

# Strategic Injustice in AI-Assisted Goal Setting: The Marginalization of Social Objectives

Mikolaj Pindelski

Warsaw School of Economics, Poland

[mikolaj.pindelski@sgh.waw.pl](mailto:mikolaj.pindelski@sgh.waw.pl)

**Abstract:** The study investigates how artificial intelligence driven decision making influences the prioritization of traditional financial objectives versus sustainability related goals. Using a modified “cake sharing” model in the NetLogo simulation environment, it was examined three objectives: improving ROI, increasing revenue, reducing costs, and ensuring sustainability of sales activities. There were run 500 iterations in the simulation phase. The simulation paired the sustainability oriented goal (SDG) with each financial objective. There was observed potential systematic discrimination against non-financial targets. Input data were derived from CD Projekt’s 2024 financial and non-financial reports, analyzed with ChatGPT to recommend strategic priorities before integration into the NetLogo framework. Findings reveal that AI generated recommendations favor financial objectives, particularly ROI, leading to marginalization of sustainability goals. It confirmed the hypotheses that AI supported goal setting can reinforce a bias. While the model is simplified and limited to four objectives under controlled conditions, the results underscore the risk of relying on generative AI in strategic planning. The study highlights the need for managers to critically assess algorithmic assumptions in AI-supported decision making.

**Keywords:** SDG, LLM, Goals discrimination, Artificial Intelligence, OpenAI, Strategic discrimination

---

## 1. Introduction

The growing application of artificial intelligence and AI driven large language models across both individual enterprises and broader economic systems is becoming increasingly evident. Numerous studies point to the displacement of human labor by algorithmic and AI-driven systems (Ming-Hui, Rust, 2018; Frey & Osborne, 2013, 2017; Frank et.al. 2019) and the transformative influence of AI on economic processes (Agrawal et al., 2019). The AI-driven forecasting and performance optimization techniques (Mehendale et al., 2019) and demand prediction systems (Bandara et al., 2019) support the dominance of financial objectives in AI-supported strategies primarily increase of revenue (Khalimonchuk, Pozovna, 2024; Trivedi & Patel, 2020). This trend draws attention to an important research question: does the use of artificial intelligence in strategic goal setting result in a consistent preference for traditional financial outcomes, potentially at the expense of broader, socially-oriented goals such as those aligned with the Sustainable Development Goals (SDGs)?

The aim of this study is to examine the effects of setting different strategic objectives and the way they evolve over time. It is to be checked by applying the agent-based simulations in the NetLogo software environment (Tisue & Wilensky, 2004). Specifically, four objectives were modeled:

1. maximizing return on investment (ROI),
2. increasing revenue,
3. reducing costs,
4. promoting sustainability in sales, treated here as the SDG-related objective.

To test for potential bias or discrimination in strategic recommendations, the SDG objective was paired in turn with each of the three financial objectives: ROI, revenue, and cost reduction. Each pair was simulated across 500 iterations to observe how the prioritization of one objective affected the progression of the other. The simplified input parameters were defined for each scenario, after which financial and non-financial data of the CD Projekt company were entered into ChatGPT. The reports are open and available on the www of the company. The model was then prompted to select the preferred goal within each pairing.

Following this, a NetLogo analytical framework was applied to accelerate the simulation process and track the evolution of indicators over time. Once initial results were obtained, they were reintroduced into ChatGPT to generate updated strategic guidance for a subsequent simulation round. This iterative approach was designed to capture how generative AI might influence the strategic trajectory in environments with competing goals. The integration of ChatGPT with NetLogo has already shown promising outcomes in similar applications (Chen et al. 2024; Martinez et al., 2024).

Through this approach, the study aims to identify whether and to what extent the use of generative AI in strategic goal setting leads to the systematic marginalization of socially valuable objectives when paired with dominant financial metrics.

## 2. Motivation and Research Goals

Organisations are increasingly incorporating artificial intelligence systems, algorithmic frameworks, and advanced analytical tools into a wide range of operational tasks, including forecasting, decision-making, and resource allocation. These technologies are applied in diverse areas such as creditworthiness assessment, recruitment, and strategic planning. However, despite their perceived objectivity, algorithms are not free from bias. Growing evidence points to the phenomenon of “algorithmic injustice,” whereby systems assumed to be neutral may in fact exhibit systematic discrimination. That may happen against specific social groups, types of decisions, or strategic priorities. Such outcomes may arise not from deliberate intent but from the structure of training data and the modelling assumptions embedded in the algorithms. This issue is particularly relevant in the context of strategic management. Beyond operational decisions such as targeting specific customer segments or optimizing operations, algorithms are now influencing strategic decisions, including the formulation of strategic goals. When AI systems are tasked with goal setting, they frequently prioritise objectives associated with economic performance, such as revenue growth or ROI, potentially at the expense of social or sustainability related goals. The risk arises when AI is designed or trained to maximise a single, quantifiable target, potentially leading to the marginalisation of other important objectives.

The aim of this study is to examine how AI influenced strategic target decisions evolve over time, and to what extent prioritising one objective may systematically suppress the development of others. In particular, the study explores whether the use of generative artificial intelligence, such as OpenAI’s large language models, introduces bias in the strategic goal setting process by favouring specific targets over others. To investigate this issue, two research hypotheses are proposed:

*H1: Prioritising one strategic goal over another will lead to a significant decline in the performance of the non prioritised objective over time.*

*H2: The use of generative AI, such as OpenAI’s ChatGPT, in strategic goal-setting may result in biased decisions that systematically discriminate against certain objectives.*

These hypotheses are tested through simulations conducted in the NetLogo platform, using a modified version of the “cake-sharing” model. The model is based on the premise that it is not possible to simultaneously maximise all objectives indefinitely. A condition analogous to the Pareto efficiency principle. Once a Pareto optimal point is reached, further improvement in one objective necessarily entails trade-offs in the others. The simulation assumes a finite and constant resource pool (e.g., 100%), which must be allocated among competing objectives. The core question is whether the AI-driven system will seek to maintain a Pareto-optimal balance across objectives over time, or whether it will consistently prioritise the maximisation of a single AI-recommended goal, to the detriment of others. For example, an aggressive focus on revenue growth may limit the achievement of social impact targets or even undermine cost-efficiency goals. By analysing how these trade-offs play out over time particularly those combining traditional strategic metrics (ROI, revenue, cost) with sustainability-related targets (SDG). The study seeks to uncover patterns of algorithmic preference that may affect long-term strategic balance within the strategