higher automation level. The “decision logic” measured here is therefore the combined effect of the forward-looking NPV, its hurdle rate and its irreversibility, taken as one package against the myopic rule.

**Figure 5:** Gini-index, skill-wage premium and unemployment rate over time of the ULC rule and the mean-field NPV with employment protection off (MF-off), both without institutions. Average over 50 seeds with a ± 1 standard-deviation band.**Table 6:** Paired comparison of the decision-logic step: ULC versus the mean-field protection-off run (MF-off), both without institutions. Difference = ULC – MF-off; positive means a higher value under ULC. Complete matrix in Appendix B.

Outcome	Difference (ULC - MF-off)	95%-CI	d_z	Sign. horizon
Gini	+0.161	[0.154, 0.167]	6.92	100%
Skill-wage premium	+0.065	[0.059, 0.071]	3.02	99%
Unemployment	+0.168	[0.161, 0.174]	6.81	100%

5.3. Experiment 3 - Institutions: employment protection on/off

The third experiment isolates the second step of the dissection, the employment protection institutions, by turning them on and off with the forward-looking logic held fixed: the calibrated baseline (protection on) against the mean-field protection-off run (MF-off) of Experiment 2. Switching *employment_protection* off removes the severance payment (the severance term drops out of every NPV), the chain regulation (no flex-to-permanent conversions) and the build-up of permanent contracts (all workers stay flex with no tenure tracked) in one move, while leaving everything else untouched: the minimum wage, the wage curve, the AI cost path, demand and the matching process are identical in the two runs. The comparison therefore isolates the aggregate effect of component C7 (section 3.12) and nothing else.

Switching protection on lowers all three outcomes relative to the protection-off run, but by a small-to-medium and, crucially, time-varying amount (Figure 6, Table 7). The effect size is 0.71 (d_z) on the Gini and unemployment and 0.90 on the skill-wage premium, the largest of the three because the premium varies little across seeds.

The temporal pattern is the distinctive feature. The paired permutation test marks the trajectories as significantly different over 34–35% of the horizon for the Gini and unemployment and 29% for the premium, and that the significant parts are concentrated at both ends: the simultaneous band excludes zero early and again late, with an insignificant middle (the band plots are in Appendix B). The model logic explains this shape, but not through a brake that re-strengthens over time. Severance enters the adoption decision as a one-off upfront cost that falls only on permanent workers and scales with their tenure (section 3.12); it therefore raises the hurdle to proactively automate a task that is currently staffed by permanent workers. This brake is strongest at the start, for two reasons that coincide at the warm-start: the inherited permanent workers carry high tenure, so the severance owed on their tasks is large (around €40, Figure 27), and the AI rental cost is still falling, so many adoption decisions are still marginal and a severance increment is enough to tip them. The protected economy therefore automates and displaces slightly less in the first phase, and starts with lower inequality and unemployment.

Crucially, this brake does not re-strengthen later; it fades. Once the AI cost reaches its floor, early and identically in both runs (section 3.10), automating becomes so cheap that the severance increment no longer flips any decision, and the per-step automation gap stabilises at a small, roughly constant margin. The severance owed per task does not accumulate but declines and then flattens (Figure 27), and the average tenure of permanent workers stabilises later rather than growing without bound (Figure 28), because exogenous separations cap it; the permanent-contract stock and the flex-to-permanent conversion flow are likewise flat across the second half. The late re-opening of the gap is therefore not a contract-dynamics effect.

What re-opens it is the Keynesian demand feedback (component M1, section 3.13) amplifying the small early difference. The extra automation without protection lowers the wage bill, which feeds the demand intercept (subsection 3.11.2); lower demand lowers the price, and when average cost approaches the price the expansion gate closes and firms stop hiring proactively. The diagnostics show exactly this loop: the number of expansion-gate blocks rises steeply without protection while staying flat with it (Figure 29), and the demand level drifts down faster in the protection-off run (Figure 30), dragging price, production and profits with it. Because the closed gate suppresses precisely the re-hiring of the cheapest, most automatable workers, the divergence concentrates in low-skill employment (around 73% versus 79% by the end, Figure 31); the unemployment gap widens to five or six percentage points, more than the stable automation gap of about two points of the workforce can account for on its own, the remainder being re-absorption foregone through the demand channel. The AI cost is not the driver: it is on its floor and flat throughout this phase (section 3.10). The institutional wedge is thus a small but genuine early brake that the model's own demand feedback amplifies into a persistent level difference, an interpretation developed in subsection 6.2.5.




Figure 6: Gini-index, skill-wage premium and unemployment rate over time with employment protection on (baseline) and off, both in the mean-field NPV mode. Average over 50 seeds with a ± 1 standard-deviation band.

Table 7: Paired comparison of the institutional step: employment protection off (MF-off) versus the baseline (protection on), both mean-field. Difference = off – baseline; positive means a higher value with protection off. Complete matrix and the permutation band plots in Appendix B.

Outcome	Difference (off - baseline)	95%-CI	d _z	Sign. horizon
Gini	+0.017	[0.011, 0.024]	0.71	35%
Skill-wage premium	+0.010	[0.007, 0.012]	0.90	29%
Unemployment	+0.017	[0.011, 0.024]	0.69	34%

5.4. Experiment 4 - AI costs trajectory

Having set up and decomposed the headline gap, the analysis returns to the technology axis. The fourth experiment holds the mean-field logic fixed and varies the trajectory of the AI-costs over the horizon: a decreasing trajectory (the baseline), a flat trajectory and a rising trajectory. The trajectory determines for a large part how far the automation reaches, and in a similar direction for all outcomes: a decreasing trajectory gives the highest inequality and unemployment, a rising trajectory brings the automation to a stop and stable, and the flat trajectory is in between (Figure 7). The difference in effect is large. Notably, the flat trajectory is not stabilised at the end of the horizon while the rising and decreasing are stable (Table 8). Against the decreasing baseline the rising trajectory differs strongly and over almost the whole horizon ($d_z \approx 10$ on the Gini and unemployment, significant on 99%), and the flat trajectory differs moderately ($d_z \approx 3.5$, 99%); on the skill-wage premium the separation is smaller, and somewhat larger for the flat than for the rising trajectory ($d_z \approx 3.7$, 94% versus ≈ 2.0 , 78%).

The trajectory is set using component C5 and works in two places. In the direct cost comparison of C1 the lower AI costs lower the cost per unit (for AI), which causes more tasks to be able to be automated, in the NPV of component C4 a decreasing trajectory causes higher expected savings so investments, even with a hurdle rate, become profitable. The macro-feedback M1 determines the plateau (together with the AI productivity): the more there is automated, the lower the wage bill and the demand (C6) and the sooner the expansion gate (C1) shuts, which explains why the decreasing trajectory finally plateaus while the flat one is rising.




Figure 7: Gini-index, Skill-wage premium and Unemployment rate over time of the three cost trajectories (decreasing/baseline, flat, rising). Average over 50 seeds with a ± 1 standard deviation band.

Table 8: Paired comparison of the baseline (A = Decreasing) with the flat and rising trajectory. Difference = A-B, positive means a higher value for the decreasing trajectory. Complete matrix in Appendix B.

Configuration	Outcome	Difference (A-B)	95%-CI	d_z	Sign. horizon
Flat	Gini	+0.088	[0.081, 0.095]	3.53	99%
Flat	Skill-wage premium	+0.046	[0.043, 0.050]	3.67	94%
Flat	Unemployment	+0.088	[0.081, 0.095]	3.47	99%
Rising	Gini	+0.148	[0.144, 0.152]	10.26	99%
Rising	Skill-wage premium	+0.021	[0.018, 0.024]	2.02	78%
Rising	Unemployment	+0.157	[0.153, 0.161]	9.89	99%

5.5. Experiment 5 - Comparative advantage

Experiment 4 moved the trajectory of AI costs; the fifth experiment turns to where that cheaper AI bites, varying the non-routine multiplier on the values 0.50 (AI is clearly weaker for non-routine tasks), 0.85 (baseline), 1.15 and 1.50 (AI is clearly stronger for non-routine tasks). The Gini-index and the unemployment rate increase with the value of the multiplier (Figure 8): stronger AI on non-routine tasks leads to more automation and higher inequality and unemployment. The skill-wage premium responds non-monotonically. Only when AI is clearly weaker on non-routine tasks (0.50) does the premium sit clearly below the baseline; for 1.15 and 1.50 it hardly differs from the baseline. So the direction of the inequality effect is clear, but whether it is combined with polarisation or compression depends on the multiplier (Table 9).

The multiplier is in the AI-productivity function of component C1 (section 3.6) and sets the productivity of AI on non-routine tasks relative to routine tasks, which are the fixed reference. Through the cost comparison of C1 it moves which tasks are profitable to automate, which continues in the adoption decision of component C4. Therein lies the non-monotonicity of the premium: a low multiplier shrinks the set of non-routine tasks worth automating, so the routine tasks performed by both skill groups (but with a preference for the low-skilled because of the wage level) make up a larger share of the automation that remains, while a high multiplier extends automation into the non-routine tasks mainly performed by the highly skilled. Which skill group bears the displacement shifts with the multiplier, and that shift then determines, through the composition of employment (C2) and the wage curve per skill level (C3), whether the premium widens or narrows.

The employment rates per skill level make this mechanism visible. In both Figure 9a and Figure 9b the curves are ordered by the multiplier, low (0.50) on top and high (1.50) at the bottom, so a higher multiplier lowers employment for both skill groups. The scale effect therefore dominates the composition effect: raising the multiplier scales AI productivity on non-routine tasks while leaving routine tasks at their reference level, which lifts average AI productivity across the task set, widens the range of cost-competitive tasks (C1), raises adoption (C4) and cuts total labour demand (C2). The low-skilled (Figure 9a) span a wide range, from $\approx 90\%$ (low) to $\approx 74\%$ (high); notably they remain highest under the low multiplier even though the routine tasks they perform then make up a larger share of what is automated, because the total volume of automation is far smaller. The high-skilled (Figure 9b) sit structurally lower ($\approx 40\text{--}50\%$) and are hit hardest by the high multiplier, which makes AI competitive on their non-routine tasks ($\approx 40\%$ versus $\approx 46\%$ at the baseline). This also explains the non-monotonic premium. Both employment rates move in the same direction, but not in the same proportion: between the high and the low multiplier the low-skilled gain some forty workers against twenty for the high-skilled, so the low-skill wage rises relatively more than the high-skill wage through the wage curves of C3. The premium therefore compresses at the low multiplier (Table 9), while at 1.15 and 1.50 the two groups move closely enough together for the ratio to stay near the baseline.




Figure 8: Gini-index, Skill-wage premium and Unemployment rate over time of the four configurations (low, baseline, medium, high). Average over 50 seeds with a ± 1 standard deviation band.




Figure 9: Employment rate over time per skill group for the four configurations (low, baseline, medium, high). Average over 50 seeds with a ± 1 standard deviation band.

Table 9: Paired comparison of the baseline ($A = 0.85$) with the low, medium and high comparative advantage. Difference = $A-B$, positive means a higher value for the baseline. Complete matrix in Appendix B.

Configuration	Outcome	Difference (A-B)	95%-CI	d_z	Sign. horizon
Low (0.50)	Gini	+0.082	[0.079, 0.086]	5.83	100%
Low (0.50)	Skill-wage premium	+0.044	[0.042, 0.046]	5.11	92%
Low (0.50)	Unemployment	+0.082	[0.078, 0.086]	5.53	100%
Medium (1.15)	Gini	-0.026	[-0.031, -0.021]	-1.51	64%
Medium (1.15)	Skill-wage premium	-0.007	[-0.013, 0.000]	-0.26	14%
Medium (1.15)	Unemployment	-0.028	[-0.032, -0.023]	-1.59	66%
High (1.50)	Gini	-0.049	[-0.058, -0.041]	-1.53	100%
High (1.50)	Skill-wage premium	+0.013	[0.001, 0.024]	0.30	5%
High (1.50)	Unemployment	-0.056	[-0.064, -0.048]	-1.83	100%

5.6. Experiment 6 - Minimum wage

The remaining experiments turn to the Dutch institutional levers. The sixth experiment varies the minimum wage over four values: no floor (0), a low level (10), the Dutch statutory value of 14.71 (baseline) and a high level (20). Only a higher minimum wage has a remarkable effect, removing or lowering the floor to 10 hardly differs from the baseline as visible in Figure 10. For a higher minimum wage the Gini and the unemployment are higher than the baseline, while the skill-wage premium is lower (Table 10).

The minimum wage is the wage floor component in C3 (section 3.8). A high floor raises the cost per unit for human labour for low skilled (routine) tasks in the cost comparison of C1, because of that in the NPV calculation of component C4 automation becomes profitable faster, the channel of technical change (Acemoglu, 2002). For the employed a higher floor compresses the wage gap, but the displaced workers become unemployed, with an income of 0 in the Gini calculation. Because of that the unemployment rate and Gini rise together. The fact that the other values have no effect shows that the floor only becomes a bottleneck in the cost comparison when it is above the market wage of low-skilled workers.




Figure 10: Gini-index, Skill-wage premium and Unemployment rate over time of the four configurations (base, no minimum wage, low minimum wage, high minimum wage). Average over 50 seeds with a ± 1 standard deviation band.

Table 10: Paired comparison of the baseline ($A = 14.71$) with the three alternative minimum wage levels. Difference = $A-B$, positive means a higher value for the baseline. Complete matrix in Appendix B.

Configuration	Outcome	Difference (A-B)	95%-CI	d_z	Sign. horizon
No minimum (0)	Gini	-0.001	[-0.006, 0.004]	-0.05	3%
No minimum (0)	Skill-wage premium	+0.010	[0.007, 0.012]	1.07	8%
No minimum (0)	Unemployment	-0.003	[-0.008, 0.003]	-0.13	6%
Low (10)	Gini	-0.003	[-0.009, 0.003]	-0.13	4%
Low (10)	Skill-wage premium	+0.009	[0.006, 0.012]	0.79	8%
Low (10)	Unemployment	-0.005	[-0.011, 0.002]	-0.20	6%
High (20)	Gini	-0.030	[-0.035, -0.025]	-1.73	90%
High (20)	Skill-wage premium	+0.038	[0.035, 0.042]	2.78	95%
High (20)	Unemployment	-0.034	[-0.039, -0.029]	-1.85	95%

5.7. Experiment 7 - Severance payment

Experiment 3 toggled employment protection as a whole; the seventh experiment isolates its main internal margin, the severance rate, over no payment (0), the statutory one-third (baseline) and a high level (1.0) (Figure 11, Table 11). The two directions are strongly asymmetric. Removing the severance payment entirely (0 against the statutory rate) raises the Gini and the unemployment rate by a small-to-medium amount (d_z around 0.8, significant on about 35% of the horizon), almost the same difference as turning the whole institution off in Experiment 3 (d_z around 0.7, 35%). Raising the rate above the statutory level (1.0 against 1/3), by contrast, barely changes the outcomes ($|d_z|$ below 0.35, and near zero on the Gini and unemployment).

Both directions point the same way because the severance payment is the only cost channel of component C7 (section 3.12): it enters the NPV of C4 as a one-off upfront cost, and only for permanent workers. With the rate at zero the chain regulation and the permanent contracts still exist but impose no cost, so the run behaves almost like the institution-free MF-off of Experiment 3, which is why removing the rate reproduces nearly the whole on/off effect. The saturation in the other direction, where a higher rate adds nothing, shows that at the statutory level the redundancy cost is already small relative to the AI savings and only rarely the decisive factor: tripling it tips no further decisions, while removing it lifts the brake on the decisions where it was pivotal. The institutional effect of Experiment 3 is thus, to first order, the effect of the statutory severance payment itself.




Figure 11: Gini-index, Skill-wage premium and Unemployment rate over time of the three configurations (base, no severance, high severance). Average over 50 seeds with a ± 1 standard deviation band.

Table 11: Paired comparison of the baseline ($A = 1/3$) with the two alternative severance rates. Difference = $A-B$, positive means a higher value for the baseline. Complete matrix in Appendix B.

Configuration	Outcome	Difference (A-B)	95%-CI	d_z	Sign. horizon
No (0)	Gini	-0.015	[-0.020, -0.010]	-0.82	35%
No (0)	Skill-wage premium	-0.006	[-0.008, -0.003]	-0.65	10%
No (0)	Unemployment	-0.016	[-0.021, -0.011]	-0.83	34%
High (1.0)	Gini	-0.000	[-0.005, 0.004]	-0.01	17%
High (1.0)	Skill-wage premium	-0.003	[-0.005, -0.001]	-0.34	5%
High (1.0)	Unemployment	+0.000	[-0.005, 0.005]	0.01	18%

5.8. Experiment 8 - Transition probability

The eighth experiment stays on the same severance channel from the contract-mix side, varying the probability p_{convert} that a flex worker upon reaching the chain limit is converted to a permanent contract, over 0.0, 0.25, the baseline 0.5, 0.75 and 1.0. The effect on the three primary outcomes is small over the complete horizon, with a consistent direction (Figure 12, Table 12): less conversion, so a larger part being flex workers, goes together with a slightly higher inequality and unemployment rate, while more conversion than the baseline gives negligible results.

The transition probability determines, through the chain regulation in C7, the contract composition. More conversion means more permanent contracts and thus a higher future severance obligation, which in the NPV of component C4 will slow down the proactive automation. The mechanism is the same as experiment 7, and the outcome is thus consistent: because the severance costs are small with respect to the AI savings, a different contract mix hardly changes the aggregated outcomes.




Figure 12: Gini-index, skill-wage premium and unemployment rate over time of the five transition-probability configurations (base, low, no, high and complete conversion). Average over 50 seeds with a ± 1 standard deviation band.

Table 12: Paired comparison of the baseline ($A = 0.5$) with the two alternative transition probabilities. Difference = $A - B$, positive means a higher value for the baseline. Complete matrix in Appendix B.

Configuration	Outcome	Difference (A-B)	95%-CI	d_z	Sign. horizon
None (0.0)	Gini	-0.011	[-0.015, -0.006]	-0.67	24%
None (0.0)	Skill-wage premium	-0.001	[-0.003, 0.001]	-0.12	8%
None (0.0)	Unemployment	-0.012	[-0.016, -0.007]	-0.70	26%
Low (0.25)	Gini	-0.009	[-0.014, -0.004]	-0.47	3%
Low (0.25)	Skill-wage premium	-0.003	[-0.005, -0.000]	-0.28	4%
Low (0.25)	Unemployment	-0.009	[-0.014, -0.004]	-0.47	3%
High (0.75)	Gini	+0.003	[-0.001, 0.007]	0.20	2%
High (0.75)	Skill-wage premium	-0.002	[-0.004, -0.000]	-0.30	4%
High (0.75)	Unemployment	+0.003	[-0.000, 0.007]	0.24	4%
Complete (1.0)	Gini	+0.003	[-0.002, 0.008]	0.16	12%
Complete (1.0)	Skill-wage premium	-0.004	[-0.006, -0.001]	-0.42	9%
Complete (1.0)	Unemployment	+0.004	[-0.002, 0.009]	0.20	17%

5.9. Across the experiments

Ordered by effect size, the eight experiments fall into a clear hierarchy. The largest and most persistent differences come from the technology axis: replacing the myopic ULC logic with the forward-looking NPV decision (Experiment 1's baseline-to-ULC gap, $d_z \approx 9.8$ on the Gini and unemployment) and varying the AI-cost trajectory (Experiment 4, d_z up to 10.3) each move the outcomes by large margins over the full horizon. The decomposition (Experiments 2 and 3) places these on opposite ends of the hierarchy: the forward-looking decision logic (Experiment 2) ranks with the large technology levers, while the employment-protection institutions (Experiment 3) rank among the small ones. The comparative advantage of AI across tasks (Experiment 5) is the next strongest lever; it sets the direction of the inequality effect, with the scale effect on total labour demand dominating the composition effect between skill groups. The Dutch institutions sit lower, but not uniformly. Turning employment protection on or off as a whole (Experiment 3) produces a small-to-medium effect ($|d_z|$ up to 0.90 on the premium) that is significant at both the start and the end of the horizon, so the institution is a modest but genuinely persistent brake rather than a negligible one. Within that institution the severance payment carries the effect: removing it (Experiment 7) reproduces almost the whole on/off difference of Experiment 3, whereas raising it above the statutory level adds nothing and the conversion probability (Experiment 8) barely moves the outcomes. The institutional effect is thus, to first order, the statutory severance payment: present at its statutory level, but neither growing with a higher rate nor depending much on the contract mix. The minimum wage (Experiment 6) is the one institution with a clear effect, and only once it rises above the market wage of low-skilled workers: there it compresses the wage premium for those who keep their jobs while pushing the displaced into unemployment. One internal mechanism recurs across all eight experiments: the relative price of AI, set by its cost path (C5, which in the baseline falls exogenously towards a floor below labour cost) and its task-level productivity (C1), governs how far automation reaches through the NPV decision (C4), and the labour-market institutions bend that path only when they enter the cost comparison; employment protection enters early, through the severance payment on occupied tasks, and the Keynesian demand feedback then amplifies that small early difference into the visible late level gap. The discussion of the results follows in chapter 6.

6

Discussion

This chapter moves from output to understanding: for each mechanism it establishes what drives the result inside the model and then, explicitly separated, what the mechanism claims about the labour market itself, what would have to be true for that claim to hold, and where it could be checked against reality. chapter 7 then moves from understanding to action: what follows for Dutch policy, what should be monitored, and what should be researched next.

The order is as follows. A summary of the main findings (section 6.1) is followed by the interpretation per mechanism (section 6.2), a reflection on what the agent-based approach did and did not contribute (section 6.3), an explicit treatment of the role of uncertainty (section 6.4), and the limitations (section 6.5), which close with what those limitations mean for reading the results as a statement about the Netherlands. Throughout, the model is treated as a scenario-exploratory instrument rather than a point forecast: it describes directions, feedbacks and sensitivities around a calibrated baseline, and the numbers it produces are not predictions of Dutch wage inequality or unemployment.

6.1. Summary of the main findings

The eight experiments and the OFAT analysis (Appendix G) together produce a consistent picture. The common thread is that the outcomes, the Gini index of worker income, the skill-wage premium between high- and low-skilled workers, and the unemployment rate, are driven mainly by the relative price and the comparative advantage of AI, and much less by the Dutch contract institutions. The distribution of displacement across skill groups, which skill group bears the burden of automation, is treated throughout as the mechanism that drives these two headline outcomes, wage inequality and unemployment, rather than as a separate outcome in its own right.

Experiment 1: Endogenous NPV-adoption versus the ULC-logic of Upreti & Sridhar (2024). The short-sighted ULC-decision leads to structurally higher inequality and unemployment, and the gap grows over the horizon instead of closing; on inequality and unemployment the difference is large and holds over the whole horizon, on the wage structure it is considerably smaller (Table 5). What the two configurations share is the shape of the path: both rise steeply and then flatten. The forward-looking rule does not change the character of the trajectory but how far it runs before the brakes bind, because the profitability threshold, the expected wage decline through displacement and the Keynesian feedback work together. Experiments 2 and 3 decompose this further.

Experiment 2: Decision logic (the forward-looking decision). With the institutions switched off on both sides, the forward-looking NPV automates far less than the myopic ULC rule, lowering inequality and unemployment over the whole horizon by a wide margin, with a smaller separation on the wage structure (Table 6). This step accounts for the large majority of the headline gap: it is the forward-looking investment logic, far more than the institutions, that does the tempering.

Experiment 3: Employment protection (on/off). Turning the Dutch employment-protection regime on, against an otherwise identical run with it off, lowers inequality, unemployment and the skill-wage

premium, but only by a small margin: the two trajectories run within each other's variation bands for most of the horizon and separate visibly only late (Figure 6, Table 7). The effect is statistically robust at both the start and the end of the horizon, with an insignificant middle, but it is a small difference rather than a different trajectory. Protection slows the first wave of automation through the severance cost on occupied tasks, the per-step gap then settles once the AI price floors, and the late re-opening is the Keynesian demand feedback amplifying that early difference, not a re-strengthening severance brake. The institution is therefore a modest but persistent brake; the spread across seeds widens considerably towards the end of the horizon, so the short-run brake is the robust part of this finding.

Experiment 4: AI cost trajectory. The course of the AI costs does not shift one curve up or down, it decides which of three qualitatively different futures the economy ends in (Figure 7, Table 8). Under a rising cost path automation comes almost to a halt and inequality and unemployment stay flat at close to their starting level for the entire horizon. Under the decreasing baseline path they climb steeply and then settle at a much higher plateau. The flat path lies between the two but, unlike the other two, has not stabilised by the end of the horizon: it is still rising when the simulation stops. This is the sharpest contrast the experimental programme produces.

Experiment 5: Comparative advantage. The non-routine multiplier sets how far automation reaches across the task set, and with it the level of inequality and unemployment (Table 9): a higher multiplier raises both, a low multiplier holds both down markedly. Because the multiplier scales AI productivity on non-routine tasks while leaving routine tasks at their reference level, it widens the automatable set rather than only redirecting it, so both skill groups lose employment together: the scale effect (how many tasks in total become cheap enough to automate) dominates the composition effect (which of the two skill groups those tasks belong to). The skill-wage premium therefore responds non-monotonically: it compresses clearly only under the low multiplier, where the low-skilled gain proportionally more employment than the high-skilled, while at higher values it stays close to the baseline.

Experiment 6: Minimum wage. Only a high minimum wage has a noticeable effect (Figure 10, Table 10): removing or lowering the floor is nearly indistinguishable from the baseline, while a high floor raises inequality and unemployment early on and lowers the skill-wage premium throughout. Notably, the inequality gap it opens narrows again towards the end of the simulation period, so the high floor accelerates the transition more than it changes its destination. The minimum wage only bites once it rises above the market wage of the low-skilled and then becomes a bottleneck in the cost comparison.

Experiment 7: Severance payment and Experiment 8: Transition probability. The severance rate has a strongly asymmetric effect (Table 11). Removing the severance payment entirely reproduces almost the whole on/off effect of Experiment 3, so the institution's effect is essentially the severance payment itself; but raising the rate above the statutory level adds nothing, and the conversion probability (Experiment 8) leaves the trajectories practically indistinguishable (Figure 12). It is therefore the presence of the statutory severance payment, not a higher rate or a particular contract mix, that carries the institution's (small) effect. This is in itself a meaningful outcome, interpreted in detail below.

OFAT-sensitivity analysis (Appendix G). The Gini index is the most sensitive outcome and is dominated by parameters that set the automation level and therewith the unemployment rate. Besides the number of firms (partly a scale variable), those are mainly the AI parameters: the baseline AI productivity (ϕ_{base}), the long-run cost floor (k_{ai_floor}), the investment cost (I_{base}), the complexity scaling and the planning horizon ($T_{horizon}$). The contract- and protection parameters ($severance_rate$, $chain_limit$, $p_convert$) are at the bottom of the ranking.

6.2. Interpretation of the results

Each mechanism below is treated in two steps. First, what drives it inside the model: which components interact to produce the pattern reported in chapter 5. Second, and set apart under outside the model, what the mechanism claims about the labour market itself: which of its components have real-world counterparts, what would have to be true for the claim to hold, and where it could be checked or falsified. The second step is where the model is deliberately let go of; chapter 7 then takes the further step from understanding to action.

6.2.1. The relative price of AI shapes the outcomes

The most powerful explanatory variable in the model is the price of AI relative to labour, which outweighs every institution. Experiment 4 shows that both the level and the trajectory of the AI costs strongly influence the automation level, and therewith the unemployment and the inequality. The OFAT analysis points the same way at parameter level: after the scale parameters, whose ranking position mixes a model-technical with an economic effect (section 6.5), the baseline AI productivity (ϕ_{base}) and the long-run cost floor (k_{ai_floor}) move the Gini index most, with the planning horizon ($T_{horizon}$) close behind (Appendix G). A higher AI productivity or a lower cost floor makes automation more attractive, raises unemployment and pushes the Gini up, because the unemployed have an income of zero in the calculation.

Conceptually this dynamic runs through the Keynesian feedback. More automation raises unemployment, which through the wage curve presses wages down; lower wages lower the wage bill and thus aggregate demand, which presses the price and erodes firm profitability. When average cost approaches the price, the expansion gate closes: firms no longer expand proactively. In the baseline firms operate relatively close to that profitability frontier, so this brake activates regularly. This is precisely the mechanism that distinguishes the mean-field NPV from the ULC-logic, which misses this slowdown and therefore overshoots to a much higher automation level.

It is worth being explicit about how much of this ordering is a result and how much is built into the model. Three architectural choices push towards technology dominance before any experiment is run: the AI-cost curve is exogenous and decoupled from adoption, and in the baseline declines towards a floor that lies below the cost of labour (assumptions C5.2 and C5.4 in Appendix H); automation is all-or-nothing displacement, with no augmentation or reinstatement to return labour demand (assumptions A0.12 and A0.8); and the baseline is pattern-calibrated so that firms operate near the profitability frontier (assumption A0.13). Given a technology engineered to become and remain cheaper than labour in a displacement-only world, it is close to inevitable that the price of AI outweighs an institution priced as a modest one-off severance wedge. The robust contribution is therefore less the bare ranking “price ahead of institutions” than the demonstration that even a forward-looking, institutionally-frictioned economy still automates heavily once AI is this cheap. Whether the ordering survives a flat or rising AI-cost path, or an augmentation channel, is left open and recommended for further research (section 7.4).

Outside the model. The claim is narrower than “technology beats institutions”, and the model states it in terms a firm would recognise. An automation decision weighs three quantities against each other: the upfront cost of implementing and integrating AI on a task, $I(x)$; the running cost of using it afterwards, k_{ai} , against the wage of the person doing the work; and the one-off cost of ending that employment, $S(x)$. Automation follows only when the discounted stream of running-cost savings exceeds all three. What the model finds is that this decision is settled by the first two: the severance wedge is small enough that it flips only the decisions that were already marginal. That is a claim about magnitudes, checkable against real implementation budgets and real transition payments. It also locates where the public discussion misleads. The price that falls fastest and is quoted most often, the cost of inference, is only part of the running cost and none of the upfront cost; implementation and integration do not fall at the same rate, and in some early deployments they have proved large enough to reverse the business case altogether (section 6.4). Whether AI is cheaper than labour is therefore an empirical question about implementation cost in a specific firm and a specific task. The model contains that quantity as I_0 , and had to pattern-calibrate it (Appendix C) precisely because the realised cost of putting AI on a task is not systematically observed.

6.2.2. Automation is a level effect, not a growth effect

A feature visible in almost every figure of chapter 5 deserves to be named explicitly, because it is easy to read past and it changes how the results should be understood. The outcome paths are not trends. Inequality and unemployment rise steeply in the first decades and then flatten, settling at a plateau that they hold for the remainder of the eighty-year horizon (Figures 3 and 7). AI adoption in this model produces a one-off shift from one level to another; it does not set the economy on a permanently worsening path.

The plateau is not an artefact but the model reaching an equilibrium in the cost comparison. Automation

continues only while there remain tasks on which AI is cheaper than the worker performing them. Each task automated removes a worker from employment, which through the wage curve pushes wages down, which makes the remaining human-performed tasks cheaper to keep. Meanwhile the AI cost path approaches its floor and stops falling. The two prices converge: at the margin, the cost of a human and the cost of AI come to break even, and at that point no further task clears the investment threshold. The economy settles at whatever level of automation, unemployment and inequality that break-even point implies. Where the break-even sits depends on the parameters, which is why the different scenarios plateau at different heights, and why the rising-cost path in Experiment 4 never leaves its starting level at all: there the two prices never cross. The flat-cost path is the instructive exception, still climbing when the simulation stops, because a constant AI price against slowly falling wages keeps opening marginal tasks without ever settling.

This matches what the wider literature expects of a general-purpose technology of this kind. Acemoglu (2025) argues that when AI's microeconomic effects run through task-level cost savings, its aggregate consequences follow a version of Hulten's theorem: the gains are determined by the share of tasks affected and the average cost saving on them, a one-time increment to the level of productivity rather than a permanent increase in its growth rate. The mechanism in this model is the labour-market counterpart of that argument, and it arrives at the same shape from the opposite direction: a bounded set of exposed tasks yields a bounded, one-off displacement effect.

One consequence for interpretation follows immediately: the horizon matters less than it appears. Extending the simulation beyond eighty years would not change the outcomes much, because the system has already found its equilibrium, and the analytically interesting behaviour is concentrated in the transition rather than in the end state.

Outside the model. If this shape is right, the public question is badly posed. The debate asks whether AI will cause mass unemployment, as though the answer were a direction; the model says the direction is not the issue, because the process ends by itself. What matters is how large the one-off step is, and how long an economy takes to absorb it, and those are separate questions with separate answers. The size of the step is set by where the two costs come to break even, which depends on how far AI capability extends across tasks and what it costs in practice to deploy. The absorption time depends on retraining, mobility and the creation of new work, none of which this model contains. A labour market could therefore face a large step that it absorbs well, or a modest step that it absorbs badly, and only the second is really a policy failure.

This reframes what the institutions examined here are for. Instruments that shorten or soften the transition and instruments that move the break-even point itself are doing different jobs: the first buys time to absorb a step whose size is fixed, the second changes the size of the step. Dutch dismissal protection belongs firmly in the first category, which is why it delays without changing the destination; a price instrument would belong in the second. The historical record is consistent with this division rather than proof of it. Earlier general-purpose technologies also displaced labour in bounded waves rather than pushing unemployment permanently upward, and the hardship came mainly from the pace and the incidence of the reallocation. The analogy should be drawn carefully, though: those transitions were absorbed in large part through the creation of new tasks, a channel this model does not represent. History therefore speaks to the absorption side of the problem and says little about where the break-even point will sit this time.

6.2.3. Comparative advantage sets how far automation reaches

Experiment 5 sits in the task-based tradition of Autor, Levy and Murnane (2003) and Acemoglu and Restrepo (2018, 2019), in which technology is allocated to tasks rather than to jobs. The non-routine multiplier is the model's version of that idea: it fixes AI productivity on routine tasks and scales it on non-routine tasks, so a low value means AI is weak on non-routine work rather than strong on routine work.

The result does not reproduce Upreti and Sridhar (2024), and the difference is instructive. They find that automating routine tasks raises wage inequality while automating non-routine tasks lowers it. Here, a higher multiplier raises both the Gini index and unemployment, and the wage premium compresses at the low multiplier rather than the high one. The reason is that their comparison shifts automation

between task types at a given level, whereas this design also changes how much is automatable at all: scaling one side while holding the other fixed raises average AI productivity, so a higher multiplier widens the automatable set instead of merely redirecting it. Both skill groups then lose employment together. The scale effect dominates the composition effect, but partly by construction, which is why this result should be read as a property of the design as much as a finding about comparative advantage.

Outside the model. The transferable claim is that what matters is where AI is strong, not how strong it is. The model implements that on one axis, routine versus non-routine, with the high-skilled placed on non-routine work. Generative AI does not respect that axis. Its strengths lie in drafting, summarising, coding and standard analysis: cognitively routine work performed largely by well-paid, well-educated people. Much of the work that resists it, care, cleaning, installation and hospitality, is manual, non-routine and low-paid (Eloundou et al., 2024). Read literally, the model points automation at the bottom of the wage distribution; corrected for this mismatch, the same mechanism points considerably higher up. What survives the translation is the structural point: when a technology becomes profitable across a broader set of tasks, both skill groups lose employment together rather than one gaining at the other's expense. The damage then runs through unemployment rather than through the wage gap among those still working, which makes the wage premium a poor early-warning indicator in reality as well as in the model.

The underlying mechanism has direct empirical support, though from the previous wave of automation. Acemoglu and Restrepo (2022) attribute between 50 and 70 per cent of the change in the US wage structure over four decades to the relative wage decline of workers specialised in routine tasks in rapidly automating industries. Task displacement is therefore not a modelling device but a measured driver of realised inequality. What is open is whether generative AI hits the same groups. Gmyrek et al. (2023) place the highest exposure in clerical occupations in high-income countries, which sit in the middle rather than at the bottom of the wage distribution. If that holds, routine displacement continues but climbs the wage ladder, and a model that ties automation to low-skilled routine work by construction will locate the damage lower down than reality does.

6.2.4. The minimum wage as the sole binding institution

Of the policy instruments examined, the minimum wage is the only one with a clear, albeit asymmetric, effect (Experiment 6). The explanation lies in the cost comparison of component C1: a high wage floor raises the unit cost of human labour for low-skilled routine tasks, so that automation becomes profitable for more tasks in the NPV calculation, the channel of induced technical change of Acemoglu (2002). For workers a higher floor compresses the wage gap, but displaced workers become unemployed with an income of zero, so unemployment and Gini rise together. That lower levels of the floor have no effect shows that the minimum wage only binds once it rises above the market wage of the low-skilled; this matches the Dutch situation in which the minimum wage binds people at the lower end of the labour market.

Outside the model. This is the model's clearest example of induced technical change (Acemoglu, 2002): raising the price of a specific input makes replacing it more attractive. Whether the mechanism is currently active in the Netherlands is an empirical question, and the model deliberately cannot answer it, because it has one sector and one low-skilled wage while the real question is occupation-specific. What can be said is what would have to be true for the mechanism to bite: the wage floor would have to sit above what the market would otherwise pay for that work, and the work would have to be technically substitutable. Both conditions must hold together, which is why the effect appears only at the high setting in the model and disappears entirely at the low ones. That conditionality is the finding, not the magnitude, and it offers one reading of why studies of minimum-wage effects on automation reach mixed results: they average over occupations in which the two conditions do and do not coincide. It also connects to a broader fiscal point. Acemoglu, Manera and Restrepo (2020) show that the relative tax treatment of labour and capital can by itself tilt firms towards automation, and a statutory wage floor is one more input into that same relative price. A wage floor and a capital-tax allowance push the automation decision in the same direction.

6.2.5. Why the contract institutions hardly bite

The contract institutions, the severance payment (*transitievergoeding*) and the chain regulation, produce an effect with three features. It is small in absolute terms but statistically robust (Experiment 3). It is essentially the severance payment: removing that alone (Experiment 7) reproduces almost the whole on/off difference, while raising the rate above the statutory level adds nothing and the contract mix (Experiment 8) barely matters. And it is persistent, significant at both ends of the horizon.

Three mechanisms account for this, worked out with the cost example and the diagnostic series in Appendix B (subsection B.5.1). The severance payment is small relative to the upfront implementation cost, so it can only flip adoption decisions that were already marginal, which is why the effect saturates at the statutory rate. The modelled contract build-up, in which workers start as flex and acquire entitlement only at the chain limit (assumption C7.4 in Appendix H), damps it further, and here the model understates the real institutional asymmetry. The persistence, finally, is not a re-strengthening brake: severance owed per task peaks early and then declines, and what re-opens the gap late in the horizon is the Keynesian demand feedback (subsection 3.11.2) amplifying a small early difference in displacement.

Employment protection is therefore neither a purely short-run instrument nor a strong structural one: it buys transition time early and, because the demand feedback carries that difference forward, also modestly lowers the long-run level of automation, inequality and unemployment. It slows the journey and shifts the destination a little. The size of that long-run shift depends on how close firms operate to the profitability frontier and is heavy-tailed across seeds, so the short-run delay is the more robust half of the conclusion (section 6.5).

This nuances the hypothesis from the introduction, drawn from Acemoglu (2002) and the productivity results of Vergeer and Kleinknecht (2010): rigidities that keep labour costs high and predictable would steer investment towards labour-saving technology, and in the Dutch dual model AI would replace the flexible layer first. The steering effect is present in direction, since a higher wage floor does accelerate automation, but contract protection is too weak relative to the AI price to redirect adoption by more than the modest margin of Experiment 3. The flex-first sequence is moreover a structural assumption here rather than an emergent outcome, so the model can only partly test the original hypothesis.

Outside the model. The reason the institution barely moves the outcome is a matter of proportion, and the proportion is real rather than modelled. A Dutch transition payment amounts to a third of a monthly salary per year of service: for a worker with a few years of tenure, some weeks of pay. The investment it is set against is the cost of automating that person's work, software, integration, process redesign and the internal capacity to run it. If those two are an order of magnitude apart in practice, as they are in the model, then employment protection cannot plausibly be the variable that decides whether a task is automated, and no reasonable adjustment of the statutory rate would change that. This is a claim anyone can check against real severance costs and real implementation budgets, and it does not depend on the calibration.

Two consequences deserve emphasis because they run against how the instrument is usually discussed. The first is that the protection works through timing rather than prevention: in the model, firms do not dismiss protected workers to automate, they wait for someone to leave and then do not replace them. Dismissal protection therefore does not save the job; it decides when the job disappears and gives the person holding it the chance to leave on their own terms. The second is that the effect saturates. Because severance can only tip decisions that were already marginal, once the statutory level has captured those, a higher level finds nothing left to tip. Strengthening dismissal protection in response to AI would therefore be expected to yield little, which is an unusual conclusion to draw from a model built to take Dutch institutions seriously, and is worth stating precisely because it was not the expected one.

6.3. Reflection on the agent-based approach

An agent-based model is an expensive instrument: it must be built, verified, calibrated and run over thousands of seeds before it produces a single comparable number. What did that buy, beyond what a simpler or aggregate model would have produced?

The answer starts with the investment rule, because that rule is what the whole model is organised around, and it does not work on averages. Whether automating a given task clears the hurdle depends on the complexity of that task, on the wage of the specific worker performing it, on whether that worker holds a permanent or a flexible contract, and on how much tenure has accumulated. Two firms facing an identical AI price therefore reach opposite conclusions about otherwise comparable tasks, because their task composition and their workforce differ. What the model holds at any moment is not one decision but a distribution of decisions, of which only part sits close enough to the threshold to be moved. This matters most for the institutional experiments: a severance payment, a wage floor or a contract rule does not enter as an economy-wide cost but is recalculated inside every firm against its own tasks and its own workers, so the same instrument bites differently across the population. An aggregate formulation collapses that distribution into a single representative decision, and a representative firm either automates or does not.

The consequence is a fairly precise division of what does and does not require this approach. Where the economy ends up is a threshold condition, and thresholds can be evaluated on averages: a model without heterogeneous firms could plausibly locate roughly the same plateau, since that plateau is where the running cost of AI and the cost of labour come to break even. This is visible in the results themselves. In Experiment 1 the baseline and the ULC configuration trace the same trajectory and differ in how high the plateau sits (Figure 3); the same holds for Experiment 2 and, more strongly, for the institutional levers of Experiments 3, 7 and 8, where the paths overlap within their variation bands for most of the horizon (Figures 6, 11 and 12). Even Experiment 4, substantively the sharpest result in the programme, is a threshold being crossed or not: under a rising cost path automation never clears the hurdle and the economy does not move for eighty simulated years, while under the falling baseline it climbs and settles much higher (Figure 7). The price of AI selects which future the model ends in, but an aggregate model with the same investment rule could show that switch.

The route to the plateau is a different matter, and that is where the institutions live. How quickly automation gets there, and whether a policy accelerates or delays it, depends on which specific decisions are near the threshold at which moment, and that is case-specific by construction. An average severance cost per worker cannot represent it: the entitlement is a step function of individual history rather than a smooth cost. Under the chain regulation a worker carries no severance obligation for thirty-six months and then, on conversion, carries one that grows with tenure. Whether automating a task is cheap or expensive therefore depends on when the person occupying it was hired and whether they have separated in the meantime.

Two results follow from this and would not survive the averaging. The first is the saturation of the severance effect (subsection 6.2.5): the statutory rate already captures the decisions that were marginal, so raising it finds none left to move. That result presupposes a distribution of decisions with a margin in it. The second is the temporal signature of Experiment 3, significant at both ends of the horizon with an insignificant middle, which traces a heterogeneous warm-start cohort working itself out: high-tenure workers whose severance bites early, who then separate, leaving a workforce whose entitlement rebuilds from zero. A representative-agent formulation has no cohort to work out and would show a constant gap or none. A third contribution is the spread across runs: the variation bands are what reveal that the late employment-protection gap rests on a minority of runs in which the demand loop runs away, while the early brake appears in almost all of them, a distinction invisible to a model reporting one mean path.

The claim should not be stretched further than that. No deterministic benchmark was run alongside the ABM, so the comparison is argued rather than demonstrated. The argument on which the original choice rested, that heterogeneous firms and workers, bounded rationality, institutional friction and an endogenous decision rule are the features representative-agent models set aside (Fagiolo & Roventini, 2017; Tesfatsion, 2006), was made before the model was run (section 2.3). It justifies the choice without measuring what the choice delivered. What can be said with confidence is narrower and, for a thesis about policy instruments, more useful: the destination may be within reach of a simpler model, but the transition, which is what the institutions act on, is not.

Two results, finally, ran against expectation, and both were made measurable by the design rather than assumed by it. The hypothesis taken from the induced-technical-change literature was that the Dutch employment-protection institutions would be a first-order force steering adoption. They are not: their

effect is an order of magnitude smaller than the technology levers, it saturates at the statutory rate, and what persistence it has runs through a demand-side feedback rather than through the severance stock itself. Without the on/off decomposition this could have been asserted in either direction. The second case is more technical: the compression of the skill-wage premium, which detailed inspection exposed as an initialisation artefact rather than an AI effect (Appendix B, section B.5). That is not a finding about the labour market, but it is a reminder that a model of this size generates patterns that look like results and are not, and that separating the two is part of what the method costs.

6.4. The role of uncertainty

The adoption decision in this model is forward-looking, but it is forward-looking under certainty. The mean-field NPV firm discounts a known savings stream over a known horizon at a known rate: there is no variance around the expected AI-cost path, no possibility that the technology underdelivers, and no option value of waiting. Risk enters only through the fixed hurdle rate $\eta = 0.15$, a catch-all for managerial caution (Appendix C).

The discount rate deserves particular scrutiny, because it is the one input to that calculation that is both uncertain and consequential in principle: it is anchored on a published cost of capital of about 8.2% per year for the technology sector, a defensible but not a measured choice for the firms in this model (Appendix C). The sensitivity analysis allows the question to be answered directly. Varying the rate by half in each direction, an annual cost of capital of roughly 4% to 12%, moves the final Gini index from about 0.45 to about 0.36 around a baseline of 0.40, and moves the share of tasks automated by less than two and a half percentage points. The direction is the expected one: a firm that discounts the future more heavily values the savings from automation less, so it automates somewhat less and unemployment ends lower. The magnitude, however, is modest against the parameters that matter most. Over a comparable range the planning horizon moves the Gini from 0.27 to 0.51 and the baseline AI productivity from 0.21 to 0.49, three to four times as much. The conclusions of this thesis are therefore not driven by the choice of discount rate; what matters far more is how far ahead firms are assumed to look, and how productive AI is assumed to be on the task, and the latter is precisely the quantity with no measured counterpart. Making the decision forward-looking was the central step beyond the myopic ULC rule, and it proved to be the single strongest moderating force in the model; but the uncertainty that surrounds such an investment in reality was deliberately left out. The direction of that omission can be signed. Under genuine uncertainty and irreversible investment, waiting has value: a risk-aware firm confronted with the same price path would adopt later and less than the model's firms do. The model therefore, if anything, overstates the speed of adoption for a given cost path, and understates the relative grip of institutions whose main contribution is to buy time.

A second layer concerns not what the model leaves out but how confidently its parameters are known. section 6.5 sets out the individual modelling choices; the point that belongs here is the one that a list of separate limitations tends to hide, namely that the parameters are not independent of one another. The baseline is not a set of separately measured values but a set calibrated against each other until the model produced coherent behaviour, with the autonomous demand level, the investment cost, the AI cost scale and the wage curve tuned jointly. That makes the parameter set internally consistent but tightly demarcated: moving away from this point is not a small disruption of a known system, and the interactions between parameters, which are precisely where the risk of a jointly calibrated set sits, are the one thing a one-factor-at-a-time analysis cannot see.

Within that uncertainty, not everything is equally shaky, and it is worth separating the two ends. The most uncertain inputs are the productivity parameters: how productive AI is relative to a human on a given task, and how much more effective a high-skilled worker is than a low-skilled one on non-routine versus routine work. These numbers do not exist in any measured form. They were set to produce plausible behaviour, and they sit at the top of the sensitivity ranking, which is an uncomfortable combination: the outcomes are most sensitive to exactly the quantities that are least anchored. Any statement in this thesis about how much automation occurs inherits that weakness in full.

At the other end sits the contribution that does not depend on the calibration. Replacing an exogenous or myopic adoption rule with an explicit, forward-looking investment decision is a structural change to how the model works, not a parameter setting. It introduces a threshold that must be cleared, a

horizon over which savings are discounted, and an expectation about future wages that responds to the model's own unemployment; those mechanisms operate under any parameterisation, and they will temper adoption relative to a myopic rule wherever they are placed. The direction of that effect is therefore robust in a way the size of it is not, and it is the part of this thesis that should generalise beyond the baseline. The same holds, more weakly, for the qualitative mechanisms: that a wage floor accelerates automation once it binds, that a severance payment delays rather than prevents, and that its effect saturates rather than scaling, are consequences of how the decision rule is built rather than of where the parameters were set.

Read as a whole, then, the honest summary is this: under the conditions constructed here, these are the results. The conditions are documented, the reasoning behind them is traceable, and the mechanisms are more transferable than the magnitudes. What a reader should not take from this thesis is a number.

The deepest uncertainty, however, sits in the model's most powerful variable. Experiment 4 shows that the AI-cost trajectory does not merely scale the outcomes but selects between qualitatively different futures, and that trajectory is precisely the real-world quantity that nobody can forecast today. The baseline assumes a cost that falls towards a floor below the cost of labour. Real-world signals point in both directions: raw inference prices have fallen steeply, but the firm-side cost of adoption, integration, supervision and error correction has not fallen at the same pace. Acemoglu (2025) reaches the aggregate version of the same conclusion, finding the macroeconomic productivity gains of AI over the coming decade to be non-trivial but modest once task-level cost savings are taken seriously rather than assumed, and early deployments are already surfacing in which AI turned out to cost more than the workers it replaced.⁴ A flat or rising effective AI cost is therefore a live possibility rather than a theoretical curiosity, and in the model it is exactly the scenario in which automation nearly halts and the institutions retain their grip the longest.

The consequence for reading this thesis is that the results form a map of conditional futures, not a forecast. Findings that hold across the price paths, the tempering effect of forward-looking adoption, the double edge of the minimum wage, the saturation of the severance payment, deserve the most weight. Findings tied to the cheap-AI regime, above all the dominance of technology over the institutions, are conditional on a cost path that may not materialise. section 7.4 turns this observation into a research agenda.

6.5. Limitations

The interpretation above stands or falls with a number of assumptions and modelling choices. The most important of these are grouped in Appendix H: here the choices that most strongly steer the outcomes and their interpretation are highlighted.

6.5.1. Pattern calibration and baseline dependence

The most fundamental limitation is that the baseline of the model is pattern-calibrated: the parameters in Appendix C are chosen so that the model produces qualitatively plausible and interesting patterns, an expansion gate that closes, plausible levels of inequality and unemployment, rather than being statistically estimated on data. This has two consequences. First, all presented results are conditional on this baseline; they hold in the neighbourhood of this specific parameter point, and all experiments and the OFAT analysis vary around that baseline. Second, the calibration itself can steer the findings: that firms operate close to their profitability frontier and the gate closes regularly is partly a consequence of the chosen demand and cost parameters. In reality firms may have more slack and not sit permanently near the profitability frontier, which would change the strength of the Keynesian brake. The pattern calibration is therefore made explicit in the assumption list (assumption A0.13 in Appendix H).

6.5.2. The inequality effects are an upper bound

Several assumptions cause the model to over- rather than underestimate inequality. There is no reinstatement or reskilling channel: the two skill groups are exogenous and fixed, without schooling or mobility between them (assumptions A0.8 and C1.5), so the compensation of displaced labour through

⁴See for instance J. Green, "AI costs more than the people it replaced", *Forbes*, 2 July 2026, <https://www.forbes.com/sites/jemmagreen/2026/07/02/ai-costs-more-than-the-people-it-replaced/>.

new tasks is absent. Automation is moreover an all-or-nothing substitution per task without augmentation (assumption A0.12), while a large part of current generative AI is deployed complementary to labour; the model therefore captures only the displacement effect and not the productivity or wage gains of human-AI collaboration. Finally there is no government or fiscal redistribution beyond the minimum wage and employment protection (assumption A0.3), so the reported Gini is a gross measure, before taxes and benefits, that overstates net inequality in practice. Together these choices imply that the reported inequality levels should be read as an upper bound and not as a central estimate.

6.5.3. Are jobs just a bundle of tasks?

The entire analysis rests on a premise inherited from the task-based tradition (Acemoglu & Restrepo, 2019; Autor, 2013; Autor et al., 2003): that a job can be represented as a set of separable tasks, each of which can be performed by a human or by a machine, so that automation proceeds task by task and a job is what remains when the automated tasks are removed. That premise is what makes the model tractable, and it deserves to be questioned rather than assumed, because it is not obviously true.

A hairdresser makes the problem concrete. The occupation decomposes into cutting and styling, washing, advising the customer, scheduling appointments, handling payment, ordering stock, cleaning the salon and maintaining the relationship that makes the customer return. A task-level exposure assessment would find a substantial share of that list automatable: scheduling, payment and stock ordering are routine administrative work, and cleaning is being mechanised independently. On a naive reading of the task share, a large fraction of the job is at risk. Nothing of the sort follows. The task that defines the occupation, cutting hair to a specific head in real time, is manual, non-routine and personal, and automating everything around it does not remove the job; it makes the hairdresser more productive at the part that cannot be removed, and reduces the need for supporting staff rather than for hairdressers.

Three distinct problems are visible in that example, and all three apply to the model used here. First, tasks carry unequal weight within an occupation. A job is not a list of interchangeable activities but a structure with a core and a periphery, and a 40 per cent automatable task share tells us almost nothing about a 40 per cent risk of displacement unless we know which 40 per cent. Second, and following directly from this, automating peripheral tasks is often complementary to the core: it is the augmentation channel arriving through the back door, and a purely task-counting model registers it as displacement. Third, tasks are not independent or fixed in number. Automating one task changes what the neighbouring tasks are: some disappear because they only existed to support the automated one, some merge, and new ones appear (supervising the tool, checking its output, handling the cases it cannot). The bundle is reshaped, not simply shortened. To this can be added a residual that task inventories systematically miss, the coordination, judgement, accountability and simple presence that hold a job together and are rarely written down as tasks at all.

The model inherits these problems in a specific and, in one respect, uncomfortable form. The Leontief structure (assumption C1.1) treats every task as a potential bottleneck, which means all tasks within a firm are equally essential. That is the exact opposite of the hairdresser case, where one task carries the occupation and the rest support it. The assumption is conservative in one direction, since a firm cannot simply drop a task it finds inconvenient, and too coarse in the other, since it weighs the defining task and the administrative one identically. Combined with the absence of task creation (assumption A0.8), this means the model can capture how far automation reaches across a fixed task set, but not how the set itself is reorganised in response.

Why keep the approach nonetheless. The task level is where automation decisions are actually taken, and it is the only level at which usable data exists: the exposure scores this thesis draws on (Gmyrek et al., 2023) are constructed task by task, precisely because occupation-level judgements proved unreliable. The main alternative, classifying whole occupations as automatable or not, is coarser and occupations are treated as all-or-nothing when in practice only parts of them were exposed. Moving to the task level was a correction of that error, not a further simplification.

The honest position is therefore narrower than either extreme. The task decomposition is the best available operationalisation and remains defensible for what this thesis actually does, which is to compare scenarios against one another on an identical task structure: because every configuration inherits the same simplification, it largely cancels in the paired comparisons that carry the results. What the

approach cannot support is the step from a task share to a statement about which occupations disappear, and no claim of that kind is made here. Jobs, in short, are more than bundles of tasks; but for the question of how institutions and prices shift the boundary of what gets automated, treating them as bundles is a simplification whose cost falls mainly on the levels rather than on the comparisons.

6.5.4. The skill-task split is too coarse

One of the model's most consequential simplifications is structural, and it sits at the heart of the distributional results. The model ties the routine/non-routine task distinction to the skill distinction along a single axis: routine tasks are skill-neutral and only non-routine tasks favour the high-skilled (assumption C1.8), so the two skill groups sort cleanly into routine and non-routine work. Reality is more tangled, and the dividing line is too coarse. There are high-skilled routine tasks, much routine cognitive professional work, from standard radiological reads to bookkeeping and document review, that generative AI exposes precisely because they are routine, and there are low-skilled non-routine tasks, cleaning, care and many manual service jobs, that resist automation precisely because they are non-routine and manual. By tying skill to the non-routine dimension alone, the model can misplace where automation lands across the skill groups, plausibly understating the exposure of high-skilled routine work and overstating that of low-skilled non-routine work. The limited cross-skill matching (assumption C2.4) softens but does not remove this, since the underlying productivity advantage runs along the non-routine dimension alone. Because the skill-wage premium and the incidence of unemployment across the two groups are exactly the distributional outcomes this thesis reports, they should be read with this coarseness in mind.

The premium carries a second caveat that follows from the same simplification. It falls from roughly 1.8 to about 1.3 in every scenario, including those in which almost nothing is automated and those in which automation is switched off entirely, so the drift is not an AI effect but an adjustment of the model's own starting point: the warm start allocates workers at near-equal wages, and the wages that allocation implies then pull apart far enough to make most high-skill placements too expensive, after which ordinary turnover repairs the mismatch (Appendix B, section B.5). The level of the premium therefore carries the initialisation with it and should not be read as a statement about Dutch wage differentials. The comparisons are unaffected, since every configuration starts from the same point and the common adjustment cancels in the paired tests of chapter 5; what those tests measure is how much more or less the premium compresses under a different AI-cost path, comparative advantage or institution.

6.5.5. Structural and functional simplifications

A number of further simplifications colour the outcomes. The Leontief technology (assumption C1.1) makes every task a potential bottleneck and ties automation, employment and output tightly together; almost all dynamics run through this mechanism, while firms in practice can partly skip or combine tasks. The AI learning curve is exogenous and decoupled from adoption (assumptions C5.2 and C5.4): AI becomes cheaper inevitably, driven by time rather than by cumulative adoption, which strongly influences the timing of automation and makes AI structurally cheaper than labour in the long run. The model has one homogeneous good and one sector (assumption A0.2), while the Dutch dual story is strongly sectoral; the outcomes hold for an average economy, not for specific sectors such as hospitality or industry. And as discussed in subsection 6.2.5, the contract structure (assumption C7.4, new workers always start as flex) diverges from reality and damps the institutional asymmetry. Relatedly, CLA wage rigidity is captured only as a symmetric partial-adjustment speed λ and not as downward nominal rigidity (assumption C3.2 and section 3.8): The wage institution that best defines the polder model is treated as a delayed reaction. This likely understates how firmly CLAs hold nominal wages up during downturns, and therefore how strongly that rigidity could, through induced technical change, accelerate substitution toward AI.

A further simplification of the worker lifecycle colours the contract dynamics. Workers leave a firm only through the exogenous separation process (assumption C2.2); there is no retirement or demographic turnover (assumption A0.6). Two consequences follow for the late-horizon employment-protection effect analysed in subsection 6.2.5. Its size depends on how close firms sit to the profitability frontier, so it inherits the limitation of the pattern calibration (assumption A0.13) rather than of the lifecycle. And it should be read as an upper bound, because the model omits precisely the channels that would damp a

runaway demand spiral in reality: fiscal stabilisers (A0.3), reinstatement and reskilling (A0.8 and C1.5) and human-AI augmentation (A0.12).

6.5.6. Limitations of the sensitivity analysis

The sensitivity analysis in Appendix G is local and one-factor-at-a-time: each parameter is varied separately around the baseline. Interaction effects between parameters therefore remain invisible, and the results hold only for the immediate neighbourhood of the baseline. A global variance-based test, such as a Sobol analysis, would expose interaction and non-linear effects across the full parameter grid, but fell outside the scope of this research because of the required computational effort. In addition the scale parameters (the number of agents) mix a model-technical scale effect with an economic effect, so their high place in the ranking should be read with that caveat. The analysis is thus an exploratory check of which assumptions dominate the outcomes, in line with the exploratory setup, and not a full uncertainty quantification.

6.5.7. What these limitations mean for reading the results

Taken together, the limitations point in one direction, and it is worth stating plainly what that means for anyone reading this thesis as a statement about the Dutch labour market rather than about a simulation. The absolute levels should be discarded entirely. An economy in which more than a third of the workforce is out of work is not a forecast for the Netherlands; it is what a model produces when labour can be displaced but never reinstated, when nobody retrains, when there is no unemployment benefit or tax system redistributing anything, and when the entire economy consists of one sector making one good. Each of those omissions pushes the reported inequality and unemployment upward, and none of them pushes back. The reported levels are an upper bound on a world that does not exist.

What survives the translation is thinner but more useful: the ordering of the forces, the direction in which each lever pushes, and the conditions under which a lever works at all. That the price of AI relative to labour dominates the contract institutions is a statement about relative magnitudes that does not depend on the absolute levels being right. That a wage floor accelerates automation once it binds, that a severance payment delays rather than prevents, and that the effect of a severance payment saturates rather than scaling with its size, are all mechanisms rather than magnitudes, and mechanisms travel better than numbers. And precisely because the model excludes augmentation and reinstatement, the one thing it cannot tell us is the thing that may matter most in practice: whether generative AI, in the Netherlands, will mostly replace workers or mostly make them more productive. Everything in this thesis is conditional on the first of those two worlds.

With the interpretation, the reflection on the approach, the role of uncertainty and the limitations in place, chapter 7 answers the research questions and translates the findings to the Dutch labour market and its policy debate.

7

Conclusions

This chapter concludes the thesis. It first answers the research questions (section 7.1), staying close to what the model demonstrably shows. It then deliberately steps away from the model: section 7.2 translates the findings to the real Dutch labour market and its policy debate, and section 7.3 draws the lessons for the modelling literature. section 7.4 sets out the research agenda, ordered by importance rather than enumerated, and section 7.5 closes.

7.1. Answering the research questions

Sub-question 1 (which mechanisms of task-based AI automation are decisive in the Dutch context) was answered in chapter 2, and sub-question 2 (how the model of Upreti and Sridhar was methodologically adapted and calibrated) across the conceptual model (chapter 3) and its implementation (Appendix A), with the explicit answer in section 3.14. The results of chapter 5, interpreted in chapter 6, answer the remaining sub-questions.

Sub-question 3 asked how endogenizing the investment decision affects the simulated causal relationship between technology use and both wage inequality and unemployment compared to exogenous scenarios, and which feedback loops come into play between wage levels, labour scarcity and adoption speed. The endogenous, forward-looking NPV adoption clearly tempers automation, and with it both inequality and unemployment, relative to the short-sighted ULC approach (Experiment 1). Three loops do the work. The mean-field expectation makes firms anticipate that economy-wide automation raises unemployment, which presses wages down through the wage curve and therefore lowers the savings expected from automating further. The Keynesian demand feedback erodes profitability, because displacement lowers the wage bill and with it demand and price. The profitability gate then halts proactive expansion once average cost approaches price. Together these loops give the relationship its characteristic shape: inequality and unemployment rise and then settle at a plateau, because automation continues only until the running cost of AI and the cost of the remaining labour come to break even (subsection 6.2.2). The simulated relationship is therefore a one-off level shift rather than a continuing deterioration, which is the clearest difference from an exogenous adoption scenario, where nothing in the model restrains the process.

How much of this tempering the loops themselves account for is settled by the decomposition of the headline comparison (Experiments 2 and 3): the large majority is the forward-looking investment logic, and only a small part the institutional layer. The institutions are taken up under sub-question 4, together with their effectiveness as instruments.

Sub-question 4 asked about the expected effectiveness of different policy instruments in mitigating wage inequality and unemployment, and how the Dutch labour-market institutions reinforce or weaken the feedback loops set out above. Taking the loops first: the institutions bend them in direction but not in magnitude, and one works against them. Employment protection reinforces the braking loops, because the severance payment raises the upfront cost on occupied tasks. The minimum wage weakens them, because raising the cost of low-skilled labour makes automation profitable for more tasks and so feeds

the displacement channel rather than damping it. Effectiveness as instruments depends on which of the two outcomes is targeted, and none of the instruments examined improves both. A high minimum wage is the only one with a clear effect on the wage structure, and it is double-edged: it compresses the premium for those who keep their jobs while pushing unemployment and the Gini up (Experiment 6), so it mitigates inequality among workers at the cost of inequality between the employed and the displaced. Employment protection mitigates in the intended direction but weakly, and mainly through timing: it slows the first wave of displacement and, through the demand loop, carries that early difference forward into a modest long-run gap. Raising the transition payment above the statutory rate adds nothing and the contract mix barely matters (Experiments 3, 7 and 8, with the mechanism in subsection 6.2.5). Its value is the transition time it buys rather than the outcome it prevents, and that value is only realised if the time is used. Because the outcomes are governed above all by the relative price of AI, the instrument with the most leverage is not a labour-market institution at all but price-oriented policy, such as a tax on AI use: it acts on the variable that sets where the economy ends up, rather than on the speed at which it arrives there.

The main question, to what extent endogenous AI adoption within the Dutch labour market influences wage inequality and unemployment, and which policy instruments can mitigate these effects, can thus be answered as follows. Endogenous, forward-looking AI adoption still raises inequality and unemployment over the horizon in almost all scenarios, through the displacement of low-skilled routine tasks, but it does so markedly less than the short-sighted, myopic adoption rule of Upreti and Sridhar (2024): making the adoption decision forward-looking is itself the single strongest moderating force, because the anticipated wage decline, the Keynesian demand feedback and the profitability gate together hold automation back (Experiment 1). The magnitude of the remaining increase is governed above all by the relative price of AI, while the direction of the skill-premium effect depends on the comparative advantage of AI across tasks. The classical Dutch labour-market institutions offer only limited grip on this process: employment protection buys transition time without reshaping the outcome the way the price of AI does. The minimum wage is the one institution with a clear effect, but a double-edged one: a higher floor compresses the wage premium of those who keep their jobs, yet by raising the cost of low-skilled labour it makes automation profitable for more tasks and therefore pushes unemployment and the Gini up rather than down. The most promising lever is therefore not a classical labour-market institution but price-oriented policy, such as a tax on AI use, precisely because the relative price of AI dominates the outcomes; the institutions mainly shift the timing and, a little, the level. These conclusions are exploratory and conditional on the calibrated baseline.

7.2. Implications for the Dutch labour market

This section translates the findings to the labour market the model was built to say something about, and it marks clearly where demonstration ends and informed judgement begins. The model is displacement-only, gross-income, one-sector and pattern-calibrated, so what follows should be read as direction and mechanism, never as magnitude.

7.2.1. How displacement would surface

The model's firms rarely dismiss protected workers in order to automate; they wait for ordinary turnover and automate the vacated task, because that avoids the severance payment altogether. If Dutch firms behave in anything like this way, AI-driven displacement will not, in the first instance, look like waves of dismissals of permanent staff. It is more likely to surface as vacancies that are quietly not refilled, as hiring freezes for starters, and as flexible contracts that are not renewed, with the adjustment burden falling once again on the flexible shell and on labour-market entrants (CBS, 2025a). The practical consequence is a monitoring one. The indicators that would show this early are the total number of open vacancies and how long they stay open, the share of vacancies that are closed without being filled, the unemployment and hiring rates among starters and young workers, and the renewal rates of temporary contracts. Dismissal statistics, the figures a labour-market debate instinctively reaches for, would be the last place it appears: by the time displacement is visible there, the process would already be well underway.

7.2.2. Augmentation is the largest open question

The model captures only substitution, and its inequality outcomes are upper bounds for that reason alone. The exposure data on which this thesis draws point the other way: Gmyrek et al. (2023) conclude that the overwhelming effect of generative AI will be to augment occupations rather than to automate them, with the largest impact in high-income countries because of their higher share of clerical employment. Adding that channel would change the results substantially: unemployment would end up lower than reported here, though plausibly still elevated, and the effect would differ strongly across sectors. It would also shift who is affected, since clerical work is a major source of female employment and the incidence is therefore strongly gendered, a dimension a two-skill-group model cannot see at all.

The decisive question is not whether AI makes individual workers more productive, it plainly does for a great deal of computer-based office work, but what firms do with that productivity. Two responses to the same technology point in opposite directions. A firm can keep its staff and let everyone become more efficient, in which case augmentation raises output without displacing anyone. Or it can conclude that two heavy AI users now cover what previously required a whole team, in which case augmentation is substitution wearing a friendlier name. The technology does not decide between these; the firm does. This is why the same tool can be reported as job-creating in one organisation and job-destroying in another, and it is a source of uncertainty about the real-world counterpart of this model.

Who captures the productivity gain is a related question with, in my assessment, a less balanced answer. The most likely beneficiary is the employer. Firms carry the cost of the technology, but if a smaller number of workers can do the work of a larger number, ordinary supply and demand does not obviously push up the wages of those who remain, even though their individual output has risen: the workers displaced from the same occupation are the competition for the jobs that are left. Programming is the clearest current example. If AI substantially reduces the number of programmers a firm needs, the ones who stay are unlikely to be paid markedly more for having become more productive, because there are now more programmers than positions. Some of the gain will reach consumers through prices, but the distributional default is that it accumulates capital rather than labour. That is speculation rather than a model result, and it is offered as such; a serious answer would require modelling the augmentation channel and the bargaining over its surplus explicitly, which is why it heads the research agenda in section 7.4.

Taking a position on this is unavoidable, since the model cannot, and mine is the following. Augmentation will become the normal case, and increasingly so: for most office and knowledge work the practical result of the current wave is a worker who does more, not a worker who is gone. Outright replacement will happen at the edges, and is being experimented with there now, in narrowly scoped, high-volume, well-specified work where the output is easy to check. If the technology keeps developing at anything like its recent pace, that edge will move inward and replacement will become steadily more common, but it will move unevenly: the jobs that will hold out longest are the physical ones, where the constraint is not cognition but hands, presence and an unstructured environment. That produces an uncomfortable asymmetry with the model's own structure. This thesis, following the routine/non-routine tradition, points automation at low-skilled routine work; my own expectation is that it arrives first and hardest in cognitive work well above the wage floor, and that manual work, including much low-paid manual work, is the more defensible position for the time being. The task-based reflection of subsection 6.5.3 points in the same direction, and it is one of the clearest respects in which the simulated world and the expected one diverge.

One caution belongs with this. Firms may currently be more eager to adopt AI than the realised productivity gains justify. Acemoglu (2025) makes the quantitative version of this point: working through a task-based model, and using existing estimates of task-level exposure and cost savings, he finds the aggregate productivity and GDP effects of AI over the next decade to be non-trivial but modest, and far smaller than the transformation the public debate assumes. Generative AI is nonetheless being marketed and discussed as something close to a transformation of all knowledge work, and adoption decisions taken in that atmosphere are not the coolly discounted investment calculations this model assumes. If part of the current adoption wave is enthusiasm rather than calculations, then displacement could run ahead of the genuine cost advantage for a time, and some of it could later be reversed, signifying a negative hurdle rate. The model, whose firms are boundedly rational but never fashionable,

cannot represent this at all.

7.2.3. What employment protection does and does not buy

In the model the severance payment slows automation early and, through the demand feedback, leaves a modest lasting trace; raising it above the statutory level adds nothing. The translation is twofold. On the one hand, the Dutch employment-protection is worth more as an automation policy than it seems, not because it prevents displacement but because it buys time precisely in the phase when the technology is new, and time is the scarce resource for retraining, task redesign and demographic absorption. On the other hand, it would be a mistake to respond to AI by making dismissal more expensive. The saturation result suggests the brake is already as strong as it usefully gets at the statutory rate, and in reality a much higher severance would carry costs the model does not price, from reduced permanent hiring to a still larger flexible shell. The institution's value lies in the time it buys, and that value is only realised if the time is used, which the protection itself does not arrange.

7.2.4. The minimum wage: not binding yet

The model's warning is specific: once the wage floor rises above what low-skilled labour earns in the market, every further increase strengthens the business case for automating exactly the tasks those workers hold. Whether that warning is currently live in the Netherlands is a separate question, and the answer is probably not. The work generative AI is best at, computer-based office and administrative work, is generally paid well above the statutory minimum, so the floor is not what makes automation attractive there. The mechanism is conditional rather than absent, and it becomes relevant in two ways. It applies today to specific occupations where a wage close to the floor coincides with high substitutability, a combination that is occupation-specific and therefore invisible in a one-sector model such as this one. And it becomes broadly relevant if the minimum wage continues to rise, or if AI capability extends further down the wage distribution, or both. The policy implication is not to hold the floor down; it is that in those specific segments, wage-floor policy and automation policy are the same policy and should be weighed together rather than decided in separate files.

7.2.5. Taxing AI: the strongest instrument and the hardest to build

The model points at price-oriented policy because the relative price of AI dominates its outcomes. That conclusion does not rest on the model alone, and it arrives in a debate that is already well developed. Acemoglu, Manera and Restrepo (2020) argue that the US tax system is biased against labour and in favour of capital, and that this bias has itself promoted inefficiently high levels of automation; on their estimates, moving to optimal taxation of the two factors would raise employment by around four per cent. Their point is not that automation should be discouraged but that the existing fiscal treatment silently subsidises it, which reframes an AI tax as the removal of a distortion rather than the introduction of one. Guerreiro, Rebelo and Teles (2022) reach a complementary conclusion from an optimal-taxation model with endogenous skill choice: a sustained fall in automation costs produces a large rise in inequality, and it is optimal to tax robots while the current generation of routine workers is still in the labour force, with the optimal tax falling to zero once those workers have retired. Zhang (2019) shows that whether such a tax narrows or widens the skilled–unskilled wage gap depends on the substitution elasticities in the robot-producing sector, so the distributional case is conditional rather than automatic.

Three things follow for this thesis. First, the direction the model points in is the direction this literature already points in, which is reassuring for a mechanism that was not built to produce it. Second, and more usefully, the Guerreiro et al. result gives the present findings a sharper interpretation than the model alone supports: if the optimal robot tax is transitional, aimed at the cohort that cannot retrain, then it belongs in the same category as the employment protection examined here, an instrument that buys time for a specific generation rather than a permanent correction. That the model finds Dutch dismissal protection doing exactly that, buying time rather than changing the destination, suggests these two instruments are substitutes for the same job, and that the price instrument is simply the stronger of the two. Third, there is a fiscal argument that runs alongside the labour-market one: if AI displaces labour at scale, public finances are hit from both sides at once, since social-security expenditure rises while the wage-based revenue that funds it falls, a displaced worker paying little income tax and an algorithm paying none.

The goal of such a tax should not be to stop AI adoption, and framing it that way invites the obvious

objection that it would sacrifice productivity growth. The goal is narrower and more defensible: to keep hiring a human, or a human working with AI, competitive against replacing that human outright. It is a correction of a relative price, not a brake on a technology.

The design problems are formidable and should not be minimised. AI use is genuinely hard to measure: there is no natural unit, and a tax authority cannot reasonably audit how much inference a firm bought and for what. There is a real efficiency trade-off, because if the productivity gains are large, taxing them carries a cost that must be set against the labour-market benefit, and where the balance lies is an empirical question nobody has yet answered. And a tax imposed by one country on a globally supplied technology invites relocation and puts that country at a disadvantage; a national AI tax without international coordination risks harming the country that introduces it more than the practice it targets. None of this makes the instrument unworkable, but it does mean the gap between the model's clean price lever and a usable tax is wide, and that closing it is itself a research agenda. What can be said is that of everything examined here, this is the instrument that would actually change the outcome.

7.3. Academic implications

Three lessons follow for research that models technological change.

The first concerns how adoption itself is represented. In most existing models a technology spreads either because the modeller assumes it does, or because firms draw it in a research lottery. Letting firms weigh the cost of the technology against the cost of the workers it would replace, over a horizon they care about, changes the outcome substantially: adoption is slower and its distributional consequences are milder than the same model produces under a short-sighted rule. That difference turned out to be larger than everything the Dutch institutions do combined. If a model is used to say anything about the pace of technological change, how the adoption decision is specified is not a detail of implementation but one of the most important modelling choices.

The second concerns institutions. Adding dismissal protection, a wage floor, contract types and collective wage-setting to a model of automation is worth doing, but this research suggests a boundary condition that is easy to miss: an institution only becomes visible in the outcomes when the cost it imposes is large relative to the cost difference driving the technology. Dutch employment protection is not large in that comparison, and no reasonable adjustment of it would make it so. Institutional detail, in other words, does not automatically buy institutional relevance, and a model that adds it without checking the relative magnitudes may be reporting a mechanism it has effectively switched off.

The third is methodological. Simulating individual firms and workers is expensive, and this thesis found that for most of its experiments the extra detail did not change the shape of the answer, only its height. Where it did earn its cost was in three specific places: where a single assumption switched the economy between qualitatively different futures rather than scaling one outcome, where holding every chance event fixed made it possible to isolate the contribution of one institution at a time, and where the spread across repeated runs revealed that half of a result rested on a handful of extreme cases. That is a usable rule for future work: heterogeneity pays where outcomes are discontinuous, where the counterfactual matters, and where the distribution carries information the average hides.

7.4. Recommendations for further research

From the findings and limitations an agenda for further research follows. Three priorities are set out in detail, because they carry substantially more weight than the remainder; three further extensions follow more briefly.

The single most consequential extension is an **augmentation and reinstatement channel**. Everything this model reports is conditional on automation being all-or-nothing displacement without new task creation, schooling or mobility between skill groups; it is for exactly this reason that the inequality outcomes must be read as upper bounds. Adding human-AI collaboration within tasks, and a reinstatement channel through which displaced labour returns via new tasks or reskilling, would turn the upper bound into something closer to a central estimate, would allow the model to capture the productivity and wage gains of complementary AI, and could plausibly change the sign of the wage effects for the groups most exposed. No other single change would do as much for the real-world validity of the conclusions.

The second priority is a systematic treatment of **uncertainty**, in two complementary directions. Around the model, the local OFAT analysis should be widened to a global, variance-based exploration of the parameter space (a Sobol or comparable design), which would expose interactions and non-linearities and test how far the conclusions travel from the pattern-calibrated baseline. Within the scenarios, the policy instruments should be re-run against the technology futures the model already isolates: crossing the institutional levers (employment protection, the minimum wage, the severance rate) factorially with the AI-cost trajectories of Experiment 4, including the flat and rising paths, and with the comparative advantage of Experiment 5. The price of AI is the scenario-determining factor of this thesis and, as argued in section 6.4, the single most uncertain quantity in the real world; whether the headline ordering “price of AI ahead of the institutions” survives a future in which AI does not become structurally cheaper is exactly the question this design would answer.

The third priority is the **design of a price-based instrument**, because it is the only lever examined here that would move the outcome rather than the timing, and it is the one furthest from being usable. The economic case is established: the existing tax treatment of labour relative to capital already tilts firms towards automation (Acemoglu et al., 2020), and optimal-taxation analysis suggests a tax while the exposed cohort is still working (Guerreiro et al., 2022). What is missing is everything between that result and a policy. Three questions in particular need answering, and none of them is addressed by the existing literature or by this thesis. What is the taxable unit, given that AI use has no natural measure and a firm’s spending on it is neither observable nor stable? How can an instrument distinguish substitution from augmentation, so that it does not tax the channel that raises productivity at the same rate as the one that displaces workers, a distinction this thesis identifies as decisive but cannot itself model? And what happens under national introduction, given that the technology is supplied globally and its use is easy to relocate? Answering these is a research programme in public economics rather than in simulation, but the labour-market case for undertaking it is what this thesis contributes.

Beyond these three priorities, three further extensions would each remove a specific limitation. Sector heterogeneity, splitting the one-sector economy into sectors with their own task profiles and contract mixes, would do justice to the strongly sectoral Dutch dual structure and would allow the occupation-specific conditions identified in section 7.2 to become visible. An endogenous AI price, coupling the cost curve to cumulative adoption, would replace the exogenous decay that now partly builds the headline result into the model. And an alternative contract structure, with direct permanent hiring and transition payments for long-running flexible contracts, would test whether the institutional asymmetry becomes more visible when the modelled contract dynamics stop damping it.

7.5. Conclusion

This research extended the task-based model of Upreti and Sridhar (2024) with an explicit, forward-looking investment decision on AI adoption, with the Dutch labour-market institutions and with a Keynesian demand feedback. The central lesson is that making that decision endogenous changes the outcomes meaningfully: automation and inequality are tempered relative to a short-sighted adoption logic, driven by the expected wage decline, the demand feedback and the profitability gate. A decomposition of the headline comparison shows that the large majority of the gap between the myopic rule and the calibrated baseline is the forward-looking investment logic itself, and only a small part the Dutch institutions. At the same time the relative price of AI proves by far the most important force behind the outcomes, while the typically Dutch contract institutions hold only limited grip on the process: they slow the journey and shift the destination a little, rather than reshaping it.

These conclusions must be read as exploratory. The model is an instrument to explore directions, feedbacks and sensitivities around a pattern-calibrated baseline, not a forecast of future Dutch wage inequality, and its most decisive variable, the price path of AI, is precisely the one the real world has not yet settled. Its value lies in making visible which mechanisms matter and which uncertainties dominate: the price dynamics of AI above all, the institutions far less, and the difference between a displacing and an augmenting technology most of all still to be determined, in reality as much as in the model.

If this thesis supports one recommendation above the others, it is to be prepared. The model shows an economy in which the decisive variable, the relative price of AI, can move fast and can select between futures that are not merely worse or better but structurally different; and it shows institutions that mainly

buy time rather than change destinations. Policy, by contrast, is typically made after a problem has become measurable, and the monitoring indicators identified in section 7.2 suggest that in this case the measurement will lag the process. That asymmetry, a technology that can move in years and a policy apparatus that reacts in electoral cycles, is the practical risk this research points at. It does not require believing that artificial general intelligence is imminent to take that seriously; it requires only accepting that the price path is uncertain and that the instruments which work, according to this model, are the ones that act on that price and on the time available to adjust. Both take longer to build than to need.

Statement on the use of AI

For the writing of this thesis I used AI tools. In this statement I explain which tools I used, for what purposes, and how I kept control and responsibility for the final result.

Used tools. During the process, from proposal to completion, I used two tools: ChatGPT (OpenAI) and Claude (Anthropic). In the initial phase, during the proposal and literature research, I mainly used ChatGPT. Once the modelling process started, Claude was used more because of its model capabilities.

Purposes for which I used the tools. The use of AI was supportive; it was not used to generate complete thesis content. Concretely, I used AI to rewrite my own texts, make them more fluent, and change my Dutch formulations into correct academic English; to help develop and refine my ideas; to review the paper in the later stages for inconsistencies, grammar, flow, logic, layout and clarity; to check my Python code, make it more readable, add clearer comments and debug where necessary; to write parts of the code on my instructions, mainly the analysis, plotting and sensitivity scripts that process and visualise the model output; to create the Streamlit dashboard and plot results; to create the cover page image; to convert originally written Word text into LaTeX, help with the layout in LaTeX, create list of abbreviations, create list of tables & figures and clickable references in LaTeX.

Method and control. In all cases the direction, specification and verification were mine. For the writing, my own text was the core and AI acted as reviewer, sparring partner and editor. For the code, I wrote and own the core model, while for the surrounding analysis I specified what each script had to compute and let AI write larger portions of it, after which I checked the logic and verified the results before using them. I gave it proposals and only implemented changes after I had explicitly reviewed and approved them. In that way, I maintained control over the content and message in every review round. Because I read and processed everything myself, no fabricated or unverified references ended up in this paper.

Responsibility. The research questions, the theoretical decisions, the conceptual and computational model, the analysis of the results, the discussion and the conclusion are my own work. AI helped me write and program more efficiently and in better English, but the intellectual decisions and final responsibility for this work are completely mine.

Bjarne Timmer
The Hague, June 2026


References

- Acemoglu, D. (2002). Directed Technical Change. *The Review of Economic Studies*, 69(4), 781–809. <https://doi.org/10.1111/1467-937X.00226>
- Acemoglu, D. (2025). The simple macroeconomics of AI. *Economic Policy*, 40(121), 13–58. <https://doi.org/10.1093/epolic/eiae042>
- Acemoglu, D., & Autor, D. H. (2011). Skills, Tasks and Technologies: Implications for Employment and Earnings. In O. Ashenfelter & D. Card (Eds.), *Handbook of Labor Economics* (pp. 1043–1171, Vol. 4B). Elsevier. [https://doi.org/10.1016/S0169-7218\(11\)02410-5](https://doi.org/10.1016/S0169-7218(11)02410-5)
- Acemoglu, D., Manera, A., & Restrepo, P. (2020). Does the US Tax Code Favor Automation? *Brookings Papers on Economic Activity*, 2020(1), 231–300.
- Acemoglu, D., & Restrepo, P. (2018). The Race between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment. *American Economic Review*, 108(6), 1488–1542. <https://doi.org/10.1257/aer.20160696>
- Acemoglu, D., & Restrepo, P. (2019). Automation and New Tasks: How Technology Displaces and Reinstates Labor. *Journal of Economic Perspectives*, 33(2), 3–30. <https://doi.org/10.1257/jep.33.2.3>
- Acemoglu, D., & Restrepo, P. (2022). Tasks, Automation, and the Rise in U.S. Wage Inequality. *Econometrica*, 90(5), 1973–2016. <https://doi.org/10.3982/ECTA19815>
- Autor, D. H. (2013). The “task approach” to labor markets: An overview. *Journal for Labour Market Research*, 46(3), 185–199. <https://doi.org/10.1007/s12651-013-0128-z>
- Autor, D. H., Levy, F., & Murnane, R. J. (2003). The Skill Content Of Recent Technological Change: An Empirical Exploration. *The Quarterly Journal of Economics*, 118(4), 1279–1333. <https://doi.org/10.1162/003355303322552801>
- AWVN. (2025). Eindevaluatie cao-jaar 2024: Loon blijft dominant. Retrieved June 1, 2026, from <https://www.awvn.nl/cao/publicaties/evaluatie-cao-jaar-2024/>
- Balsmeier, B., & Woerter, M. (2019). Is this time different? How digitalization influences job creation and destruction. *Research Policy*, 48(8), 103765. <https://doi.org/10.1016/j.respol.2019.03.010>
- Blanchflower, D. G., & Oswald, A. J. (1995). An Introduction to the Wage Curve. *Journal of Economic Perspectives*, 9(3), 153–167. <https://doi.org/10.1257/jep.9.3.153>
- Brynjolfsson, E., Mitchell, T., & Rock, D. (2018). What Can Machines Learn, and What Does It Mean for Occupations and the Economy? *AEA Papers and Proceedings*, 108, 43–47. <https://doi.org/10.1257/pandp.20181019>
- Caloia, F. G., Parlevliet, J., & Mastrogiacomo, M. (2023). Staggered wages, unanticipated shocks and firms’ adjustments. *Journal of Macroeconomics*, 76, 103521. <https://doi.org/10.1016/j.jmacro.2023.103521>
- CBS. (2024a). Flexwerk in Nederland en de EU. Retrieved April 30, 2026, from <https://www.cbs.nl/nl-nl/dossier/dossier-flexwerk/flexwerk-in-nederland-en-de-eu>
- CBS. (2024b). Het aanbod van arbeid - De arbeidsmarkt in cijfers 2024 | CBS. Retrieved June 2, 2026, from <https://longreads.cbs.nl/de-arbeidsmarktincijfers-2024/het-aanbod-van-arbeid>
- CBS. (2024c). Hoeveel flexwerkers zijn er? Retrieved April 30, 2026, from <https://www.cbs.nl/nl-nl/faq/flexwerk/hoeveel-flexwerkers-zijn-er->
- CBS. (2025a). *Flexwerkers en de coronacrisis*. Centraal Bureau voor de Statistiek. Retrieved August 13, 2026, from <https://www.cbs.nl/nl-nl/faq/corona/bevolkingsgroepen/flexwerkers-en-de-coronacrisis>
- CBS. (2025b). Wat voor werk doen flexwerkers? Retrieved April 30, 2026, from <https://www.cbs.nl/nl-nl/faq/flexwerk/wat-voor-werk-doen-flexwerkers->
- CBS. (2025c, August). Fewer workers changing jobs. Retrieved June 2, 2026, from <https://www.cbs.nl/en-gb/news/2025/33/fewer-workers-changing-jobs>

- CBS. (2025d, November). Kenmerken van bedrijven die AI-technologie gebruiken. Retrieved December 4, 2025, from <https://www.cbs.nl/nl-nl/maatwerk/2025/44/kenmerken-van-bedrijven-die-ai-technologie-gebruiken>
- CBS. (2025e, December). 2. Kenmerken van de werknemer. Retrieved June 1, 2026, from <https://www.cbs.nl/nl-nl/longread/aanvullende-statistische-diensten/2025/monitor-loonverschillen-mannen-en-vrouwen-2024/2-kenmerken-van-de-werknemer>
- CBS. (2026, April). Arbeidsdeelname; onderwijsniveau. Retrieved June 1, 2026, from <https://www.cbs.nl/nl-nl/cijfers/detail/85266NED>
- Corden, W. M. (1978). Keynes and the others: Wage and price rigidities in macro-economic models. *Oxford Economic Papers*, 30(2), 159–180. <https://www.jstor.org/stable/2662886>
- Cottier, B., Rahman, R., Fattorini, L., Maslej, N., & Owen, D. (2024, May). How much does it cost to train frontier AI models? Retrieved May 4, 2026, from <https://epoch.ai/blog/how-much-does-it-cost-to-train-frontier-ai-models>
- Cyert, R. M., & March, J. G. (1963). *A Behavioral Theory of the Firm*. Prentice-Hall.
- Dawid, H., Gemkow, S., Harting, P., van der Hoog, S., & Neugart, M. (2014, January). Agent-Based Macroeconomic Modeling and Policy Analysis: The Eurace@Unibi Model. Retrieved April 30, 2026, from <https://papers.ssrn.com/abstract=2384391>
- Dawid, H., & Neugart, M. (2023). Effects of technological change and automation on industry structure and (wage-)inequality: Insights from a dynamic task-based model. *Journal of Evolutionary Economics*, 33(1), 35–63. <https://doi.org/10.1007/s00191-022-00803-5>
- De Beer, P. T., & De Sauvage Nolting, R. (2024). *Balanceren op de cao-helling: Oorzaken van de daling van de cao-dekkingsgraad en wat eraan te doen is* (Research Report No. 27). De Burcht, Wetenschappelijk Bureau voor de Vakbeweging. Amsterdam. Retrieved April 30, 2026, from https://wp.kennisbanksocialeinnovatie.nl/wp-content/uploads/2024/07/Publicatie_27_Balanceren_op_de_cao-helling-1.pdf
- De Nederlandsche Bank. (2024, October). Met slimmer werken personeelstekorten te lijf. Retrieved May 4, 2026, from <https://www.dnb.nl/algemeen-nieuws/achtergrond-2024/met-slimmer-werken-personeelstekorten-te-lijf/>
- Dixit, A. K., & Pindyck, R. S. (1994). *Investment under Uncertainty*. Princeton University Press. <https://doi.org/10.2307/j.ctt7sncv>
- Dosi, G., Fagiolo, G., & Roventini, A. (2010). Schumpeter meeting Keynes: A policy-friendly model of endogenous growth and business cycles. *Journal of Economic Dynamics and Control*, 34(9), 1748–1767. <https://doi.org/10.1016/j.jedc.2010.06.018>
- Dosi, G., Pereira, M. C., Roventini, A., & Virgillito, M. E. (2019). What if supply-side policies are not enough? The perverse interaction of flexibility and austerity. *Journal of Economic Behavior & Organization*, 162, 360–388. <https://doi.org/10.1016/j.jebo.2018.11.026>
- Dosi, G., Pereira, M. C., Roventini, A., & Virgillito, M. E. (2022). Technological paradigms, labour creation and destruction in a multi-sector agent-based model. *Research Policy*, 51(10), 104565. <https://doi.org/10.1016/j.respol.2022.104565>
- Eloundou, T., Manning, S., Mishkin, P., & Rock, D. (2024). GPTs are GPTs: Labor market impact potential of LLMs. *Science*, 384(6702), 1306–1308. <https://doi.org/10.1126/science.adj0998>
- Evans, G. W., & Honkapohja, S. (2001). *Learning and Expectations in Macroeconomics*. Princeton University Press. <https://www.jstor.org/stable/j.ctt7s6t9>
- Fagiolo, G., & Roventini, A. (2017). Macroeconomic Policy in DSGE and Agent-Based Models Redux: New Developments and Challenges Ahead. *Journal of Artificial Societies and Social Simulation*, 20(1), 1. <https://doi.org/10.18564/jasss.3280>
- Felten, E., Raj, M., & Seamans, R. (2021). Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses. *Strategic Management Journal*, 42(12), 2195–2217. <https://doi.org/10.1002/smj.3286>
- Frey, C. B., & Osborne, M. A. (2017). The future of employment: How susceptible are jobs to computerisation? *Technological Forecasting and Social Change*, 114, 254–280. <https://doi.org/10.1016/j.techfore.2016.08.019>
- Georgieff, A., & Hyee, R. (2022). Artificial Intelligence and Employment: New Cross-Country Evidence. *Frontiers in Artificial Intelligence*, 5, 832736. <https://doi.org/10.3389/frai.2022.832736>

- Gmyrek, P., Berg, J., & Bescond, D. (2023). *Generative AI and jobs: A global analysis of potential effects on job quantity and quality* (ILO Working Paper No. 96). International Labour Organization. Geneva. <https://doi.org/10.54394/FHEM8239>
- Goos, M., Manning, A., & Salomons, A. (2014). Explaining job polarization: Routine-biased technological change and offshoring. *American Economic Review*, 104(8), 2509–2526. <https://doi.org/10.1257/aer.104.8.2509>
- Grimm, V., Revilla, E., Berger, U., Jeltsch, F., Mooij, W. M., Railsback, S. F., Thulke, H.-H., Weiner, J., Wiegand, T., & DeAngelis, D. L. (2005). Pattern-oriented modeling of agent-based complex systems: Lessons from ecology. *Science*, 310(5750), 987–991. <https://doi.org/10.1126/science.1116681>
- Guerreiro, J., Rebelo, S., & Teles, P. (2022). Should Robots Be Taxed? *The Review of Economic Studies*, 89(1), 279–311. <https://doi.org/10.1093/restud/rdab019>
- Hagemann, H. (2019). “Economic Possibilities for our Grandchildren”. In R. W. Dimand & H. Hagemann (Eds.), *The Elgar Companion to John Maynard Keynes* (pp. 162–167). Edward Elgar. <https://doi.org/10.4337/9781788118569.00035>
- Hannan, M. T., & Freeman, J. (1984). Structural inertia and organizational change. *American Sociological Review*, 49(2), 149–164. <https://doi.org/10.2307/2095567>
- Hernandez, D., & Brown, T. B. (2020). *Measuring the Algorithmic Efficiency of Neural Networks*. arXiv: 2005.04305 [cs]. <https://doi.org/10.48550/arXiv.2005.04305>
- KPMG. (2025). *Cost of Capital Study 2025* (Industry Report). Retrieved August 17, 2026, from <https://assets.kpmg.com/content/dam/kpmgsites/ch/pdf/cost-of-capital-study-2025.pdf>
- Lavoie, M., & Stockhammer, E. (2013). Wage-led growth: Concept, theories and policies. In M. Lavoie & E. Stockhammer (Eds.), *Wage-led growth: An equitable strategy for economic recovery* (pp. 13–39). Palgrave Macmillan UK. https://doi.org/10.1057/9781137357939_2
- Maslej, N., Fattorini, L., Perrault, R., Gil, Y., Parli, V., Kariuki, N., Capstick, E., Reuel, A., Brynjolfsson, E., Etchemendy, J., Ligett, K., Lyons, T., Manyika, J., Niebles, J. C., Shoham, Y., Wald, R., Walsh, T., Hamrah, A., Santarlasci, L., ... Oak, S. (2025). Artificial Intelligence Index Report 2025. <https://doi.org/10.48550/ARXIV.2504.07139>
- Meier, I., & Tarhan, V. (2007, January). Corporate Investment Decision Practices and the Hurdle Rate Premium Puzzle. <https://doi.org/10.2139/ssrn.960161>
- OECD. (2019). *OECD Input to the Netherlands Independent Commission on the Regulation of Work* (tech. rep.). Organisation for Economic Co-operation and Development. Retrieved April 30, 2026, from https://www.eerstekamer.nl/overig/20190621/_oecd_input_to_the_netherlands/document
- SER. (2025, May). *AI en werk: Samen naar een werkende toekomst met AI* (Advies No. 25/03). Sociaal-Economische Raad. Den Haag. <https://www.ser.nl/nl/adviezen/ai-en-werk>
- Simon, H. A. (1955). A behavioral model of rational choice. *The Quarterly Journal of Economics*, 69(1), 99–118. <https://doi.org/10.2307/1884852>
- Storm, S. (2017). The New Normal: Demand, Secular Stagnation, and the Vanishing Middle Class. *International Journal of Political Economy*, 46(4), 169–210. <https://doi.org/10.1080/08911916.2017.1407742>
- Svanberg, M., Li, W., Fleming, M., Goehring, B., & Thompson, N. (2024, January). Beyond AI exposure: Which tasks are cost-effective to automate with computer vision? <https://doi.org/10.2139/ssrn.4700751>
- Tesfatsion, L. (2006). Agent-based computational economics: A constructive approach to economic theory. In L. Tesfatsion & K. L. Judd (Eds.), *Handbook of computational economics* (pp. 831–880, Vol. 2). Elsevier. [https://doi.org/10.1016/S1574-0021\(05\)02016-2](https://doi.org/10.1016/S1574-0021(05)02016-2)
- Upreti, A., & Sridhar, V. (2024). Effect of automation of routine and non-routine tasks on labour demand and wages. *IIMB Management Review*, 36(4), 289–308. <https://doi.org/10.1016/j.iimb.2024.09.001>
- Vergeer, R., & Kleinknecht, A. (2010). The impact of labor market deregulation on productivity: A panel data analysis of 19 OECD countries. *Journal of Post Keynesian Economics*, 33(2), 371–408. <https://doi.org/10.2753/PKE0160-3477330208>
- Webb, M. (2019, November). The Impact of Artificial Intelligence on the Labor Market. <https://doi.org/10.2139/ssrn.3482150>
- Wet arbeidsmarkt in balans. (2021, January). <https://wetten.overheid.nl/BWBR0042307/2021-01-01>

- Wet minimumloon en minimumvakantiebijslag. (2026, February). Retrieved April 8, 2026, from <https://wetten.overheid.nl/BWBR0002638/2026-02-04>
- Wet op het algemeen verbindend en het onverbindend verklaren van bepalingen van collectieve arbeidsovereenkomsten. (2023, June). Retrieved April 30, 2026, from <https://wetten.overheid.nl/BWBR0001987/2023-06-01>
- Wet werk en zekerheid. (2020, January). <https://wetten.overheid.nl/BWBR0035254/2020-01-01>
- Wright, T. P. (1936). Factors Affecting the Cost of Airplanes. *Journal of the Aeronautical Sciences*, 3(4), 122–128. <https://doi.org/10.2514/8.155>
- WRR. (2021). *Opgave AI: De nieuwe systeemtechnologie* (tech. rep. No. 105). Wetenschappelijke Raad voor het Regeringsbeleid. Den Haag. https://www.wrr.nl/site/binaries/site-content/collectiondocuments/2021/11/11/opgave-ai-de-nieuwe-systeemtechnologie/WRRrapport_+Opgave+AI_De+nieuwe+systeemtechnologie_NR105WRR.pdf
- Zhang, P. (2019). Automation, wage inequality and implications of a robot tax. *International Review of Economics & Finance*, 59, 500–509. <https://doi.org/10.1016/j.iref.2018.10.013>



Implementation

This appendix describes how the conceptual model from chapter 3 is translated into code. The seven conceptual components (C1-C7) and the macro-feedback layer (M1) from chapter 3 are the organisation structure, the code uses the same identifiers. The complete source code, including all experiment and analysis scripts and the pinned dependencies, is openly available at https://github.com/Bedjarn/LabourMarket_ABM_TU.

A.1. Software stack

The model is implemented in Python (3.12.3) with the help of the Mesa agent-based-modelling framework. NumPy handles the calculations and pandas the dataframes. Matplotlib produces the figures. A separate Streamlit application offers an interactive dashboard that was used for sense-making and parameter calibration (section A.5). The sensitivity analysis is a one-factor-at-a-time (OFAT) sweep that runs the model multiple times; the results can also be viewed in the dashboard. The exact package versions used to produce the results in this thesis are listed in Table 13; they are also pinned in *requirements.txt*.

Table 13: Software dependencies and the exact versions used for the thesis results.

Package	Version	Role
Mesa	2.1.4	Agent-based-modelling framework
NumPy	1.26.4	Numerical calculations
pandas	2.1.3	Dataframes and tabular output
Matplotlib	3.8.4	Figures
Streamlit	1.32.0	Interactive dashboard
openpyxl	3.1.2	Excel export of the results
SciPy	1.13.1	Statistical tests in the validation/verification notebooks

A.2. Mapping of components to code

The model is not split in one file per component. In an agent-based model the behaviour of a component is distributed over different parts. Production- and adoption are in the producer-agents, while wage setting and the goods market are on a model level. An implementation with one file per component would break-apart tightly connected code. The implementation keeps the whole model in one script, *labour_market_model.py*, and makes the component structure explicit by using banners and inline comments: every code block has a comment with the identifier of a component ([C1]...[C7], [M1]). All the code of one component is findable by searching on an identifier.

The repository, available at https://github.com/Bedjarn/LabourMarket_ABM_TU, is structured as shown in Table 14.

Table 14: Repository structure: the location in the code of each part of the implementation.

Location in code	Role
<i>model/labour_market_model.py</i>	The complete ABM: with classes Task, Worker, Producer and LabourMarketModel, with all the components.
<i>model/base_parameters.py</i>	The baseline parameterset. The parameter groups are labelled per component.
<i>model/live_dashboard.py</i> , <i>model/dashboard_utils.py</i>	The front-end of the Streamlit-dashboard and the back-end helpers.
<i>model/run_live_dashboard.py</i>	Entry point that launches the Streamlit dashboard.
<i>model/run_model.py</i>	Dashboardless batch-runner of the four adoption modes.
<i>model/multirun_experiments.py</i>	The multi-seed engine for the experimental programme of chapter 4.
<i>model/trajectory_stats.py</i>	The paired permutation trajectory test used to compare configurations.
<i>model/analysis/</i>	Secondary analysis scripts: the OFAT-sweep and the Excel-export of the OFAT-sweep.
<i>notebooks/</i>	The verification and validation notebooks, with the saved validation figures (Appendix E), and the Gmyrek et al. (2023) data extraction.
<i>data/</i>	The input data: the AI-exposure scores (2023) and the occupation–task breakdowns.
<i>data/Experiments/</i>	The saved multi-seed experiment results (workbooks E1–E8) reported in chapter 5.
<i>README.md</i> , <i>requirements.txt</i>	<i>require-</i> Run instructions and the pinned dependencies of Table 13.

Table 15 maps the conceptual components to the code.

Table 15: Mapping of the conceptual components (C1–C7, M1) to the code.

Component	Location in code
C1 Production	<i>leontief_output()</i> , <i>leontief_bottlenecks()</i> , the class Task, and the production-methods of the Producer-class: <i>productivity_human()</i> , <i>productivity_ai()</i> , <i>produce()</i> , <i>trim_ai_capacity()</i> .
C2 Labour market	The class Worker; <i>Producer.on_worker_separated()</i> and <i>hiring_requests()</i> ; and <i>LabourMarketModel.separate_workers()</i> , <i>match_and_hire()</i> , <i>hire_worker_to_task()</i> and <i>_warm_start()</i> .
C3 Wage setting	<i>LabourMarketModel.update_wages()</i> and <i>calculate_wage_bill()</i> .
C4 AI-adoption	The unit-cost methods of the Producer-class (<i>unit_cost_human()</i> , <i>unit_cost_ai()</i> , <i>preferred_input_for_task()</i>) together with <i>calculate_npv()</i> , <i>should_automate()</i> and <i>record_automation_decision()</i> .
C5 Learning curve	<i>LabourMarketModel.ai_cost_at()</i> , the closed path of the AI-rent path.
C6 Demand	<i>LabourMarketModel.calculate_profit_income()</i> and <i>calculate_market_price()</i> .
C7 Employment protection	<i>Producer.severance_cost_for_task()</i> and <i>LabourMarketModel.process_chain_limit()</i> , together with the attributes <i>contract_type</i> and <i>tenure</i> of the Worker-class.

Component	Location in code
M1 Macro feedback	LabourMarketModel.step(), that sets the order in which C1–C7 are activated, and <code>_compute_all_stats()</code> , that registers the per-step output.

A.3. The simulation loop and the order of sub-steps

LabourMarketModel.step() is the macro-feedback layer M1. The order of the 14 numbered substeps is set and each sub-step is labelled to the component of which it triggers a mechanism. Some substeps delegate to their own method, `process_chain_limit` (C7), `separate_workers()` (C2), `update_wages()` (C3), `match_and_hire()` (C2), `calculate_market_price()` (C6), while others have short inline blocks, like moving the learning curve or raising the tenure years. Production and price setting are performed two times in a step: one after the labour market events and once at the end, so that the data of that step is registered and up-to-date. The last sub-step calculates the aggregated statistics and passes them to Mesa's data collector.

A.4. Warm-start: initialising a fully staffed economy

The warm-start is a procedure where the model on $t = 0$ is set in a realistic starting position. It is implemented in `LabourMarketModel._warm_start()` and is called in the initialisation of the model. In this step some decisions are made that affect the outcome of the experiments, each explained below.

A.4.1. What the warm-start does

The warm-start distributes all workers over the tasks of all firms following the same rule that is used in the simulation for matching: per task it is tried to get the cheapest worker from the preferred skill-pool, if all are hired the other skill pool is used as a fallback. The order in which tasks are filled is decided by the seed. If there are enough workers to fill all tasks all tasks should have one worker at $t = 0$. In the baseline there are 450 workers (200 high-skilled and 250 low-skilled) but only 400 tasks (20 firms $\times$ 20 tasks). Every task is therefore staffed at $t = 0$, while the surplus of about 50 workers remains unemployed: the economy starts at its employment ceiling, an aggregate unemployment rate of roughly 11%, which is the lowest the model can mechanically reach because labour supply is deliberately set above the number of jobs.

After the allocation `update_wages(instant=True)` is called. That sets the wage levels for high and low directly to the target wage (so no delay to the target wage with λ in subsection 3.8.2). Immediately after the warm start the initial stat is collected for an accurate $t = 0$ measurement, without wage-adjustments or separations.

A.4.2. Why the warm-start fills every job

The warm-start deliberately initialises the model with every job filled: every task receives a worker, and the matching rule contains no acceptability check that could block the assignment. This is a modelling choice rather than a convenience, and it rests on three reasons.

First, it isolates the research question. The aim of the model is to identify how AI adoption shifts the employment composition of a functioning labour market, not how a labour market converges from an arbitrary disequilibrium. If the simulation were to start with empty tasks or with a substantial pool of inactive workers, the first several dozen steps would be dominated by the absorption of that initial slack rather than by the mechanisms under study. A particularly misleading artefact is that an empty initial economy would produce a rising employment trajectory purely because the simulation begins at zero employment and workers are progressively hired against reservation wages, a pattern that has nothing to do with AI adoption dynamics but that would be very difficult to disentangle from the results in the figures. Starting with every job already filled removes that confound: any employment decline observed over the simulation horizon is unambiguously attributable to the endogenous mechanisms of the model (separations, demand contraction, automation), not to an arbitrary initial allocation.

Second, it aligns with the empirical Dutch reference point. The model is calibrated against the Dutch labour market, where the unemployment rate has remained close to its structural level in recent years. Initialising with every job filled is therefore not an extreme idealisation but a defensible upper

reference point for the starting condition against which the scenario experiments are evaluated; the structural unemployment that remains at $t = 0$ reflects a tight, fully operational market rather than cyclical slack. The warm-start also provides the natural moment to seed the tenure distribution and the permanent/flex contract mix required by component C7 (employment protection): without an initial allocation of workers to firms there is simply no substrate on which contract types and tenure can be defined, and the WAB/WWZ mechanisms would have to be initialised from a fictional zero-tenure population.

Third, the internal mechanics of the model require a populated economy at $t = 0$. With zero workers hired at the start, no firm would generate output, no firm would generate revenue, and the gate $AC < p$ that governs proactive expansion would be undefined or closed at every firm simultaneously. Every one of the twenty tasks for all firms would register as a binding Leontief bottleneck on $t = 0$, which would in turn force every firm to take an all-or-nothing investment decision over the all tasks while producing exactly nothing, a case that the NPV mechanism is not designed to evaluate sensibly. Starting with every task staffed circumvents this: firms enter the first model step with a well-defined operating output, a well-defined revenue, and a well-defined cost structure against which the adoption rules in component C4 can meaningfully be compared.

A related reason concerns wage initialisation. Calling `update_wages(instant=True)` immediately after the allocation places w_h and w_l exactly on the level implied by the realised employment counts L_h and L_l . These levels are precisely the ones against which the rest of the parameter set has been calibrated to mimic the Dutch wage structure. Without that instant adjustment, wages would have to drift toward their equilibrium over many steps through partial adjustment at speed λ , generating a curve that would feed through component C6 into demand and through component C4 into the very first automation decisions, once again contaminating the experimental window with start-up artefacts rather than the dynamics of interest.

Taken together, these reasons reflect a single methodological commitment: the experiments in this thesis are designed to study how a working labour market, one that is already operational and labour-driven by construction, finds its way toward a partially or fully automated configuration as AI is introduced. They are not designed to study how an empty economy bootstraps itself from zero employment into some mixed steady state. The warm-start makes that distinction explicit at $t = 0$.

A.5. Calibration and default values

The model uses one baseline set of parameters as defined and explained in Appendix C. Each parameter is calibrated using one of four types: mandate, data-anchored, pattern-oriented, or replication-anchored. Below the most important ones are shown.

Wage rates, monetary scale, and the time step.

The model is advanced in monthly steps (`steps_per_year = 12`). All monetary variables are expressed in the same monetary scale M (€) per time step. Because the time step is one month, w is dimensionally a wage-per-month in the model. The model equations are linearly homogeneous in M and are therefore invariant under any simultaneous rescaling of all monetary variables.

This homogeneity makes it possible to calibrate the numerical values of wage parameters on a scale recognisable for the Dutch labour market, namely that of the hourly wage. For example, the minimum wage $w_{\min} = €14.71$ is set at the value of the Dutch hourly minimum wage, the wage curves intercepts $a_h = 15$ and $a_l = 6$ with slopes $b_h = 0.10$ and $b_l = 0.06$ produce wages of around €35 and €21 respectively at full employment, which correspond to the CBS (2025e) averages for higher professional/university educated and lower educated workers measured on an hourly basis. This choice of scale is a calibration anchor, not a dimensional statement: operationally, w functions in the model as a wage-per-time step.

Under this convention, the implementation of Dutch labour market institutions remains internally consistent. The transition payment is statutory fixed at one-third of a monthly salary per year of service; in the model formula:

$$severance = severance_rate \cdot w \cdot tenure_in_years$$

with the severance rate of 1/3, this by definition gives "1/3 of one time-step wage per year of service", structurally identical to the law, because the time step is one month. For a permanent employee on the minimum wage with five years of service, the model produces a transition payment of $(1/3) \cdot 14.71 \cdot 5 \approx 24.5$ monetary units (€ in the thesis), or 1.67 time-step wages in model units, the same ratio (1.67 months of wages) that the WAB yields for this employee in reality. Here $w = 14.71$ is the per-step (monthly) wage in the model; that this number is calibrated to the Dutch hourly minimum wage is a scale convention (explained below) and does not affect the ratio.

Although actual monthly wages in the Netherlands are obviously much higher than €25 per month, the model values can be interpreted more intuitively by reading them on an hourly-wage scale. This does not affect the model's substantive results, because the monetary scale is used consistently and the analysis depends on relative rather than absolute monetary values. Converting the wage variables to a monthly full-time scale would simply require multiplying them by approximately 160 hours per month. This rescaling would change the numerical units, but not the model dynamics, ratios, or conclusions. The hourly-wage interpretation therefore serves only as a calibration convention that makes the parameters easier to understand.

A.6. Streamlit-dashboard

A Streamlit dashboard was developed alongside the model to provide an interactive front-end for parameter setting, live visualisation of the model's behaviour, and configuration of individual runs. The dashboard wraps `LabourMarketModel` with sliders and selection widgets for every parameter in `base_parameters.py`, renders the main time series (employment, wages, prices, AI adoption rates) as the simulation advances, and exposes a set of diagnostic preview tabs that show internal states which are not part of the regular output, like bottleneck distributions, investment-log entries, productivity profiles, and the realised severance per task.

It was used throughout the development process for three distinct purposes. First, for sense-making and debugging: stepping through individual model steps with the sliders in view made it possible to verify that mechanisms behaved as intended and to localise anomalies quickly when they did not. Second, for the calibration of the pattern-oriented parameters: parameters without a direct empirical counterpart were tuned by inspection until the model reproduced the six validation patterns of Appendix E. Third, for the post-hoc analysis: the dashboard was used to interpret the OFAT sensitivity results, and to design the configuration of the experiments themselves and collect output graphs of them.

The dashboard is therefore a development and analysis instrument, not the engine behind the reported experiments.

B

Mean-field decision and complete
pairwise comparisons

This appendix presents the complete pairwise comparisons between all three modes for Experiments 1 through 8, including the mutual pairs between non-baseline configurations. All configurations are run in the mean-field NPV mode (MF), except for the comparison of the three NPV modes and Experiments 1 and 2 (which involve the ULC configuration). The following applies to each row: $n = 50$ seeds; p comes from the paired permutation test with 5000 sign permutations; difference (A-B) is the average difference over the horizon; CI-low/CI-high are the bounds of a percentile bootstrap 95% confidence interval (resampled over the seeds) for that horizon-averaged difference; d_z is the standardised effect; and 'Sign. horizon' is the share of steps for which the simultaneous band excludes zero. The 'Sign.' column and the p -value come from the global sup- t permutation test, which asks whether the two trajectories differ at *any* point along the horizon; CI-low/CI-high instead bracket the *average* difference over the whole horizon. A row can therefore be globally significant ($p < 0.001$) while the confidence interval of the averaged difference still contains zero, namely when a difference that is concentrated on part of the trajectory averages out to near zero over all 960 steps. Almost all p -values are below 0.001, so the effect size and the significant horizon are the informative quantities. Because this appendix reports a large family of pairwise tests (eight experiments, three outcomes, and all configuration pairs within each), the per-test p -values are not adjusted for multiple comparisons; the conclusions accordingly rest on the effect size (d_z) and the share of the horizon that is significant rather than on any single p -value.

B.1. Comparison of the three NPV modes

The three expectation modes, mean-field, naïve and adaptive, are compared with one another in the same way as the other experiments, as described in section 4.6. All runs use the same parameters; only the NPV mode differs.

Table 16: Complete pairwise comparison for the three NPV-modes.

A	Mode A	B	Mode B	Outcome	n	p	Sign.	Difference (A-B)	CI-low	CI-high	d_z	Sign. horizon
Mean-field	MF	Naive	NV	Gini	50	< 0.001	TRUE	0.009475	0.004015	0.01489	0.48	23%
Mean-field	MF	Naive	NV	Skill-wage premium	50	0.0390	TRUE	0.003053	0.0006218	0.00545	0.36	0%
Mean-field	MF	Naive	NV	Unemployment	50	< 0.001	TRUE	0.0099	0.004177	0.01547	0.48	22%
Mean-field	MF	Adaptive	AD	Gini	50	< 0.001	TRUE	0.006227	0.001922	0.0104	0.40	10%
Mean-field	MF	Adaptive	AD	Skill-wage premium	50	0.2809	FALSE	0.002241	-2.818e-05	0.004528	0.28	0%
Mean-field	MF	Adaptive	AD	Unemployment	50	< 0.001	TRUE	0.006443	0.002172	0.01075	0.40	11%
Naive	NV	Adaptive	AD	Gini	50	< 0.001	TRUE	-0.003248	-0.007388	0.0009772	-0.22	4%
Naive	NV	Adaptive	AD	Skill-wage premium	50	0.0052	TRUE	-0.0008128	-0.002967	0.001237	-0.11	2%
Naive	NV	Adaptive	AD	Unemployment	50	< 0.001	TRUE	-0.003456	-0.007537	0.0008325	-0.23	4%

The test in Table 16 shows that the three modes are statistically different from each other: all p -values are below 0.001. At the same time the effects are small and the significant horizon is limited (highest is 23%). By contrast the difference between the Mean-field approach and the ULC-configuration has a significant horizon of 100% as seen in Table 17. The choice of expectation mode is therefore a second-order effect: it shifts the outcomes noticeably but modestly, whereas institutional expansions and investment logic determine the picture in the first order. Thus, the qualitative conclusions of this thesis are robust to the choice of expectation mode.

The ordering is moreover consistent with the theory. For both inequality and unemployment the mean-field mode yields the strongest outcome, followed by adaptive and then naïve, exactly the order in which the modes internalise the macro-feedback of their own adoption: mean-field discounts that automation depresses wages through rising unemployment, adaptive lags this trend, and naïve ignores it. The skill-wage premium barely responds to the mode, which

indicates that the expectation mode governs the level of automation rather than the wage structure between high and low-skilled workers (Figures 13, 14 and 15)



Figure 13: Gini index over the 960-step horizon for the three NPV modes (naïve, adaptive, mean-field). Mean over 50 seeds with a ± 1 standard-deviation band.




Figure 14: Skill-wage premium over the 960-step horizon for the three NPV modes (naïve, adaptive, mean-field). Mean over 50 seeds with a ± 1 standard-deviation band.




Figure 15: Unemployment rate over the 960-step horizon for the three NPV modes (naïve, adaptive, mean-field). Mean over 50 seeds with a ± 1 standard-deviation band.

B.2. Experiment 1

Table 17 reports the complete pairwise comparison; Figures 16, 17 and 18 show the corresponding trajectories.

Table 17: Complete pairwise comparison for experiment 1.

A	Mode A	B	Mode B	Outcome	n	p	Sign.	Difference (A-B)	CI-low	CI-high	d_z	Sign. horizon
MF	MF	ULC	ULC	Gini	50	< 0.001	TRUE	-0.1741	-0.1788	-0.1692	-9.86	100%
MF	MF	ULC	ULC	Skill-wage premium	50	< 0.001	TRUE	-0.0707	-0.07685	-0.06442	-3.15	97%
MF	MF	ULC	ULC	Unemployment	50	< 0.001	TRUE	-0.1814	-0.1864	-0.1763	-9.79	100%




Figure 16: Gini index over the horizon, Experiment 1 (mean-field NPV baseline versus ULC). Mean over 50 seeds with a ± 1 standard-deviation band.




Figure 17: Skill-wage premium over the horizon, Experiment 1 (mean-field NPV baseline versus ULC). Mean over 50 seeds with a ± 1 standard-deviation band.




Figure 18: Unemployment rate over the horizon, Experiment 1 (mean-field NPV baseline versus ULC). Mean over 50 seeds with a ± 1 standard-deviation band.

B.3. Experiment 2

Table 18 reports the decision-logic step of the decomposition (Experiment 2, section 5.2): the mean-field NPV with employment protection switched off, against the ULC configuration, both without institutions. Figures 19, 20 and 21 show the trajectories. The institutional step (baseline versus mean-field protection-off) is reported under Experiment 3 below.

Table 18: Decision-logic step of the Experiment 1 decomposition: mean-field NPV with employment protection off versus ULC, both without institutions.

A	Mode A	B	Mode B	Outcome	n	p	Sign.	Difference (A-B)	CI-low	CI-high	d_z	Sign. horizon
MF protection off	MF	ULC	ULC	Gini	50	< 0.001	TRUE	-0.1609	-0.1669	-0.1544	-6.92	100%
MF protection off	MF	ULC	ULC	Skill-wage premium	50	< 0.001	TRUE	-0.065	-0.07087	-0.05909	-3.02	99%
MF protection off	MF	ULC	ULC	Unemployment	50	< 0.001	TRUE	-0.1677	-0.1744	-0.161	-6.81	100%



Figure 19: Gini index over the horizon, decision-logic step (mean-field protection-off versus ULC). Mean over 50 seeds with a ± 1 standard-deviation band.




Figure 20: Skill-wage premium over the horizon, decision-logic step (mean-field protection-off versus ULC). Mean over 50 seeds with a ± 1 standard-deviation band.




Figure 21: Unemployment rate over the horizon, decision-logic step (mean-field protection-off versus ULC). Mean over 50 seeds with a ± 1 standard-deviation band.

B.4. Experiment 3

Table 19 reports the pairwise comparison of the baseline (employment protection on) with the protection-off run (MF-off), both mean-field, the institutional step of the decomposition (section 5.3) examined over the full horizon. Figures 22, 23 and 24 show the trajectories; Figures 25 and 26 show the paired permutation-test bands for the Gini index and the skill-wage premium, in which the significant mass sits at both ends of the horizon.

Table 19: Pairwise comparison for experiment 3 (employment protection on versus off, both mean-field).

A	Mode A	B	Mode B	Outcome	n	p	Sign.	Difference (A-B)	CI-low	CI-high	d_z	Sign. horizon
Protection on	MF	Protection off	MF	Gini	50	< 0.001	TRUE	-0.01711	-0.02376	-0.01055	-0.71	35%
Protection on	MF	Protection off	MF	Skill-wage premium	50	< 0.001	TRUE	-0.009558	-0.01236	-0.006652	-0.90	29%
Protection on	MF	Protection off	MF	Unemployment	50	< 0.001	TRUE	-0.01736	-0.02423	-0.01068	-0.69	34%



Figure 22: Gini index over the horizon, Experiment 3 (employment protection on versus off). Mean over 50 seeds with a ± 1 standard-deviation band.




Figure 23: Skill-wage premium over the horizon, Experiment 3 (employment protection on versus off). Mean over 50 seeds with a ± 1 standard-deviation band.




Figure 24: Unemployment rate over the horizon, Experiment 3 (employment protection on versus off). Mean over 50 seeds with a ± 1 standard-deviation band.




Figure 25: Paired permutation test for the Gini index, Experiment 3 (employment protection off versus on). The upper panel shows the two configurations (mean $\pm$ 1 sd); the lower panel shows the paired difference (A-B) with its 95% simultaneous band, with the significant steps shaded. The significant mass lies at the start and the end of the horizon.




Figure 26: Paired permutation test for the skill-wage premium, Experiment 3 (employment protection off versus on). Panels as in Figure 25.

B.5. The skill-wage premium as an initialisation artefact

The skill-wage premium w_h/w_l falls from roughly 1.8 to about 1.3 in every scenario, including those in which almost nothing is automated and those in which automation is switched off entirely. subsection 6.5.4 states the conclusion; the sequence that produces it is set out here, because it is a property of

the initialisation rather than a result about automation.

Step 1: near-equal starting wages. At $t = 0$ the two wages sit close together. The low-skill wage is on the minimum-wage floor of €14.71 and the high-skill wage on its wage-curve intercept of €15. Because high-skilled workers are more productive on non-routine tasks, at these near-equal wages they are the cheaper choice on almost every non-routine task.

Step 2: the warm-start allocation. The matching rule therefore places the high-skilled on non-routine work and the low-skilled on routine work: in the baseline run, 185 and 215 workers respectively.

Step 3: the wages that allocation implies. Only after the allocation does *update_wages(instant=True)* set wages to the level the realised employment counts imply, and those pull far apart: about €33 for the high-skilled against about €19 for the low-skilled.

Step 4: the cost comparison flips. With the wages this far apart, a high-skilled worker is only the cheaper option on the most complex non-routine tasks. On everything else the low-skilled worker has become cheaper. In the baseline this holds for 174 of the 185 jobs the high-skilled occupy, so almost their entire starting position is suddenly too expensive.

Step 5: self-repair through turnover. Whenever a worker leaves through ordinary separation, the vacated task is refilled with the now-cheaper low-skilled worker. High-skill employment falls, and because the high-skill wage is tied to high-skill employment through the wage curve, that wage slides down from about €33 towards €24 while the low-skill wage stays near €19. A shrinking top wage over a roughly stable bottom wage is, by definition, a falling premium.

Step 6: the process is self-limiting. As the high-skill wage drops, high-skilled workers become competitive on more tasks again, their employment levels off, and the premium settles around 1.3.

The drift is therefore driven by the combination of a steep high-skill wage curve ($b_h = 0.10$), near-equal starting wages, and productivity settings that give the high-skilled an advantage only on the most complex tasks. Because every configuration inherits the same initialisation, the drift is common to all scenarios and cancels in the paired comparisons reported in chapter 5, which is why the premium is read as a relative signal throughout this thesis and never as a level.

B.5.1. The full mechanism: why the institutional effect is small, saturating and persistent

This subsection works out in full the mechanism summarised in subsection 6.2.5: a parametric explanation for the size of the employment-protection effect, a structural one that damps it further, and a temporal one for its persistence.

A parametric explanation. The severance payment enters the NPV calculation as a one-off extra upfront cost, on top of the complexity-scaled investment cost $I(x) = I_{base} \cdot (1 + complexity_scaling \cdot x)$. With the calibrated values ($I_{base} = 80$; $complexity_scaling = 0.15$) the investment cost per task runs from about €92 (simplest task) to €260 (most complex task). The severance payment, by contrast, is $severance_rate \cdot wage \cdot tenure\ years$ per permanent worker, and is a fraction of that. A worked example makes the magnitudes concrete. Consider a low-skilled routine task of moderate complexity ($x = 3$) held by one permanent worker. The investment cost is $I(3) = 80 \cdot (1 + 0.15 \cdot 3) = 116$; the complexity term alone moves this cost from €92 to €260 across the complexity range, a span of roughly €168. The severance for that worker is $S = (1/3) \cdot wage \cdot tenure$: at a wage of €21 and four years of service this is about €28, and across the plausible band (wages €15–21, two to five years of service) it stays between roughly €10 and €35, smaller than the complexity component of the investment alone. Under the adoption mode a firm automates when the discounted savings exceed $(1 + \eta) \cdot I + S$, with hurdle rate $\eta = 0.15$. Adding severance raises that threshold from $1.15 \cdot 116 = 133$ to 161 (+21%); over the 36-month horizon (annuity factor ≈ 31.7 at a per-step

discount rate of 0.007) this corresponds to a rise in the required per-period saving from about €4.20 to €5.10, a wedge of roughly €0.90 against low-skilled wages of €15–21. Severance is therefore a modest one-off increment on a cost that is itself dominated by the complexity-scaled investment and by 36 periods of discounted savings. Because the severance wedge is this small, it can only flip the adoption decisions where the discounted AI savings sit just above the threshold. This explains the saturation: the statutory rate already captures those marginal tasks, so a higher rate finds no further decisions to flip, while removing the wedge releases exactly those tasks, which is why setting the rate to zero (Experiment 7) moves the outcomes almost as much as removing the whole institution (Experiment 3).

A structural explanation. The model structure damps the contract effect along a second route, and here the model diverges from reality. New employees always start in the model as flex workers with zero tenure; a permanent contract arises only through conversion once the chain regulation reaches its 36-step limit (assumption C7.4 in Appendix H). The build-up is therefore first-flex-then-permanent instead of directly-permanent. Because only permanent workers generate a severance payment and flex workers have no statutory entitlement, a task that is automated often affects workers who are still (partly) flexible, so the severance brake remains small. In reality, workers are also hired directly on permanent contracts, and long-running flexible contracts can carry a transition payment under conditions (assumption C4.8 in Appendix H); in this respect the modelled contract dynamics are a simplification that underestimates rather than overestimates the institutional asymmetry.

A small early brake that the demand feedback amplifies. The third feature, the persistence, is where an initially appealing reading has to be discarded. One might expect the severance brake to scale with the stock and tenure of permanent contracts and so to follow a strong–weak–strong path: hard early on the inherited high-tenure block, fading as those workers separate and their tasks are automated reactively at zero severance (subsection 3.12.2), then rebuilding as the chain regulation matures a fresh permanent workforce whose accumulating severance brakes automation again. The diagnostics below do not support that story. The severance owed per task does not accumulate towards the end of the horizon; it peaks early and *declines* before flattening (Figure 27), the average tenure of permanent workers *stabilises* around nine-and-a-half years rather than growing without bound (Figure 28), and the permanent-contract share and the conversion flow are flat over the second half. There is no late re-strengthening of the severance brake to re-open the gap.

What does re-open it is the model's own Keynesian demand feedback (subsection 3.11.2, section 3.13) amplifying a small early difference. This early difference is specifically the severance brake on automation, and it should not be confused with the skill-reallocation transient of subsection 6.5.4: the early displacement of high-skilled workers and the compression of the wage premium occur almost identically with and without employment protection, because they run through the wage curve and exogenous turnover rather than through severance, so they cancel in the paired comparison and do not contribute to the gap. The early automation that the severance brake slows falls moreover mainly on routine tasks, so in the first phase protection spares low-skilled routine jobs rather than high-skilled ones. The mechanism has three stages. First, a seed: severance enters the NPV as a one-off upfront cost on occupied tasks, and while it is large (the high-tenure warm-start block) and the AI price is still falling, it tips a meaningful share of the still-marginal adoption decisions, so the protected economy automates and displaces slightly less in the first phase. Second, the seed stops growing: once the AI price floors, automating is so cheap that the severance increment no longer flips any decision, the per-step automation gap settles to a small constant, and the severance brake, far from re-strengthening, fades. Third, amplification: the early extra displacement under protection-off lowers the wage bill, which lowers the demand, the price and hence firm profitability; when average cost approaches the price the expansion gate closes and firms stop hiring proactively, so the displaced, especially the cheap and easily automatable low-skilled, are not re-absorbed, which lowers the wage bill further. This loop is visible directly in the diagnostics below: expansion-gate blocks rise steeply without protection while staying flat with it (Figure 29), the demand level diverges downward in the protection-off run (Figure 30), and the gap concentrates in low-skill employment (Figure 31). The widening of the unemployment gap to five or six percentage points exceeds the roughly two points the stable automation gap can explain on its own; the remainder is re-absorption foregone through the

demand channel, not fresh automation. The AI price reaches its floor early and is flat throughout this phase, so it is not the driver. The on/off toggle of Experiment 3 and the removal of the severance payment in Experiment 7 bite almost identically, confirming that within this setup it is the severance payment, not the chain regulation or the contract build-up, that seeds the difference.

B.5.2. Diagnostics for the mechanism behind the temporal pattern

The diagnostic series below identify the mechanism behind the start-and-end significance of Experiment 3 (subsection 6.2.5), all averaged over the same 50 seeds with a ± 1 standard-deviation band. They show that the late re-opening of the gap is *not* a re-strengthening severance brake but the Keynesian demand feedback amplifying a small early automation difference. Two figures rule out the contract-dynamics reading: the average severance owed per task peaks early and then *declines* before flattening (Figure 27), and the average tenure of permanent workers *stabilises* around nine-and-a-half years rather than growing without bound (Figure 28); neither rises over the second half, so there is no late re-accumulation of severance. Three figures show the amplification instead: the number of expansion-gate blocks rises steeply without protection while staying flat with it (Figure 29), the demand level diverges downward in the protection-off run (Figure 30), and the resulting employment gap concentrates in the low-skilled (Figure 31).




Figure 27: Average severance owed per task over the horizon, Experiment 3 (only defined when employment protection is on). It peaks early on the high-tenure warm-start cohort and then declines before flattening, so the severance brake fades rather than re-strengthening. Mean over 50 seeds with a ± 1 standard-deviation band.




Figure 28: Average tenure of permanent workers over the horizon, Experiment 3 (only defined when employment protection is on). Tenure stabilises around nine-and-a-half years because exogenous separations cap it, rather than accumulating without bound. Mean over 50 seeds with a ± 1 standard-deviation band.




Figure 29: Keynesian expansion-gate blocks per step, Experiment 3 (employment protection on versus off). The count rises steeply in the protection-off run while staying flat with protection, the signature of the demand-feedback loop closing the expansion gate. Mean over 50 seeds with a ± 1 standard-deviation band.




Figure 30: Demand level over the horizon, Experiment 3 (employment protection on versus off). The protection-off run diverges downward late as its lower wage bill feeds the demand intercept. Mean over 50 seeds with a ± 1 standard-deviation band.




Figure 31: Low-skill employment rate over the horizon, Experiment 3 (employment protection on versus off). The late divergence concentrates here, around 73% without protection against 79% with it, because the closed expansion gate suppresses re-hiring of the cheapest, most automatable workers. Mean over 50 seeds with a ± 1 standard-deviation band.

B.6. Experiment 4

Table 20 reports the complete pairwise comparisons; Figures 32, 33 and 34 show the corresponding trajectories.

Table 20: All pairwise comparisons for experiment 4. Configurations: decreasing costs (baseline; initial AI cost 40, floor 20), flat (initial cost and floor both 30) and rising costs (initial cost 20, floor 40).

A	Mode A	B	Mode B	Outcome	n	p	Sign.	Difference (A-B)	CI-low	CI-high	d_z	Sign. horizon
Decreasing costs (Base)	MF	Flat	MF	Gini	50	< 0.001	TRUE	0.08835	0.08144	0.09508	3.53	99%
Decreasing costs (Base)	MF	Flat	MF	Skill-wage premium	50	< 0.001	TRUE	0.04613	0.04268	0.04955	3.67	94%
Decreasing costs (Base)	MF	Flat	MF	Unemployment	50	< 0.001	TRUE	0.08811	0.08112	0.09498	3.47	99%
Decreasing costs (Base)	MF	Rising costs	MF	Gini	50	< 0.001	TRUE	0.1478	0.144	0.1518	10.26	99%
Decreasing costs (Base)	MF	Rising costs	MF	Skill-wage premium	50	< 0.001	TRUE	0.02122	0.01842	0.02415	2.02	78%

Decreasing costs (Base)	MF	Rising costs	MF	Unemployment	50	< 0.001	TRUE	0.1568	0.1526	0.1612	9.89	99%
Flat	MF	Rising costs	MF	Gini	50	< 0.001	TRUE	0.05941	0.05297	0.06595	2.58	98%
Flat	MF	Rising costs	MF	Skill-wage premium	50	< 0.001	TRUE	-0.02491	-0.02802	-0.02157	-2.15	80%
Flat	MF	Rising costs	MF	Unemployment	50	< 0.001	TRUE	0.06871	0.06235	0.07514	2.86	98%




Figure 32: Gini index over the horizon, Experiment 4 (decreasing, flat and rising AI-cost trajectories). Mean over 50 seeds with a ± 1 standard-deviation band.




Figure 33: Skill-wage premium over the horizon, Experiment 4 (decreasing, flat and rising AI-cost trajectories). Mean over 50 seeds with a ± 1 standard-deviation band.




Figure 34: Unemployment rate over the horizon, Experiment 4 (decreasing, flat and rising AI-cost trajectories). Mean over 50 seeds with a ± 1 standard-deviation band.

B.7. Experiment 5

Table 21 reports the complete pairwise comparisons; Figures 35, 36 and 37 show the corresponding trajectories.

Table 21: All pairwise comparisons for experiment 5. Non-routine multiplier by label: low = 0.50, base = 0.85, medium = 1.15, high = 1.50.

A	Mode A	B	Mode B	Outcome	n	p	Sign.	Difference (A-B)	CI-low	CI-high	d_z	Sign. horizon
Base	MF	nr multiplier low	MF	Gini	50	< 0.001	TRUE	0.0823	0.07869	0.08642	5.83	100%
Base	MF	nr multiplier low	MF	Skill-wage premium	50	< 0.001	TRUE	0.04409	0.04175	0.04649	5.11	92%
Base	MF	nr multiplier low	MF	Unemployment	50	< 0.001	TRUE	0.08224	0.07842	0.08639	5.53	100%
Base	MF	nr multiplier medium	MF	Gini	50	< 0.001	TRUE	-0.02633	-0.03102	-0.02141	-1.51	64%
Base	MF	nr multiplier medium	MF	Skill-wage premium	50	< 0.001	TRUE	-0.006506	-0.01341	0.0003906	-0.26	14%
Base	MF	nr multiplier medium	MF	Unemployment	50	< 0.001	TRUE	-0.02782	-0.03248	-0.02295	-1.59	66%
Base	MF	nr multiplier high	MF	Gini	50	< 0.001	TRUE	-0.04944	-0.05848	-0.04083	-1.53	100%

Base	MF	nr multiplier high	MF	Skill-wage premium	50	< 0.001	TRUE	0.01273	0.0009712	0.02371	0.30	5%
Base	MF	nr multiplier high	MF	Unemployment	50	< 0.001	TRUE	-0.05603	-0.06446	-0.04794	-1.83	100%
nr multiplier low	MF	nr multiplier medium	MF	Gini	50	< 0.001	TRUE	-0.1086	-0.1121	-0.1053	-8.52	100%
nr multiplier low	MF	nr multiplier medium	MF	Skill-wage premium	50	< 0.001	TRUE	-0.05059	-0.0571	-0.0439	-2.11	94%
nr multiplier low	MF	nr multiplier medium	MF	Unemployment	50	< 0.001	TRUE	-0.1101	-0.1135	-0.1068	-8.95	100%
nr multiplier low	MF	nr multiplier high	MF	Gini	50	< 0.001	TRUE	-0.1317	-0.1403	-0.1239	-4.42	100%
nr multiplier low	MF	nr multiplier high	MF	Skill-wage premium	50	< 0.001	TRUE	-0.03136	-0.04223	-0.02022	-0.78	90%
nr multiplier low	MF	nr multiplier high	MF	Unemployment	50	< 0.001	TRUE	-0.1383	-0.1466	-0.1309	-4.90	100%
nr multiplier medium	MF	nr multiplier high	MF	Gini	50	< 0.001	TRUE	-0.0231	-0.02964	-0.01737	-1.04	81%
nr multiplier medium	MF	nr multiplier high	MF	Skill-wage premium	50	< 0.001	TRUE	0.01923	0.01204	0.02614	0.76	52%
nr multiplier medium	MF	nr multiplier high	MF	Unemployment	50	< 0.001	TRUE	-0.02821	-0.03459	-0.02256	-1.29	92%




Figure 35: Gini index over the horizon, Experiment 5 (non-routine multiplier 0.50, 0.85, 1.15, 1.50). Mean over 50 seeds with a ± 1 standard-deviation band.




Figure 36: Skill-wage premium over the horizon, Experiment 5 (non-routine multiplier 0.50, 0.85, 1.15, 1.50). Mean over 50 seeds with a ± 1 standard-deviation band.




Figure 37: Unemployment rate over the horizon, Experiment 5 (non-routine multiplier 0.50, 0.85, 1.15, 1.50). Mean over 50 seeds with a ± 1 standard-deviation band.

B.8. Experiment 6

Table 22 reports the complete pairwise comparisons; Figures 38, 39 and 40 show the corresponding trajectories.

Table 22: All pairwise comparisons for experiment 6. Minimum wage by label: no minimum wage = 0, low = 10, base = 14.71, high = 20.

A	Mode A	B	Mode B	Outcome	n	p	Sign.	Difference (A-B)	CI-low	CI-high	d_z	Sign. horizon
Base	MF	No minimum wage	MF	Gini	50	< 0.001	TRUE	-0.0009451	-0.006041	0.004177	-0.05	3%
Base	MF	No minimum wage	MF	Skill-wage premium	50	< 0.001	TRUE	0.009538	0.007102	0.01206	1.07	8%
Base	MF	No minimum wage	MF	Unemployment	50	< 0.001	TRUE	-0.002555	-0.007791	0.002645	-0.13	6%
Base	MF	Low minimum wage	MF	Gini	50	< 0.001	TRUE	-0.002948	-0.009151	0.003129	-0.13	4%
Base	MF	Low minimum wage	MF	Skill-wage premium	50	< 0.001	TRUE	0.008743	0.005694	0.01178	0.79	8%
Base	MF	Low minimum wage	MF	Unemployment	50	< 0.001	TRUE	-0.00465	-0.01106	0.001647	-0.20	6%
Base	MF	High minimum wage	MF	Gini	50	< 0.001	TRUE	-0.03003	-0.03458	-0.02513	-1.73	90%

Base	MF	High minimum wage	MF	Skill-wage premium	50	< 0.001	TRUE	0.03824	0.03461	0.04206	2.78	95%
Base	MF	High minimum wage	MF	Unemployment	50	< 0.001	TRUE	-0.0339	-0.03861	-0.02857	-1.85	95%
No minimum wage	MF	Low minimum wage	MF	Gini	50	0.1734	FALSE	-0.002003	-0.007591	0.003834	-0.10	0%
No minimum wage	MF	Low minimum wage	MF	Skill-wage premium	50	0.3337	FALSE	-0.0007945	-0.003706	0.002084	-0.08	0%
No minimum wage	MF	Low minimum wage	MF	Unemployment	50	0.3015	FALSE	-0.002094	-0.007871	0.003541	-0.10	0%
No minimum wage	MF	High minimum wage	MF	Gini	50	< 0.001	TRUE	-0.02908	-0.03324	-0.02452	-1.80	81%
No minimum wage	MF	High minimum wage	MF	Skill-wage premium	50	< 0.001	TRUE	0.0287	0.02444	0.03325	1.77	98%
No minimum wage	MF	High minimum wage	MF	Unemployment	50	< 0.001	TRUE	-0.03135	-0.03595	-0.02664	-1.88	80%
Low minimum wage	MF	High minimum wage	MF	Gini	50	< 0.001	TRUE	-0.02708	-0.03182	-0.0219	-1.48	72%
Low minimum wage	MF	High minimum wage	MF	Skill-wage premium	50	< 0.001	TRUE	0.02949	0.02509	0.03403	1.80	98%
Low minimum wage	MF	High minimum wage	MF	Unemployment	50	< 0.001	TRUE	-0.02925	-0.03427	-0.02386	-1.54	72%



Figure 38: Gini index over the horizon, Experiment 6 (minimum wage 0, 10, 14.71, 20). Mean over 50 seeds with a ± 1 standard-deviation band.




Figure 39: Skill-wage premium over the horizon, Experiment 6 (minimum wage 0, 10, 14.71, 20). Mean over 50 seeds with a ± 1 standard-deviation band.




Figure 40: Unemployment rate over the horizon, Experiment 6 (minimum wage 0, 10, 14.71, 20). Mean over 50 seeds with a ± 1 standard-deviation band.

B.9. Experiment 7

Table 23 reports the complete pairwise comparisons; Figures 41, 42 and 43 show the corresponding trajectories.

Table 23: All pairwise comparisons for experiment 7. Severance rate by label: no severance = 0, base = 1/3, high = 1.0.

A	Mode A	B	Mode B	Outcome	n	p	Sign.	Difference (A-B)	CI-low	CI-high	d_z	Sign. horizon
Base	MF	No severance	MF	Gini	50	< 0.001	TRUE	-0.01516	-0.02037	-0.01017	-0.82	35%
Base	MF	No severance	MF	Skill-wage premium	50	< 0.001	TRUE	-0.005694	-0.008141	-0.003357	-0.65	10%
Base	MF	No severance	MF	Unemployment	50	< 0.001	TRUE	-0.01582	-0.02114	-0.01056	-0.83	34%
Base	MF	High severance	MF	Gini	50	< 0.001	TRUE	-0.0001791	-0.004862	0.004493	-0.01	17%
Base	MF	High severance	MF	Skill-wage premium	50	< 0.001	TRUE	-0.0029	-0.005203	-0.0006029	-0.34	5%
Base	MF	High severance	MF	Unemployment	50	< 0.001	TRUE	0.0002505	-0.004656	0.00524	0.01	18%
No severance	MF	High severance	MF	Gini	50	< 0.001	TRUE	0.01498	0.009449	0.02083	0.72	30%

No severance	MF	High severance	MF	Skill-wage premium	50	< 0.001	TRUE	0.002793	2.368e-05	0.005577	0.27	7%
No severance	MF	High severance	MF	Unemployment	50	< 0.001	TRUE	0.01607	0.01035	0.02189	0.75	30%



Figure 41: Gini index over the horizon, Experiment 7 (severance rate 0, 1/3, 1.0). Mean over 50 seeds with a ± 1 standard-deviation band.




Figure 42: Skill-wage premium over the horizon, Experiment 7 (severance rate 0, 1/3, 1.0). Mean over 50 seeds with a ± 1 standard-deviation band.




Figure 43: Unemployment rate over the horizon, Experiment 7 (severance rate 0, 1/3, 1.0). Mean over 50 seeds with a ± 1 standard-deviation band.

B.10. Experiment 8

Table 24 reports the complete pairwise comparisons; Figures 44, 45 and 46 show the corresponding trajectories.

Table 24: All pairwise comparisons for experiment 8. Conversion probability by label: no conversion = 0, low = 0.25, base = 0.50, high = 0.75, complete = 1.0.

A	Mode A	B	Mode B	Outcome	n	p	Sign.	Difference (A-B)	CI-low	CI-high	d_z	Sign. horizon
Base	MF	Low conversion	MF	Gini	50	0.0044	TRUE	-0.008547	-0.01366	-0.003658	-0.47	3%
Base	MF	Low conversion	MF	Skill-wage premium	50	< 0.001	TRUE	-0.002622	-0.005213	-0.0001722	-0.28	4%
Base	MF	Low conversion	MF	Unemployment	50	0.0014	TRUE	-0.009018	-0.01432	-0.003805	-0.47	3%
Base	MF	No conversion	MF	Gini	50	< 0.001	TRUE	-0.01081	-0.0151	-0.006371	-0.67	24%
Base	MF	No conversion	MF	Skill-wage premium	50	< 0.001	TRUE	-0.0009077	-0.00297	0.001254	-0.12	8%
Base	MF	No conversion	MF	Unemployment	50	< 0.001	TRUE	-0.01176	-0.01621	-0.007007	-0.70	26%
Base	MF	High conversion	MF	Gini	50	0.0066	TRUE	0.002754	-0.0008244	0.006564	0.20	2%

Base	MF	High conversion	MF	Skill-wage premium	50	< 0.001	TRUE	-0.002345	-0.004457	-0.0001622	-0.30	4%
Base	MF	High conversion	MF	Unemployment	50	0.0050	TRUE	0.003386	-0.0002283	0.007346	0.24	4%
Base	MF	Complete conversion	MF	Gini	50	< 0.001	TRUE	0.003062	-0.002354	0.008198	0.16	12%
Base	MF	Complete conversion	MF	Skill-wage premium	50	< 0.001	TRUE	-0.003909	-0.00649	-0.001459	-0.42	9%
Base	MF	Complete conversion	MF	Unemployment	50	< 0.001	TRUE	0.003939	-0.001627	0.009255	0.20	17%
Low conversion	MF	No conversion	MF	Gini	50	< 0.001	TRUE	-0.002259	-0.007223	0.002539	-0.13	9%
Low conversion	MF	No conversion	MF	Skill-wage premium	50	< 0.001	TRUE	0.001714	-0.0005216	0.004143	0.20	5%
Low conversion	MF	No conversion	MF	Unemployment	50	< 0.001	TRUE	-0.002743	-0.007846	0.002567	-0.15	9%
Low conversion	MF	High conversion	MF	Gini	50	< 0.001	TRUE	0.0113	0.006505	0.01615	0.66	19%
Low conversion	MF	High conversion	MF	Skill-wage premium	50	< 0.001	TRUE	0.0002762	-0.002238	0.002835	0.03	6%
Low conversion	MF	High conversion	MF	Unemployment	50	< 0.001	TRUE	0.0124	0.007616	0.01722	0.70	21%
Low conversion	MF	Complete conversion	MF	Gini	50	< 0.001	TRUE	0.01161	0.005423	0.01781	0.52	26%
Low conversion	MF	Complete conversion	MF	Skill-wage premium	50	< 0.001	TRUE	-0.001287	-0.004529	0.001803	-0.11	8%
Low conversion	MF	Complete conversion	MF	Unemployment	50	< 0.001	TRUE	0.01296	0.006578	0.01942	0.56	27%
No conversion	MF	High conversion	MF	Gini	50	< 0.001	TRUE	0.01356	0.0102	0.01686	1.10	51%
No conversion	MF	High conversion	MF	Skill-wage premium	50	< 0.001	TRUE	-0.001438	-0.003415	0.0004384	-0.20	9%
No conversion	MF	High conversion	MF	Unemployment	50	< 0.001	TRUE	0.01515	0.01157	0.01859	1.18	56%
No conversion	MF	Complete conversion	MF	Gini	50	< 0.001	TRUE	0.01387	0.008504	0.01911	0.70	40%
No conversion	MF	Complete conversion	MF	Skill-wage premium	50	< 0.001	TRUE	-0.003001	-0.005919	-9.572e-05	-0.28	12%
No conversion	MF	Complete conversion	MF	Unemployment	50	< 0.001	TRUE	0.0157	0.01001	0.02118	0.77	46%
High conversion	MF	Complete conversion	MF	Gini	50	0.0014	TRUE	0.000308	-0.005305	0.005386	0.02	2%
High conversion	MF	Complete conversion	MF	Skill-wage premium	50	< 0.001	TRUE	-0.001563	-0.004537	0.001325	-0.14	5%
High conversion	MF	Complete conversion	MF	Unemployment	50	< 0.001	TRUE	0.0005536	-0.005043	0.006089	0.03	3%



Figure 44: Gini index over the horizon, Experiment 8 (transition probability 0, 0.25, 0.5, 0.75, 1.0). Mean over 50 seeds with a ± 1 standard-deviation band.




Figure 45: Skill-wage premium over the horizon, Experiment 8 (transition probability 0, 0.25, 0.5, 0.75, 1.0). Mean over 50 seeds with a ± 1 standard-deviation band.




Figure 46: Unemployment rate over the horizon, Experiment 8 (transition probability 0, 0.25, 0.5, 0.75, 1.0). Mean over 50 seeds with a ± 1 standard-deviation band.

B.11. Illustration of the permutation test (Experiment 4)

Below figures illustrate how the paired permutation test works visually. The upper panel shows the two configurations (average ± 1 sd), the panel below shows the paired difference (A-B) with the 95%-band, where the yellow marked steps are the significant steps (significant horizon) (Figures 47, 48 and 49).




Figure 47: Permutation test, decreasing versus flat AI-cost trajectory for unemployment rate. $p < 0.001$, significant on 99% of the horizon, $d_z = 3.47$




Figure 48: Permutation test, decreasing versus flat AI-cost trajectory for the skill-wage premium. $p < 0.001$, significant on 94% of the horizon, $d_z = 3.67$





Figure 49: Permutation test, decreasing versus rising AI-cost trajectory for the skill-wage premium. $p < 0.001$, significant on 78% of the horizon, $d_z = 2.02$



Parameter calibration

The model has roughly forty parameters. During development many were set heuristically, by adjusting them until the model produced plausible behaviour. For the analysis to be defensible, every parameter is given an explicit basis. This appendix documents how the baseline parameter set, the values in *base_parameters.py* against which every scenario in this thesis is compared, is chosen, and shows that the resulting economy is internally consistent.

C.1. Calibration strategy

Each parameter is assigned to one of four tiers, and the tier determines what kind of justification is required.

Mandate parameters are fixed by the research design and are not calibrated: the number of firms, and the two switches that define the institutional scenario under study (*employment_protection* and *ai_irreversible*).

Data-anchored parameters map directly onto an observable Dutch statistic or onto statutory law. Their values are taken from various sources. A few are data-informed rather than directly identified: the available statistic is only an analogue of the model quantity, so it disciplines the value without pinning it down exactly. These are labelled accordingly in the table.

Structural / literature-informed parameters define the shape and relative scale of the task-based production functions. They do not map directly onto a single observable Dutch statistic. Their values are therefore treated as structural modelling assumptions: they preserve comparability with the task-based model of Upreti and Sridhar (2024), while being interpreted as stylised parameters rather than as a direct replication target. Their plausibility is assessed through the model's qualitative validation patterns (Appendix E) and sensitivity analysis (Appendix G).

Pattern-calibrated parameters have no single empirical counterpart. They are chosen so that the model reproduces the qualitative direction of the six validation patterns (Appendix E).

C.2. Three structural decisions

Three decisions set the scales on which every other parameter is expressed; the rest of the calibration inherits from them.

C.2.1. Wage scale: euro-per-hour-equivalent units

Wages, the AI rental cost and the upfront investment cost are all denominated in euro-per-hour-equivalent units. The wage curve is scaled so that, at full employment, the high-skilled target wage is approximately EUR 35 per hour and the low-skilled target approximately EUR 21 per hour, close to the CBS gross hourly wages for HBO/WO-educated and lower-educated employees respectively. The statutory minimum hourly wage of EUR 14.71 is the floor. With these values the minimum wage binds for

low-skilled labour only when low-skilled employment is slack, and never for high-skilled labour, which mirrors the Dutch situation in which the minimum wage is binding almost exclusively at the lower end of the labour market.

C.2.2. AI cost scale

The AI rental cost is matched to the wage scale. The initial value ($k_{ai} = 40$) is chosen so that at the start of the simulation AI on a simple routine task is marginally more expensive than minimum-wage low-skilled labour; this leaves room for an adoption curve to develop as the cost falls. The cost decays towards a floor at 50% of its initial value ($k_{ai_floor} = 20$), with a half-life of roughly 35 months. That decline is deliberately slower than the fall in raw inference cost reported by the AI Index (Maslej et al., 2025): the model's AI cost represents the firm-side cost of adoption, integration, fine-tuning, governance and support, which does not collapse at the pace of per-token API prices.

C.2.3. Discount rate

Because one model step is one month, the NPV calculation needs a per-month discount rate. The value used is a monthly rate of 0.007, consistent with an annual weighted average cost of capital in the 8–9.5% range that KPMG (2025) reports for the technology sector.

Why this anchor. The discount rate represents the opportunity cost of capital of the firm making the automation investment. The published WACC of the technology sector is used as the closest available anchor for the risk profile of an AI investment specifically, rather than a economy-wide average WACC: the asset being priced is a technology asset, even when the adopting firm is not a technology firm. This is a calibration convention; a broader-market WACC would be somewhat lower and would make automation marginally more attractive.

What the value does in the model. The discount rate sets how heavily the future savings of automation weigh against the upfront cost. Over the 36-month planning horizon the discounted-savings (annuity) factor is $\sum_{t=1}^{36} (1+r)^{-t} = (1 - (1+r)^{-36})/r \approx 31.7$ at $r = 0.007$. A higher rate lowers the present value of the savings stream, so fewer tasks clear the adoption hurdle and automation is slower and lower; a lower rate does the opposite. Because the horizon is short, however, the decision is not very sensitive to the rate within a plausible band: halving the rate to 0.0035 raises the factor to about 33.8 and increasing it by half to 0.0105 lowers it to about 29.9, a swing of roughly $\pm 6\%$ around the baseline. The OFAT analysis (Appendix G) varies $r_discount$ over exactly this band; the adoption decision is governed far more strongly by the planning horizon $T_horizon$ and the investment-cost parameters than by the discount rate itself.

What the rate does not capture. The rate discounts a savings stream that the firm treats as certain: the mean-field expectation is evaluated without variance, so no risk premium for uncertainty about the AI-cost path or the technology's performance is priced into r itself. That caution is instead absorbed, crudely, by the separate hurdle rate $\eta = 0.15$ on the investment cost, consistent with evidence that firms apply hurdle rates well above their cost of capital (Meier & Tarhan, 2007). The implications of leaving uncertainty out of the discounting are discussed in section 6.4.

C.3. Steady-state

The absolute levels of price and output are not literal euro and quantity figures. Output is a dimensionless Leontief composite and wages are in euro-per-hour-equivalent units; the labour share is therefore a ratio rather than a direct counterpart of the CBS labour income share. The calibration tunes the autonomous-demand level A_base so that this ratio, and the operating margin, sit in an economically sensible region.

Demand is governed by the two-class consumption function $A = A_base + \gamma W + \gamma_pi Pi$, where W is the lagged wage bill and Pi the lagged aggregate operating profit, and the price clears the inverse-demand curve $p = A - B*Q$. A_base is set to 100 so that, at the no-AI warm start, firm revenue covers the wage bill at a positive operating margin.

C.4. The investment-cost magnitude (I_base)

The upfront investment cost deserves separate attention because an earlier value made the model degenerate. The NPV adoption rule automates a task only when its net present value exceeds the hurdle $\eta_{\text{hurdle}} * I$, with the task-level cost growing as $I(x) = I_{\text{base}} * (1 + 0.15 * x)$ over complexity x . At $I_{\text{base}} = 200$ the cost ran from 230 at the simplest task to 500 at the most complex, against discounted per-task savings of a few hundred over the 36-month horizon, the hurdle cleared only at low complexity, and adoption stalled at complexity 2 or below regardless of the rest of the calibration. At $I_{\text{base}} = 80$ the cost runs from 92 to 200, and the hurdle clears across a larger complexity range once the AI cost has decayed for a few years.

C.5. The baseline parameter set

Table 25 lists every parameter, its baseline value, its calibration tier, and the basis for the value. The component identifiers (C1-C7) refer to the conceptual components of chapter 3.

Table 25: Baseline parameter set: value, calibration tier and basis for each parameter. The component identifiers (C1–C7) refer to chapter 3.

Parameter	Value	Tier	Basis for the value
Scale and timing			
n_producers	20	Mandate	Number of firms. Fixed by the research design.
n_high_skilled	200	Data-anchored	High-skilled worker supply. With n_low_skilled this gives a 44% high-skilled share, the HBO/WO share of the Dutch active labour force (CBS, 2026).
n_low_skilled	250	Data-anchored	Low-skilled worker supply; the complementary share of roughly 56% (MBO or lower) (CBS, 2026).
seed	42	Technical	Fixes the random draw for reproducibility; replication experiments sweep a contiguous block of seeds.
steps_per_year	12	Data-anchored	Monthly time step. Tenure and the chain regulation are measured in months under the WAB and the WWZ.
Wage-setting (C3)			
a_h	15.0	Data-anchored	High-skilled wage intercept. With b_h, the full-employment target is about EUR 35/hour, matching the HBO/WO mean gross hourly wage (CBS, 2025e).
b_h	0.10	Data-anchored	High-skilled slope, set jointly with a_h.
a_l	6.0	Data-anchored	Low-skilled wage intercept. With b_l, the full-employment target is about EUR 21/hour, matching the CBS lower-educated mean (CBS, 2025e).
b_l	0.06	Data-anchored	Low-skilled slope, set jointly with a_l.
w_min	14.71	Data-anchored	Statutory minimum hourly wage (“Wet minimumloon en minimumvakantiebijslag”, 2026).
lam	0.08	Data-anchored	Wage partial-adjustment speed. Gives a wage half-life of about 8 months, consistent with the roughly 17-month average Dutch CAO duration (AWVN, 2025) and with evidence on Dutch wage stickiness (Calaia et al., 2023).
Production and productivity (C1)			

Parameter	Value	Tier	Basis for the value
a_prod	0.1	Structural / literature-informed	Human productivity scale in the exponential task-productivity function. The value is treated as a stylised scale parameter rather than an empirical estimate, it fixes the relative productivity gradient over task complexity in the dimensionless Leontief production system.
xi_prod	0.5	Structural / literature-informed	Cross-skill productivity dampening on non-routine tasks. The value represents a moderate productivity penalty for low-skilled workers on non-routine tasks; the direction follows the task-based literature on high-skilled comparative advantage in non-routine work, while the magnitude is stylised.
phi_base	3.0	Structural / literature-informed	Baseline AI productivity scale. The value sets the relative productivity of AI before complexity and task-type corrections; it is a structural scale parameter rather than an empirical estimate, and its economic effect depends jointly on wages, AI costs, investment costs and the NPV adoption mode.
phi_decay	0.04	Structural / literature-informed	Complexity penalty in AI productivity. The value imposes a gradual productivity decline as task complexity increases, consistent with the assumption that more complex tasks are less directly automatable.
nr_multiplier	0.85	Structural / literature-informed	Baseline multiplier for AI productivity on non-routine tasks. A value below one represents a conservative routine-biased baseline. Because generative AI may also affect non-routine cognitive work, this parameter is interpreted as scenario-defining.
AI learning curve (C5)			
k_ai	40.0	Pattern-calibrated	Initial AI rental cost, set so that at the start of the run AI on a simple routine task is marginally more expensive than minimum-wage low-skilled labour.
k_ai_floor	20.0	Pattern-calibrated	Long-run rental-cost floor, 50% of the initial cost. Inference cost falls steeply, but firm-side adoption cost (integration, governance, support) does not collapse to zero.
k_ai_decay	0.020	Pattern-calibrated	Cost-decline speed; a half-life of about 35 months, deliberately slower than raw inference cost (Maslej et al., 2025) because it represents firm-side adoption cost.
AI adoption and investment (C4); separation and matching (C2)			
p_evaluate	0.10	Pattern-calibrated	Per-step probability that a firm re-evaluates automation for an affected task. Reflects bounded rationality and organisational inertia (Hannan & Freeman, 1984; Simon, 1955).
delta	0.013	Data-informed	Monthly exogenous separation probability. The value is loosely informed by CBS evidence that 3.8% of employees changed employer in Q2 2025, corresponding to roughly 1.3% per month. Because job-to-job mobility is not identical to task-level separations in the model, the parameter is treated as data-informed rather than directly observed (CBS, 2025c).
mismatch_theta	1.2	Pattern-calibrated	Cross-skill substitution tolerance; keeps cross-skill hiring available in slack markets without erasing the skill distinction.
ai_irreversible	True	Mandate	Research-design switch; implements the real-options lock-in argument.

Parameter	Value	Tier	Basis for the value
I_base	80.0	Pattern-calibrated	Svanberg et al. (2024) show that economic AI adoption depends on technical feasibility together with fixed implementation costs, scale, and the relation between automation costs and labour costs. The value I_base = 80 is therefore treated as a pattern-calibrated implementation-cost parameter: it is chosen to keep adoption economically selective while avoiding the degenerate case in which only the simplest tasks are ever automated.
complexity_scaling	0.15	Pattern-calibrated	Investment grows with complexity as $I(x) = I_base * (1 + 0.15 * x)$, a moderate cost gradient.
r_discount	0.007	Data-anchored	Per-step discount rate, from an annual WACC in the 8–9.5% range that KPMG (2025) reports for the DACH technology sector, converted to a monthly rate and rounded off.
T_horizon	36	Pattern-calibrated	NPV planning horizon and irreversibility lock-in, aligned with the 36-month chain limit so investment and contract cycles are commensurate.
eta_hurdle	0.15	Pattern-calibrated	The value is motivated by evidence that firms often apply hurdle rates above their cost of capital, reflecting uncertainty and managerial caution (Meier & Tarhan, 2007).
omega_adapt	0.15	Pattern-calibrated	Smoothing weight for adaptive wage expectations. The use of adaptive updating is grounded in the adaptive-learning literature, while the numerical value is pattern-calibrated to avoid both overly volatile and overly inert wage expectations.
Goods market and demand (C6)			
A_base	100.0	Pattern-calibrated	Autonomous (income-independent) component of the demand intercept A. Calibrated as a small floor, or buffer: a minority share of A (~15% at initialisation) so that demand tracks factor income rather than being propped up by an outsized autonomous term. This lets the price start above average cost at a positive operating margin, with the labour share in a plausible range, while still allowing the Keynesian production gate to bind once automation erodes the wage bill.
B	2	Pattern-calibrated	Slope of the inverse-demand curve $p = \max(0.01, A - B \cdot Q)$. Calibrated so that price responds to output without making demand excessively elastic.
gamma	0.05	Pattern-calibrated	Weight on wage income in the Kaleckian two-class consumption function $A = A_base + \gamma \cdot W + \gamma_pi \cdot Pi$. It plays the role of the propensity to consume out of wage income, but is not a literal dimensionless MPC: it maps the aggregate wage bill (a level, in euros) into the demand intercept (a price). Calibrated so the wage-income channel dominates A (~82% at initialisation).
gamma_pi	0.008	Pattern-calibrated	Corresponding weight on profit income ($Pi = \sum_i (P - AC_i) \cdot Q_i$). $\gamma_pi < \gamma$ preserves the wage-led tilt of Dutch demand while closing the income identity $W + Pi = P \cdot Q$. It is held below the stability bound $\gamma_pi < 1/Q$ (≈ 0.035): because profit contains $P \cdot Q$, a larger value makes the price feedback on itself and the demand–price loop diverges.
Employment protection (C7)			
employment_protection	True	Mandate	Research-design switch; the Dutch institutional setting is the object of study.

Parameter	Value	Tier	Basis for the value
severance_rate	1/3	Data-anchored	Statutory transitievergoeding: one third of a monthly wage per year of service (“Wet Arbeidsmarkt in Balans”, 2021).
chain_limit	36	Data-anchored	Statutory ketenregeling cap of 36 months (“Wet Werk En Zekerheid”, 2020).
p_convert	0.5	Pattern-calibrated	Probability that a worker reaching the chain limit is converted to a permanent contract rather than not renewed; a neutral prior, as no public conversion rate exists.
init_share_vast	0.67	Data-anchored	Initial permanent-contract share; about 5.5 million permanent of 8.2 million employees (CBS, 2024b). Self-employment is outside the model, so the employees-only ratio is the right anchor.
init_tenure_max_years	10.0	Data-anchored	Upper bound of the initial tenure draw for permanent workers; gives a mean initial tenure of about 6.5 years, consistent with the Dutch job-tenure distribution.
Firm heterogeneity (C1)			
alpha_source	“uniform”	Pattern-calibrated	Each firm’s routine-task share is drawn from a uniform distribution; a data-driven alternative using ILO occupational automation scores (Gmyrek et al., 2023) is also implemented.
alpha_min	0.25	Pattern-calibrated	Lower bound of the uniform routine-share draw.
alpha_max	0.75	Pattern-calibrated	Upper bound of the uniform routine-share draw; the resulting cross-firm mean of 0.50 approximates the economy-wide share of routine-exposed tasks.

C.6. Known simplifications and caveats

Several simplifications should be kept in mind when reading the calibration. The model’s monetary quantities are internally consistent euro-per-hour-equivalent units, not literal national-accounts figures; the labour share is a diagnostic ratio and its absolute level should not be compared directly to CBS GDP components. The model represents employees only; self-employment and owners are outside its scope. Wage rigidity is captured by a single economy-wide adjustment speed, so sectoral dispersion in CAO durations is not represented.

D

Verification

D.1. Verification

Verification tests whether the model is built as intended; validation tests whether it represents the model does what it is meant to do.

The verification has three types of tests. The first are unit tests: unit tests are performed for the main functions (Leontief, NPV calculation, severance payment, wage curve, AI learning curve). Every test checks whether a predefined input to the function gives the same output as calculated by hand. The second type concerns conservation of counts at the model level: the total number of workers should remain the same during a run, as should the number of tasks per firm, and the expenses for AI and wages should sum to the total variable costs. The third group of tests concerns the model behaviour in edge cases: for free AI, free AI investment and weak employment protection the adoption should converge to 100%, for an extremely high AI price the adoption should remain zero, for *employment_protection* = *False* and *chain_limit* = ∞ the institutional effect disappears, and for *init_share_vast* = 1 the severance component for proactive automation is maximal (Table 26).

Table 26: Summary of verification tests

Type	Test	Check	Pass/Fail
Unit test	Leontief output	Bottleneck selection for varying inputs	Pass
Unit test	NPV-calculation	Reproducible NPV calculation for all modes	Pass
Unit test	Severance payment	Severance = $(1/3) \cdot w \cdot$ tenure years per permanent worker	Pass
Unit test	Wage curve	Half-life after shock	Pass
Unit test	AI learning curve	Convergence to <i>k_ai_floor</i>	Pass
Conservation	Worker count	employed + unemployed = constant for complete run	Pass
Conservation	Task count	sum of tasks remains the same over the complete run	Pass
Conservation	Variable costs	total variable costs equal the sum of wages and AI costs, per firm and at model level	Pass
Edge case	Free AI	Adoption = 100%	Pass
Edge case	Extreme AI price	Adoption = 0%	Pass
Edge case	No employment protection	Total severance = 0 and no contract conversions	Pass
Edge case	Complete permanent workforce	Sum of severance is higher for <i>init_share_vast</i> = 1 than for other values and all contracts are permanent.	Pass

D.2. Test specifications

The verification suite is organised into three blocks, each checking a different class of expected behaviour:

1. **Unit tests** for the core computational functions (Leontief, NPV, severance pay, wage curve, AI learning curve).
2. **Conservation checks** at model level (worker, task, and cost conservation).
3. **Edge cases** that test the model response under extreme parameters (cheap/expensive AI, no institutions, fully permanent workforce).

Each test passes when its prediction holds; a failure would indicate that the implementation deviates from the specification. All tests reported below pass. The complete, executable notebook, with the exact inputs, the hand-checked expected values, and the pass/fail output for every test described below, is provided in the verification notebook (`notebooks/verification/verification_tests.ipynb`).

D.3. Block 1 - Unit Tests

D.3.1. Test 1 - Leontief Output

Prediction: $y_j = \min_x \frac{u_x}{c_x}$.

Three analytically solvable input-output pairs are tested, plus the error paths (negative input, non-positive coefficient) and the empty-input convention.

D.3.2. Test 2 - NPV Calculation (All Four Modes)

Prediction:

$$NPV = -(I + S) + \sum_{t=1}^T \frac{\Delta C_t}{(1+r)^t}$$

with $\Delta C_t = uc_t^h - uc_t^{ai}$. The result of `Producer.calculate_npv` is analytically reproduced for `npv_naive`, `npv_adaptive`, and `npv_mean_field`. For `ulc` mode, the adoption decision is confirmed to be a direct unit-cost comparison (no NPV).

D.3.3. Test 3 - Severance Pay

Prediction: For each permanent worker,

$$S_w = \text{severance_rate} \cdot w \cdot \frac{\text{tenure}}{\text{steps_per_year}}$$

and task-level severance is the sum over permanent workers only. Flexible workers generate no severance, and when `employment_protection = False`, severance is identically zero.

D.3.4. Test 4 - Wage Curve and Partial Adjustment (Half-Life)

Prediction: Under partial adjustment $w_{t+1} = \lambda \cdot w^* + (1 - \lambda) \cdot w_t$, the deviation decays as $(1 - \lambda)^t$. The half-life after a shock is

$$t_{1/2} = \frac{\ln 2}{-\ln(1 - \lambda)}.$$

The test checks that (a) the discrete recursion crosses the analytical half-life point at $\lceil t_{1/2} \rceil$, and (b) `LabourMarketModel.update_wages` behaves identically after a shock to `w_h` (with employment held fixed).

D.3.5. Test 5 - AI Learning Curve (Exponential Decay)

Prediction: $k_{ai}(t) = k_{ai, \text{floor}} + (k_{ai,0} - k_{ai, \text{floor}}) \cdot e^{-\lambda_k t}$.

The test checks (a) that the model value of `k_ai` after an arbitrary number of steps exactly satisfies this formula, and (b) that the series converges to `k_ai_floor` over sufficiently long runs.

D.4. Block 2 - Conservation Checks

D.4.1. Test 6 - Worker Conservation

Prediction: $|\text{employed}| + |\text{unemployed}| = N$ throughout the run, and the composition by skill level remains constant (no workers are created or destroyed).

D.4.2. Test 7 - Task Conservation

Prediction: The number of tasks per producer (`n_tasks`, default 20) is a structural constant; automation changes only the assignment (`automated` flag), not the cardinality of `producer.tasks`.

D.4.3. Test 8 - Cost Identity

Prediction: total variable costs_{*j*} = $\sum w_i + \sum (n_{\text{ai}} \cdot k_{\text{ai}})$ per firm, and also in aggregate at model level.

D.5. Block 3 - Edge Cases

D.5.1. Test 9 - Trivially Cheap AI => Adoption to 100%

Prediction: With `k_ai_floor` -> 0 (and `k_ai` = `k_ai_floor`), AI becomes cheaper than human labour for every task; in all adoption modes, the automation share converges to 1. Here `p_evaluate=1.0` is set so that no task remains unevaluated.

D.5.2. Test 10 - Prohibitively Expensive AI => Adoption to 0%

Prediction: When `k_ai` = `k_ai_floor` -> infinity, the AI unit cost is always higher than the human unit cost and the NPV remains negative. No task is automated.

D.5.3. Test 11 - `employment_protection` = `False` => Institutional Effect = 0

Prediction: Without employment protection, the WAB component disappears from the model: severance is identically zero, chain-regulation conversions are identically zero, and non-renewals are identically zero. A high `chain_limit` (effectively infinity) confirms that the institutional effect is inactive even without the protection flag.

D.5.4. Test 12 - `init_share_vast` = 1 => Maximum Severance Component

Prediction: With all initially employed workers on permanent contracts, the severance component for proactive automation reaches its maximum. This is tested by comparing total severance at $t = 0$ (after warm-start) with the variant where `init_share_vast` = 0.

E

Validation

This appendix sets out the validation strategy for the agent-based model introduced in Appendix A. First, the choice of pattern-oriented validation is motivated. Next, the set of six testable patterns against which the model behaviour is tested is introduced. The experimental design of the validation runs is then described, after which the results per pattern are reported, ending with a synthesis.

E.1. Why pattern-oriented validation

An agent-based model with this level of complexity is hard to validate against a single aggregate time series. Doing so would risk overfitting on one indicator while other mechanisms continued to malfunction unnoticed. Moreover, a comparison with the industrial revolution or earlier waves of automation is not sufficient, because different (routine) tasks were affected and different institutional factors played a role.

Grimm et al. (2005) formulate the principle of pattern-oriented modelling as follows: “Patterns are observations of any kind showing non-random structure and therefore containing information on the mechanisms from which they emerge.”. A model that reproduces multiple heterogeneous patterns simultaneously across a wide range of parameters earns more trust than a model that is fitted to a single indicator. This approach is commonly used in the agent-based macroeconomic literature.

The validation criterion used in this research is the direction of a trend: for each pattern, it is checked whether the model results go qualitatively in the same direction as the empirical prediction. More strongly formulated criteria, such as statistically significant results, are not used, since the model is built to reproduce the mechanisms of the Dutch labour market rather than to resemble it one-to-one.

E.2. Testable patterns

The pattern set consists of six mechanisms (Table 27), selected because they (i) are broadly supported empirically, (ii) link directly to the main mechanism of AI adoption and institutional friction, and (iii) are measurable in the output and behaviour of the model. The set is bounded by the capabilities of the model as described in the conceptual design; the model is not adjusted to test for additional patterns.

Table 27: Overview of the six validation patterns

#	Pattern	Expected behaviour
1	Profitability gate	Loss-making firms should stop proactive automation
2	AI cost	More expensive AI is adopted less often
3	Severance	Higher firing costs should slow AI adoption
4	Demand decline	Job losses should reduce aggregate demand
5	Minimum wage	A higher wage floor pushes firms to automate
6	Polarisation	AI that targets routine tasks hurts low-skilled workers most

E.2.1. Pattern 1 - Profitability gate

If production suddenly becomes extremely expensive, firms should stop proactive expansion and automation. In the model, producers only generate proactive bottleneck-expansion requests when their average cost is below the goods-market price. Once average cost rises above price, further expansion would deepen losses, so proactive automation should stop. This pattern is tested through an exogenous profitability shock that sharply increases both wages and AI rental costs. The pattern passes if producers become loss-making, the profitability gate is activated, and no new automation occurs after the shock.

E.2.2. Pattern 2 - Slower adoption rate for a more expensive AI

For a higher k_{ai_floor} or a slower k_{ai_decay} the model should show later and less AI adoption. This is a direct robustness test on the adoption logic and links to the minimum-viable-scale intuition of Svanberg et al. (2024). The test uses two separate sweeps within the Dutch institutional scenario: k_{ai_floor} over {15, 20, 25} and k_{ai_decay} over {0.005, 0.020, 0.080}. For each level, the adoption at the last time step is recorded. The test passes if the final adoption decreases with k_{ai_floor} and rises with k_{ai_decay} .

E.2.3. Pattern 3 - High severance payment slows automation down

A higher severance payment raises the cost of automation and should slow down AI adoption. The pattern is tested by varying the severance-rate parameter within the Dutch institutional scenario over {0, 1/3, 1, 5, 20}, with $p_convert = 0.95$ and $chain_limit = 2$ to ensure that enough employees hold a permanent contract for severance to have an effect. The test passes if the final adoption is lower for higher values of the severance rate.

E.2.4. Pattern 4 - Demand declines due to mass unemployment

A decrease in the employment rate should lead to lower aggregate demand. The consumption function is $A_t = A_0 + \gamma Y_{t-1} + \gamma_\pi \Pi_{t-1}$, where the wage component is the channel that pattern 4 checks. When a substantial part of the workforce loses their job, the wage bill falls immediately, which should depress aggregate demand. This in turn translates through the goods market to the price level. This pattern passes if aggregate demand and price fall in response to an employment shock.

E.2.5. Pattern 5 - Induced technical change for a higher minimum wage

The hypothesis follows from Acemoglu's (2002) induced technical change and capital-deepening results from Vergeer and Kleinknecht (2010): rigidities that keep labour costs high and predictable steer investment towards labour-saving technologies. This is checked by examining the model behaviour for different minimum wages within the Dutch institutional scenario at six levels: {0, 14.71, 25, 28, 30, 32}, where €14.71 is the Dutch minimum wage. For each level the final adoption is recorded, and the correlation between the minimum wage and adoption is calculated.

E.2.6. Pattern 6 - Polarisation dependent on AIs comparative advantage

For a lower μ_{nr} (AI is better at routine tasks), Goos et al. (2014) expect wage polarisation: high-education wages rise while low-education wages fall in relative terms. For a high μ_{nr} (AI is better at non-routine tasks (Webb, 2019)), the pattern can compress. The pattern is tested by varying the $nr_multiplier$ over {0.1, 0.85, 3.0} (with $w_min = 0$ to switch off the effects of pattern 5) and measuring two signals: the skill-wage premium (w_h/w_l) at the end of the run and the routine share of the automated tasks. The pattern holds if at least one of the two signals is stronger for a low $nr_multiplier$ than for the high $nr_multiplier$.

E.3. Experimental design validation

E.3.1. Time horizon

Every simulation runs for 960 steps. With $steps_per_year = 12$ this is 80 years of simulation time, enough to observe the relevant adoption and distributional dynamics and to let the Keynesian feedback mechanism take effect.

E.3.2. Monte Carlo replications

Per parameter configuration, $N = 12$ replications are run, each with a different seed. For each replication the stochastic components (initial α sampling, exogenous separation, evaluation draws, matching order) are resampled. Aggregated outcomes are reported as the average over the replications with a 95% confidence interval.

E.3.3. Scenario

All pattern tests are performed in the Dutch institutional scenario: employment protection turned on, active severance payments, a chain limit, a minimum wage, an initial set of permanent workers and wage stickiness.

E.4. Results per pattern

Unless noted otherwise, all validation tests were carried out in the Dutch institutional setting, with a simulation horizon of 960 steps and 12 Monte Carlo replications per parameter configuration. The complete validation runs, including the per-pattern data and the figures reproduced below, are provided in the validation notebook (`notebooks/validation/validation_tests.ipynb`). Across these six tests, all patterns pass.

E.4.1. Pattern 1 - Profitability gate

Hypothesis. If production suddenly becomes extremely expensive and firms are pushed into a loss-making position, the profitability gate should prevent proactive expansion and automation.

Operationalisation. The model is run to step 240 with exogenous separations switched off and with `p_convert` set to 1, so the warm-start allocation remains occupied by labour or already-installed AI. At the shock step, both wages and AI rental costs are raised to 10,000. This forces average costs above the goods-market price and isolates the profitability gate from vacancy recovery.

Result. The shock works exactly as intended. In 100 percent of post-shock observations, average costs exceed the goods-market price, and in 100 percent of post-shock observations average producer profit is negative. The profitability gate blocks expansion on average 20 times per post-shock step, while the total number of new automations and the number of new proactive automations both remain zero throughout the post-shock window. The post-shock automation is empty. Pattern 1 passes (Figure 50).




Figure 50: Pattern 1 (profitability gate). After an exogenous profitability shock at step 240 raises wages and AI rental costs, average cost exceeds the goods-market price and proactive automation halts. Dutch institutional scenario, 12 Monte Carlo replications.

E.4.2. Pattern 2 - Slower adoption rate for a more expensive AI

Hypothesis. If AI remains more expensive, firms should adopt it more slowly and to a lesser extent.

Operationalisation. The model varies two cost parameters separately. First, the AI cost floor is swept over 15, 20, and 25. Second, the AI learning-curve decay parameter is swept over 0.005, 0.020, and 0.080. For each configuration, final AI adoption is recorded at the end of the simulation.

Result. The pattern is clearly reproduced. Final adoption declines monotonically as the AI cost floor

risers: mean final adoption equals 0.461 at a floor of 15, 0.405 at a floor of 20, and 0.341 at a floor of 25. Conversely, adoption rises when AI prices fall more quickly over time: mean final adoption is 0.386 at a decay rate of 0.005 and 0.403 at both 0.020 and 0.080, the latter two having stabilised at the same level. The second sweep therefore flattens at the upper end, but still moves in the predicted non-decreasing direction. The figure confirms that cheaper AI leads to more adoption and that a slower decline in AI costs dampens it. Pattern 2 passes (Figure 51).




Figure 51: Pattern 2 (AI cost vs. adoption). Final AI adoption falls as the AI cost floor rises (15, 20, 25) and rises as the cost-decay rate increases (0.005, 0.020, 0.080). Dutch institutional scenario, 12 replications.

E.4.3. Pattern 3 - High severance payment slows automation down

Hypothesis. Higher severance costs should make automation less attractive and therefore reduce final AI adoption.

Operationalisation. The severance-rate parameter is varied across five values: 0, 1/3, 1, 5, and 20. To make the severance channel visible, the test uses a short chain limit and a high conversion probability, so that sufficient workers hold permanent contracts and dismissals can generate severance obligations.

Result. The model shows the expected negative relationship. Mean final adoption falls from 0.398 when severance is zero to 0.392 at the baseline one-third rate, 0.387 at a rate of 1, 0.377 at a rate of 5, and 0.337 at a rate of 20. The effect is economically modest at lower severance levels, which is consistent with the model structure: a substantial share of automation takes place in empty slots, where no dismissal payment is due. Even so, the direction is unambiguous, and very high severance further suppresses adoption. Pattern 3 passes (Figure 52).

E.4.4. Pattern 4 - Keynesian demand decline on mass unemployment

Hypothesis. A sufficiently large employment shock should reduce aggregate demand through the wage-bill channel and should, through the goods market, put downward pressure on prices.

Operationalisation. At step 240, the model imposes an exogenous unemployment shock in which 35 percent of currently employed workers are randomly selected and forced into unemployment. The macro state is recorded immediately before the shock, immediately after the shock, and for 24 follow-up steps. The test compares the immediate change in employment, aggregate demand, price, and output.

Result. The Keynesian mechanism is present in the simulations. On average, 109.9 workers are forced into unemployment, with a range from 107 to 113 across replications. The immediate change in aggregate demand equals -99.607 on average, the immediate change in the price level equals -83.582, and total output falls by -32.049. The figure shows that a large employment shock is associated with a strong decline in demand and prices. This is exactly the pattern expected from the wage-demand channel. Pattern 4 passes (Figure 53).

E.4.5. Pattern 5 - Induced technical change for a higher minimum wage

Hypothesis. A higher minimum wage raises the labour-cost floor and therefore strengthens the incentive to substitute labour with AI.




Figure 52: Pattern 3 (severance vs. adoption). Final AI adoption against the severance rate (0, 1/3, 1, 5, 20); higher severance lowers adoption. Dutch institutional scenario with a short chain limit and high conversion probability, 12 replications.




Figure 53: Pattern 4 (Keynesian demand). A 35% exogenous employment shock at step 240 lowers aggregate demand and the price level. Dutch institutional scenario, 12 replications.

Operationalisation. The minimum wage is varied over six levels: 0, 14.71, 25, 28, 30, and 32. For each level, final AI adoption is measured and the rank correlation between the minimum wage and final adoption is calculated.

Result. The relationship is strongly positive. The Spearman correlation between the minimum wage and final AI adoption is +0.97, with a p-value of 0.0012. The figure shows that adoption rises as the minimum wage is increased, with the Dutch minimum wage of €14.71 located at the start of this upward-sloping pattern. This finding is consistent with induced technical change: a higher and more binding wage floor strengthens the incentive to automate human-owned tasks. Pattern 5 passes (Figure 54).




Figure 54: Pattern 5 (induced technical change). Final AI adoption against the minimum wage (0, 14.71, 25, 28, 30, 32); adoption rises with the wage floor (Spearman $\rho = +0.97$). Dutch institutional scenario, 12 replications.

E.4.6. Pattern 6 - Polarisation dependent on AIs comparative advantage

Hypothesis. When AI has a stronger comparative advantage in routine tasks, low-skilled workers should be hit harder. This should show up either in a wider skill-wage premium or in a higher routine share of automated tasks.

Operationalisation. The non-routine productivity multiplier is varied over three values: 0.1, 0.85, and 3.0. The test records two indicators at the end of the run: the skill-wage premium and the share of automated tasks that are routine. The minimum wage is set to zero in this test so that the wage-floor mechanism from Pattern 5 does not dominate the result.

Result. The composition effect is very clear, and the wage-premium effect moves in the expected direction when comparing the routine-biased and non-routine-biased extremes. When the multiplier equals 0.1, the skill-wage premium is 1.306 and the routine share of automation is 1.000. At the value of 0.85, the premium is 1.356 and the routine share falls to 0.467. When AI favours non-routine tasks with a multiplier of 3.0, the premium falls to 1.169 and the routine share falls further to 0.319. The premium is therefore not perfectly monotonic across all three values, but the comparison between the two extremes still supports the expected direction, and the task-composition signal is especially strong. Pattern 6 passes (Figure 55).

E.5. Synthesis

In summary, the validation exercise is successful. Each pattern tests something different. Patterns 1, 3 and 5 follow fairly directly from the model equations (the profitability gate, the severance term




Figure 55: Pattern 6 (polarisation). Skill-wage premium and routine share of automated tasks for non-routine multipliers 0.1, 0.85 and 3.0. Dutch institutional scenario with $w_{\min} = 0$, 12 replications.


in the NPV, and the cost comparison): passing them confirms that the implementation behaves as specified, which is a inbetween to verification than to independent validation. Patterns 4 and 6 are more genuinely emergent. The exercise therefore establishes internal consistency across all six patterns and behavioural plausibility for the emergent subset, rather than a one-to-one empirical fit. All six patterns and the verification tests pass: the model reproduces a functioning profitability gate, greater AI adoption when AI becomes cheaper, lower adoption under stronger severance obligations, a Keynesian decline in demand after a mass-unemployment shock, more automation under a higher minimum wage, and distributional differences that depend on AI comparative advantage across task types. This does not imply that the model is a literal one-to-one representation of the Dutch labour market. However, it does show that the core mechanisms introduced in Appendix A behave in the empirically expected direction and interact coherently under the Dutch institutional scenario. In that sense, the pattern-oriented validation provides sufficient support to use the model for the policy and scenario analysis in the remainder of the thesis.

Table 28 summarises the outcome of the six pattern tests, together with the quantitative evidence behind each verdict.

Table 28: Validation results: pass/fail and supporting evidence per pattern

#	Pattern	Expected direction	Result	Evidence
1	Profitability gate	Forced $AC > p$ stops proactive automation	Pass	$AC > p$ in 100% of post-shock steps; profit < 0 in 100%; gate blocks ≈ 20 per step; new proactive automations = 0
2	AI cost vs. adoption	Cheaper AI yields more adoption	Pass	Final adoption monotone in the AI cost floor and non-decreasing in the decay rate
3	Severance vs. adoption	Higher severance does not raise adoption	Pass	Final adoption 0.398 (no severance) vs. 0.377 (high severance)
4	Demand decline	Forced unemployment lowers A and the price	Pass	$\Delta L \approx 110$ workers; $\Delta A = -99.6$; $\Delta p = -83.6$
5	Minimum wage vs. adoption	Positive rank correlation ($\rho > 0$)	Pass	Spearman $\rho = +0.97$ ($p = 0.0012$)
6	Polarisation	Low μ_{nr} widens the premium or raises the routine share	Pass	Both signals move as expected: the premium widens and the routine share of automation rises

All six patterns pass (6/6), each based on $N = 12$ Monte Carlo replications over 960 steps per parameter configuration.



Parameters

The table shows all the model parameters with the symbol, variable name (in the model), and a description (Table 29).

Table 29: Model parameters overview

Parameter	Symbol	Variable name	Description
Population			
Number of producers	J	<i>n_producers</i>	Number of producers in the model
Number of high-skilled workers	N_h	<i>n_high_skilled</i>	Number of high-skilled workers in the model
Number of low-skilled workers	N_l	<i>n_low_skilled</i>	Number of low-skilled workers in the model
seed	—	<i>seed</i>	Seed for random reproducibility
Productivity			
Human producing capacity	a	<i>a_prod</i>	Steepness of the human productivity over task complexity
Comparative advantage scaling	ξ	<i>xi_prod</i>	Cross-skilling productivity scaling
AI base productivity	ϕ_0	<i>phi_base</i>	Baseline AI productivity level, before the complexity penalty and the non-routine multiplier are applied
AI complexity decay	ϕ_1	<i>phi_decay</i>	How fast AI productivity falls with complexity
NR multiplier	μ_{nr}	<i>nr_multiplier</i>	AI multiplier on non-routine tasks
Wage setting			
Intercept high	a_h	<i>a_h</i>	Reservation wage for high-skilled
Slope high	b_h	<i>b_h</i>	Wage responsiveness to the number of high-skilled workers
Intercept low	a_l	<i>a_l</i>	Reservation wage for low-skilled
Slope low	b_l	<i>b_l</i>	Wage responsiveness to the number of low-skilled workers
Minimum wage	$w_{\min}$	<i>w_min</i>	The minimum wage for both skill levels
Adjustment speed	λ	<i>lam</i>	The speed of the wage adjustment
Goods			
Autonomous demand	A_0	<i>A_base</i>	The exogenous demand (e.g. government, export)

Parameter	Symbol	Variable name	Description
Marginal propensity of income to consume	γ	<i>gamma</i>	Strength of the Keynesian demand feedback of the wage bill
Marginal propensity of profit to consume	γ_π	<i>gamma_pi</i>	Strength of the Keynesian demand feedback of total profit
Demand slope	B	<i>B</i>	Slope of the demand curve
AI adoption			
Initial AI costs	$k_{AI,0}$	<i>k_ai</i>	The starting cost of an AI unit at $t = 0$
AI cost limit	$k_{\min}$	<i>k_ai_floor</i>	The floor/ceiling of AI costs
Speed of learning curve	λ_k	<i>k_ai_decay</i>	The decay speed of the AI curve
Evaluation probability	p_{eval}	<i>p_evaluate</i>	The probability of re-evaluating bottlenecks or separated workers
Separation probability	δ	<i>delta</i>	Probability a worker leaves a firm
Mismatch tolerance	θ	<i>mismatch_theta</i>	Tolerance for cross-skill matches.
AI irreversibility	—	<i>ai_irreversible</i>	Whether an investment must ‘sit out’ the planning horizon before its task can be re-evaluated. (True/False)
Alpha minimum	$\alpha_{\min}$	<i>alpha_min</i>	Lower limit for the α distribution.
Alpha maximum	$\alpha_{\max}$	<i>alpha_max</i>	Upper limit for the α distribution.
NPV parameters			
Investment costs AI	I_0	<i>I_base</i>	The investment cost to start using AI.
Complexity scaling	κ	<i>complexity_scaling</i>	Cost multiplier for AI investment for complexity level
Discount rate	r	<i>r_discount</i>	The discount rate per step (monthly) for the NPV calculation
Planning horizon	T	<i>T_horizon</i>	NPV forward looking periods and AI irreversibility time,
Hurdle rate	η	<i>eta_hurdle</i>	Extra hurdle for automation investment.
EMA weight	ω	<i>omega_adapt</i>	Update speed for expectations and smoothing trends.
Employment protection			
Employment protection	—	<i>employment_protection</i>	Master switch to turn employment protection on or off.
Model steps per year	—	<i>steps_per_year</i>	Timescale parameter, (default = 12, one step = one month)
Severance rate	—	<i>severance_rate</i>	Monthly-based severance payment per tenure year (default = 1/3).
Chain limit	—	<i>chain_limit</i>	Number of steps after which a flex worker is converted or laid off (default = 36).
Conversion probability	—	<i>p_convert</i>	Probability that a flex worker is converted to a permanent contract upon reaching the chain limit.
Initial permanent share	—	<i>init_share_vast</i>	Fraction of permanent workers at the start of a simulation.
Maximum initial tenure years	—	<i>init_tenure_max_years</i>	Maximum tenure years for permanent workers at warm start.



Sensitivity analysis

This appendix reports the one-factor-at-a-time (OFAT) sensitivity analysis of the model. Whereas the scenario experiments in chapter 5 test a limited set of policy and scenario axes around the baseline, this analysis systematically shows how sensitive the model outcomes are to each of the forty parameters separately. The goal aligns with the exploratory setup: rather than seeking a point estimate, it identifies which assumptions influence the outcomes most and which do not. The analysis uses the baseline configuration of section 4.2.

G.1. Setup and method

In an OFAT analysis, every parameter is varied one at a time while the others are held constant at their baseline value (Appendix C). Each numerical parameter takes five values: the baseline, two steps up and two steps down, at $\pm 25\%$ and $\pm 50\%$ of the baseline value. For each parameter with a natural limit, the values are kept within admissible bounds (probabilities in $[0, 1]$; integer counts). For a few parameters the values depart from these percentage intervals. The non-routine multiplier (`nr_multiplier`) follows the scenario range $\{0.50; 0.75; 1.00; 1.25; 1.50\}$, because this parameter is scenario-determining and the baseline (0.85) has no natural midpoint. For `alpha_min`, a grid of $\{0; 0.125; 0.25; 0.375; 0.50\}$ is used. The two toggles `employment_protection` and `ai_irreversible` are Boolean and are turned on or off rather than scaled. The complete overview, with baseline values and tested ranges, is given in Table 30.

Each configuration is repeated over the same set of seeds and run for 1000 steps. For each parameter, the replication-averaged value at the last time step is taken and compared to the baseline. This horizon is 1000 rather than the 960 steps used in the scenario experiments of chapter 5: the sensitivity analysis was run earlier in the project, before the experiment horizon was fixed, and given the long runtime of the full sweep it was not repeated. The forty-step difference is immaterial, because the AI rental cost has reached its floor and all outcomes have stabilised well before step 960, so the endpoint is the same steady state in either case.

The influence of a parameter is summarised by the maximum absolute percentage deviation from the baseline across all tested values. For each parameter, the value that differs most from the baseline is identified and expressed as a percentage. This measure is chosen because it makes parameters with different scales and units comparable and captures the largest effect within the tested set. The parameters are ordered from most to least influential in the tornado diagrams (top 20 variables) below, where the bars show the percentage deviation.

The sensitivity is reported for four outcomes: the Gini index, the skill-wage premium (w_h/w_l), and the employment rates of high-skilled and low-skilled workers respectively. The overall employment rate, the primary outcome in chapter 5, is not shown separately because it follows directly from the two group-specific rates; together, the four measures thus cover the primary outcomes of the study.

When interpreting the results, it is important to distinguish between structural economic parameters

and scale parameters. For example, the number of agents in the model ($n_producers$, $n_low_skilled$, $n_high_skilled$) partly determines the absolute scale and market relations of the simulated economy and is calibrated to produce certain patterns; their strong influence therefore partly reflects this model scale and should not be interpreted unconditionally as a structural economic effect.

G.2. Gini-index

Figure 56 ranks the parameters by their maximum effect on the Gini index.




Figure 56: OFAT drivers of the Gini index of worker income, ranked by maximum absolute deviation from baseline (%).

The Gini index is the most sensitive of the four outcomes and is dominated by parameters that influence the automation level and thus unemployment. The number of producers has the largest effect: fewer firms mean a smaller labour market and larger swings in employment. Economically more interesting are the AI parameters: the baseline AI productivity, the long-run cost floor, the investment cost and the complexity scaling together determine how many tasks are profitable to automate. A higher productivity or lower cost floor makes automation more attractive, raises unemployment and thus the Gini index (the unemployed have an income of 0).

G.3. Skill-wage premium (w_h/w_l)

Figure 57 ranks the parameters by their maximum effect on the skill-wage premium.

The skill-wage premium reacts most strongly to the number of low-skilled workers, the cross-skill productivity scaling and the wage parameters of both groups. Halving the number of low-skilled workers raises the premium by ~19%, which is largely a scale effect of the wage curve: a smaller supply of low-skilled workers pushes their wage level down relative to w_h . Economically more meaningful is the cross-skill productivity scaling (ξ_{i_prod}): a higher value makes low-skilled workers more productive on non-routine tasks, which raises their wages and lowers the skill-wage premium. The wage intercepts and slopes work in the expected direction. The minimum wage floor w_{min} compresses the premium once it becomes binding for the low-skilled.

G.4. Employment rate for high skilled

Figure 58 ranks the parameters by their maximum effect on the high-skill employment rate.

The high-skilled employment rate is determined most strongly by the number of high-skilled workers themselves, by their wage parameters and by the number of producers. The AI parameters ϕ_{i_base} , the AI cost floor and the productivity steepness (a_prod) work through the automation decision: if AI becomes more productive or cheaper, the employment rate decreases.




Figure 57: OFAT drivers of the skill-wage premium (w_h/w_l), ranked by maximum absolute deviation from baseline (%).




Figure 58: OFAT drivers of the high-skill employment rate, ranked by maximum absolute deviation from baseline (%).

G.5. Employment rate for low skilled

Figure 59 ranks the parameters by their maximum effect on the low-skill employment rate.




Figure 59: OFAT drivers of the low-skill employment rate, ranked by maximum absolute deviation from baseline (%).

For low-skilled workers, the structural parameters (e.g. `n_producers`, `n_low_skilled`) behave similarly to the high-skilled case. A higher minimum wage floor leads to more unemployment because it makes automating routine tasks more profitable. The AI parameters `phi_base`, the AI cost floor, the investment cost and the complexity scaling work in a similar direction. The contrast between high- and low-skilled, where the minimum wage affects the low-skilled almost exclusively, aligns with the Dutch situation in which the minimum wage ties people to the bottom of the labour market.

G.6. Limitations of the analysis

The analysis is local and one-at-a-time: each parameter is varied separately around a single baseline. Interaction effects between parameters are therefore not visible, and the results hold only in the neighbourhood of the baseline. A global variance-based test (such as Sobol) would reveal interactions and non-linear effects across the full parameter grid, but falls outside the scope here because of its computational cost. In addition, the scale parameters (number of agents) combine a model-technical scaling effect with an economic effect, so their ranking positions must be read with that caveat in mind. The analysis is therefore intended as an exploratory check showing which assumptions dominate the outcomes, in line with the scenario-exploring design of the study.


G.7. Complete parameter overview

Table 30 shows all forty parameters of the sensitivity analysis with their baseline value and tested domain.

Table 30: Parameters, baseline values and tested ranges. Each numeric parameter is evaluated at five points: the baseline and $\pm 25\%$ and $\pm 50\%$ around it, clipped to the parameter's admissible domain (probabilities to $[0, 1]$; counts to integers ≥ 1). * `nr_multiplier` uses the scenario range of subsection 4.4.2. † `alpha_min` uses a fixed grid $\{0; 0.125; 0.25; 0.375; 0.50\}$. The switches `employment_protection` and `ai_irreversible` are toggled off/on.

Parameter	Description	Baseline	Tested range
<code>n_producers</code>	Number of firms	20	10 – 30
<code>n_high_skilled</code>	High-skilled worker supply	200	100 – 300
<code>n_low_skilled</code>	Low-skilled worker supply	250	125 – 375
<code>steps_per_year</code>	Time steps per year	12	6 – 18

Parameter	Description	Baseline	Tested range
a_h	High-skilled wage intercept	15	7.5 – 22.5
b_h	High-skilled wage slope	0.1	0.05 – 0.15
a_l	Low-skilled wage intercept	6	3 – 9
b_l	Low-skilled wage slope	0.06	0.03 – 0.09
w_min	Statutory minimum wage	14.71	7.355 – 22.065
lam	Wage adjustment speed	0.08	0.04 – 0.12
a_prod	Human productivity steepness	0.1	0.05 – 0.15
xi_prod	Cross-skill productivity scaling	0.5	0.25 – 0.75
phi_base	Baseline AI productivity	3	1.5 – 4.5
phi_decay	AI productivity decay with complexity	0.04	0.02 – 0.06
nr_multiplier	AI multiplier on non-routine tasks	0.85	0.5 – 1.5 *
A_base	Autonomous demand	100	50 – 150
B	Inverse-demand slope	2	1 – 3
gamma	Wage-income demand weight	0.05	0.025 – 0.075
gamma_pi	Profit-income demand weight	0.008	0.004 – 0.012
k_ai	Initial AI rental cost	40	20 – 60
k_ai_floor	Long-run AI cost floor	20	10 – 30
k_ai_decay	AI learning-curve speed	0.02	0.01 – 0.03
p_evaluate	Re-evaluation probability	0.1	0.05 – 0.15
delta	Job separation probability	0.013	0.0065 – 0.0195
mismatch_theta	Cross-skill matching tolerance	1.2	1 – 1.8
alpha_min	Lower bound routine-share draw	0.25	0 – 0.5 †
alpha_max	Upper bound routine-share draw	0.75	0.375 – 1
l_base	Fixed AI investment cost	80	40 – 120
complexity_scaling	Cost growth with complexity	0.15	0.075 – 0.225
r_discount	NPV discount rate	0.007	0.0035 – 0.0105
T_horizon	NPV planning horizon	36	18 – 54
eta_hurdle	Investment hurdle rate	0.15	0.075 – 0.225
omega_adapt	Expectation smoothing weight	0.15	0.075 – 0.225
ai_irreversible	AI investment irreversible (switch)	True	off / on
employment_protection	Employment protection (switch)	True	off / on
severance_rate	Severance per tenure year	1/3	0.166667 – 0.5
chain_limit	Chain limit (ketenregeling)	36	18 – 54
p_convert	Flex-to-permanent conversion prob.	0.5	0.25 – 0.75
init_share_vast	Initial permanent-contract share	0.67	0.335 – 1
init_tenure_max_years	Max initial tenure (years)	10	5 – 15



Model assumptions

This appendix gives an overview of the assumptions made for the agent-based model. The assumptions are grouped per model component (C1-C7) plus the macro-feedback layer (M1), preceded by model-wide and scope assumptions (A0). Every assumption is given a type label and its most important implication or limitation.

Concrete parameter values and their calibration are not repeated here: these are explained in Appendix F (Parameters) and Appendix C (Parameter calibration). This appendix describes only the assumptions about the model structure and the chosen functions. At the end is a short summary of the most influential assumptions, with the technical assumptions placed in a separate subsection to keep the main list readable.

H.1. Type labels

Structural (S): the architecture of the model (which entities exist and how they relate to each other).

Behaviour (B): the decision or behavioural rules of an agent.

Function (F): the chosen mathematical function or parameterisation of a relationship.

Demarcation (D): something deliberately left outside the model (scope).

Technical (T): an implementation or numerical decision without independent meaning.

H.2. A0 - Model-wide and scope assumptions

Table 31 lists the assumptions that apply to the model as a whole and set the boundaries of the system.

Table 31: Assumptions that affect the whole model and set the boundaries of the system. They do not belong to a single component.

ID	Assumption	Type	Implication/Limitation
A0.1	Closed economy: no foreign sector, imports or exports.	D	AI effects through international trade or offshoring are not taken into account.
A0.2	One homogeneous product, one sector, one market. All firms produce the same good at one price.	S	No sectoral heterogeneity, whereas the Dutch labour market is highly sectoral (e.g. hospitality vs. industry), and the AI effect could differ strongly across sectors.
A0.3	No government beyond institutions: no taxes, social benefits or government expenditure. The government acts only through the minimum wage and employment protection.	D	No fiscal redistribution or stabilisers that dampen inequality, so the measured inequality is overstated.

A0.4	No money, credit or financial sector. Investments are not financed; the only intertemporal influence is the discount rate.	D	No finance or interest rate influences the adoption decision.
A0.5	No physical capital stock. Capital consists only of rented AI units, without depreciation or investment dynamics independent of AI.	S	Replacement investment is absent and owning AI is not an option.
A0.6	Fixed number of firms and workers. No entry of new firms or bankruptcy, and no demographics (no young entrants or retirements).	S	Long-term structural change is underestimated, and tenure can become too high in the absence of retirements.
A0.7	Constant, fully participating workforce. The unemployed remain in the pool without discouragement or outflow.	D	Participation effects of AI are out of view.
A0.8	Two exogenous, fixed skill groups (high/low); no schooling, reschooling or mobility between skill groups.	S	The reinstatement and upskilling channel is missing, so inequality effects are presumably an upper bound.
A0.9	Discrete time, one step = one month. All dynamics occur in equal monthly steps.	F	All effects within a month are summarised into a single step.
A0.10	No spatial or regional dimensions. One national labour market without regions.	D	Regional labour market differences are not considered.
A0.11	The system need not reach an equilibrium; outcomes are trajectories, not equilibria.	S	The model is scenario-exploratory, not a point estimate.
A0.12	A task is performed either entirely by humans or entirely by AI. Within a single task there is no collaboration between human and AI. Automation is an all-or-nothing substitution per task; there is no augmentation in which AI increases a worker's productivity without replacing them.	D	Excludes the augmentation or complementarity channel: in the model, AI can only replace labour, not supplement it. Consequently, the model captures only displacement effects and not the wage or productivity gains that human-AI collaboration can yield. This is a fundamental limitation, precisely because a large part of current generative AI (LLMs acting as assistants) is deployed in an augmenting manner in practice.
A0.13	Pattern-calibrated baseline: the baseline parameters (Appendix C) are calibrated to reproduce qualitatively plausible and "interesting" patterns (a Keynesian expansion gate that closes, plausible levels of inequality and unemployment) rather than being statistically estimated on data.	D/B	All results are conditional on this specific baseline; the experiments and the OFAT analysis vary parameters around it, and the calibration itself may shape the findings (for example that firms operate close to their profitability frontier and the gate closes regularly).

H.3. C1 - Production

Table 32 lists the assumptions behind the Leontief task technology and the productivity of workers and AI.

Table 32: Leontief task technology and productivity: how tasks, workers and AI generate output.

ID	Assumption	Type	Implication/Limitation
C1.1	Leontief technology: tasks are perfect complements, output = minimum over tasks. No substitution between tasks.	S	One unfulfilled task blocks the entire output: a strong bottleneck.

C1.2	Equal input coefficients; every task adds an equal amount to the final output.	F	No task weighs more heavily than another; the relative importance of tasks is ignored.
C1.3	Equal and fixed number of tasks per firm.	S	No heterogeneity in firm size or number of tasks.
C1.4	The routine share (alpha) determines the task composition and is the only source of firm heterogeneity.	S	In reality, firms differ in more than the routine/non-routine task distribution.
C1.5	Routine versus non-routine is binary and does not change over time.	S	No redefinition and no creation of tasks.
C1.6	Complexity runs from 1 to n for each task type.	F	No randomness or overlap in the complexity assignment.
C1.7	Human labour is exponential in task complexity (as in Upreti & Sridhar).	F	Highly qualified, complex work is more productive; this sets the wage-cost relationship.
C1.8	Skill differentiation applies only to non-routine tasks: routine tasks are skill-neutral, high-skilled workers have an advantage on non-routine tasks, and low-skilled workers are dampened (comparative advantage).	F	Low-skilled workers have no productivity advantage; the skill difference exists only for non-routine tasks. This is a coarse split: real high-skilled routine and low-skilled non-routine tasks do not coincide with this single axis.
C1.9	AI productivity decreases exponentially in complexity, with a separate multiplier for non-routine tasks.	F	Creates a natural automation curve in which simple routine tasks are automated first.
C1.10	Productivity is additive over workers within a task: no team effects or diminishing returns.	F	Collaboration and coordination are ignored.
C1.11	AI is infinitely scalable.	S	There is no limit or scarcity for AI.
C1.12	AI capacity is trimmed each step to the bottleneck minimum for minimal cost. No overcapacity.	B	Firms hold no buffer or reserve capacity and keep all subscription levels up to date.
C1.13	Human productivity is constant over time, while AI becomes cheaper.	S	A structural asymmetry biased towards AI; human learning-by-doing is missing.
C1.14	Labour is indivisible (no part-time work).	S	All workers work full-time.

H.4. C2 - Labour market

Table 33 lists the assumptions that govern hiring, matching, separation and vacancy handling.

Table 33: Hiring, matching, separation and vacancy handling. The rules of the labour market.

ID	Assumption	Type	Implication/Limitation
C2.1	Workers are passive: all decisions are made by firms. Employees do not negotiate or search for jobs.	S	Wage and employment dynamics are entirely demand-driven.

C2.2	Exogenous separation probability: every worker loses their job with a fixed probability per step, independent of wage, age, tenure or role.	B	Job destruction is not linked to economic conditions or individual characteristics.
C2.3	Central, costless matching: no search frictions other than the availability of the unemployed. Vacancies are filled within one step or not at all.	S	No long-term unemployment.
C2.4	Limited cross-skill substitution: the preferred skill first, then a fallback if the mismatch cost remains below a tolerance.	B	Skill differences are preserved but are not absolute.
C2.5	Hiring is driven by the bottleneck and governed by an evaluation probability. No continuous filling of vacancies.	B	Employment responds to production conditions rather than being maximised automatically.
C2.6	One full-time job per worker; no part-time work or multiple employers.	S	The Dutch part-time system is absent.
C2.7	Dismissal means zero income immediately, without a waiting period or benefits.	B	The income trap from unemployment is maximal; there is no safety net in the income distribution, so the Gini-index effects could be over- or underestimated.

H.5. C3 - Wage setting

Table 34 lists the assumptions behind the wage-supply curve per skill group and its partial adjustment.

Table 34: A wage-supply curve per skill group with partial adjustments.

ID	Assumption	Type	Implication/Limitation
C3.1	Upward-sloping supply curve per skill group. The wage rises linearly with the employment rate of that group.	F	Wages rise only through scarcity of that skill group.
C3.2	Partial wage adjustment: wages move slowly towards their target (wage rigidity as a CLA-like institution).	B	Wages respond with a delay, so shocks work through gradually.
C3.3	One country-wide wage per skill group. No firm-, sector- or task-specific wages and no individual negotiation.	S	Wage differences within a skill group remain out of view.
C3.4	A minimum wage floor for both groups.	B	The lowest wage cannot fall below the minimum wage; the law is never violated.
C3.5	The wage depends on the employment rate, not on firm profit or productivity; there is no profit-sharing.	S	No direct pass-through of productivity gains to wages.
C3.6	CLA negotiations are only implicit through rigidity and always move towards the target.	S	CLA negotiations cannot turn out 'bad' or 'good' and are heavily simplified.

H.6. C4 - AI adoption

Table 35 lists the assumptions behind the investment decision: unit costs and the NPV in the four adoption modes.

Table 35: The investment decision: unit costs and NPV in four adoption modes. This is the core contribution to Upreti & Sridhar (2024).

ID	Assumption	Type	Implication/Limitation
C4.1	Adoption is an explicit cost-benefit decision per task.	B	Unlike in Upreti & Sridhar, adoption is an economic decision rather than a lottery.
C4.2	The NPV modes use a hurdle: automate if the NPV exceeds the hurdle, looking forward to a fixed horizon at a constant discount rate.	B	Adoption requires a return margin; the horizon and discount rate drive the timing.
C4.3	Investment costs rise linearly with task complexity.	F	More complex tasks are more expensive.
C4.4	Firms anticipate the future AI cost path exactly.	B	Fully rational expectations; there is no uncertainty about AI costs.
C4.5	Firms reassess their adoption only with a fixed probability per step (bounded rationality).	B	Adoption moves gradually rather than in shocks.
C4.6	AI (ir)reversibility: in the NPV modes, a firm can reverse an investment only after the investment horizon has ended; in ULC it is always allowed.	B	Once invested, a firm is locked into AI for a fixed time.
C4.7	Adoption per task is independent: no fixed costs, economies of scale or complementarity of automation across tasks or the firm.	S	Minimum-viable-scale effects at the firm level are missing.
C4.8	Severance payments are counted as upfront costs upon automation, but only for permanent workers.	B	Flex workers may under certain circumstances also be entitled to severance, but this is ignored here.
C4.9	Risk-neutral point expectations: no uncertainty or variance in the NPV.	B	No option value in waiting; this can accelerate adoption.

H.7. C5 - Learning curve

Table 36 lists the assumptions that govern the path of the AI cost over time.

Table 36: The path of the AI cost over time.

ID	Assumption	Type	Implication/Limitation
C5.1	AI costs decrease over time following a closed, exogenous exponential decay curve with a floor.	F	AI becomes cheaper down to the floor; the curve determines the adoption timing.
C5.2	The cost curve is exogenous and driven only by time, independent of cumulative adoption or production in the model.	S	There is no endogenous learning-by-doing: adoption does not influence the AI price (no demand → price feedback).
C5.3	One national AI price, the same for all firms and tasks; no heterogeneity or negotiation over AI costs.	S	Differences in access to AI remain out of view.
C5.4	Only AI follows an exogenous cost curve; wages do not.	S	AI is set to become structurally cheaper than labour over time, which pushes adoption.

H.8. C6 - Demand

Table 37 lists the assumptions behind the goods market and the two-class consumption function that links unemployment to demand.

Table 37: The goods market and the two-class consumption function that links unemployment to demand.

ID	Assumption	Type	Implication/Limitation
C6.1	Linear inverse demand (price decreases linearly with the total quantity), one market, one price, market clearing every step.	F	No inventory build-up or rationing; everything sells at the price.
C6.2	Two-class consumption demand: the demand intercept depends on both wage income and profit income.	S	The split between wages and profits influences total demand.
C6.3	Income affects demand with exactly one step of delay.	F	The feedback is delayed by one period.
C6.4	Different propensities to consume: capital income is consumed less than wage income.	F	Redistribution from wages to profit lowers demand; demand is wage-led.
C6.5	No explicit saving, wealth or autonomous investment demand; all factor income feeds demand via the consumption propensities.	S	No wealth accumulation or independent investment cycle.
C6.6	All output is sold and the price remains positive; no unsold production.	F	No involuntary inventory build-up; demand shortages are reflected in the price only.
C6.7	Operating profit excludes one-time costs: upfront AI investment and severance pay are not included in profit income.	F	Profit income in the demand function reflects current, not one-time, cash flows.

H.9. C7 - Employment protection

Table 38 lists the assumptions on contracts, severance and the chain regulation.

Table 38: Contracts, severance and the chain regulation. The Dutch institutional rules of the model.

ID	Assumption	Type	Implication/Limitation
C7.1	Two contract types (flex and permanent), the binary Dutch distinction.	S	Intermediate forms (temporary agency work, payroll, self-employed) are reduced to two categories.
C7.2	Severance pay applies only to permanent employees and increases with salary and years of service, flexible employees have no statutory entitlement.	B	The core of the institutional asymmetry: dismissing or automating permanent work is more expensive than flexible work.
C7.3	Chain regulation: a flex worker who reaches the chain limit is converted to permanent employment with a certain probability; otherwise, the contract is not extended (provided a replacement is available; otherwise forced conversion).	B	Periodic pressure to terminate flexible contracts before conversion, or to convert them.
C7.4	New hires always start as flex workers with zero years of service.	B	Permanent contracts are created only through conversion, not upon commencement of employment.
C7.5	Non-renewal via the chain regulation does not trigger an automation evaluation, other than an exogenous separation.	B	The turnover of flex does not lead directly to automation.
C7.6	The initial contract and years of service distribution is set exogenously at the warm-start.	S	The initial state of the institute is imposed, not endogenously built up.

C7.7	Severance pay is the only explicit dismissal cost; procedural costs (UWV, subdistrict court) are only implicitly included in the compensation amount.	D	Procedural thresholds and dismissal lead times are not modelled separately.
C7.8	A main switch: if protection is switched off, contracts, compensation, and chain regulation do not work.	S	Allows for a clean comparison between the Dutch and an unregulated regime.

H.10. M1 - Macro feedback

Table 39 lists the assumptions in the layer that connects the components.

Table 39: The connection layer between all the different components.

ID	Assumption	Type	Implication/Limitation
M1.1	The sub-steps follow a fixed order. The order itself is a model choice.	S	The timing of events within a step influences the outcomes.
M1.2	The macro-feedback has a delay of exactly one time step.	F	The link between income and demand is set to one period.
M1.3	Production and price are calculated twice per step: once at the beginning (for the delayed income) and once at the end (for consistent data).	T	A consistency assumption without independent behavioural effects.
M1.4	Synchronous updates: all agents are updated at the same time without affecting each other.	S	Order effects within the population of agents are averaged out.

H.11. Technical assumptions

The following assumptions concern implementation or reproducibility and have no independent significance. They are listed for the sake of completeness (Table 40).

Table 40: Technical assumptions

ID	Assumption
T.1	A fixed seed for the random generator ensures reproducibility; all stochasticity follows pseudo-random functions.
T.2	The warm-start uses an isolated random generator, so parameter changes do not alter the starting conditions (variance reduction across experiments).
T.3	The warm-start assigns every task exactly one worker.
T.4	Numerical floors guard against division by zero. Productivity and price have a small positive minimum.
T.5	The number of AI units is rounded up, with a minimum of one.
T.6	The NPV cache may be one step out of date for unevaluated tasks; this affects only diagnostics, not behaviour.
T.7	Data are collected once per step; the t=0 state is recorded separately after the warm-start.
T.8	The Gini coefficient is calculated on all employed persons, with the unemployed counting as zero income.

H.12. Most influential assumptions

Not all assumptions carry equal weight. The following decisions shape the outcomes, and their interpretation, most strongly, and deserve explicit attention in the discussion of the results.

Perfect complementarity (Leontief) - C1.1

The minimum function makes every task a potential bottleneck and tightly links automation, employment, and output. Virtually all dynamics in the model run through this mechanism.

Exogenous, decoupled AI learning curve - C5.2 and C5.4

AI inevitably becomes cheaper along a purely time-driven path, without feedback from adoption to price. This strongly influences the timing of adoption and makes AI structurally cheaper than labour in the long run.

No augmentation: human or AI per task - A0.12

AI can only replace labour, not supplement it. The model thus captures displacement but not the complementarity with which much generative AI is deployed. This is one of the most defining simplifications and steers the outcomes towards a purely substitution narrative.

Fixed, non-re-assignable skill groups and no new tasks - A0.8 and C1.5

The compensation or reinstatement channel is disabled. Consequently, the inequality effects are likely an upper bound and not a central estimate.

Coarse skill–task mapping - C1.8

The model ties the routine/non-routine task distinction to the skill distinction along a single axis: routine tasks are skill-neutral and only non-routine tasks favour the high-skilled, so in equilibrium the low-skilled sort into routine work and the high-skilled into non-routine work. This dividing line is too coarse for reality, where the two dimensions do not coincide. There are high-skilled routine tasks — much routine cognitive professional work, from standard radiological reads to bookkeeping and document review — that generative AI exposes precisely because they are routine, and there are low-skilled non-routine tasks — cleaning, care and many manual service jobs — that resist automation precisely because they are non-routine and manual. By collapsing the routine/non-routine and the high/low-skill axes onto one another, the model can misplace where automation lands across the skill groups, plausibly understating the exposure of high-skilled routine work and overstating that of low-skilled non-routine work. The limited cross-skill matching (C2.4) softens but does not remove this, since the underlying productivity advantage runs along the non-routine dimension alone. This bears directly on the skill-premium and the distribution of unemployment across the two groups.

One homogeneous good and one sector - A0.2

The Dutch dual narrative is strongly sectoral, but sector heterogeneity is lacking. Outcomes apply to an average economy, not to specific sectors.

Passive workers and frictionless central matching - C2.1 and C2.3

Wage and employment dynamics are entirely demand-driven; search and bargaining behaviour on the supply side is absent.

Perfect, risk-neutral anticipation of AI costs - C4.4 and C4.9

Firms know the AI cost path exactly and do not value the option value of waiting, which can accelerate adoption compared to a world of uncertainty.

No government or fiscal redistribution - A0.3

The reported inequality is gross, before taxes and benefits, and thus overestimates net inequality in practice and changes the Gini-index.

Taken together, these assumptions position the model as a scenario-exploratory instrument: it explores directions, feedbacks, and sensitivity to institutions and policy, and is not intended as a point prediction of future wage inequality.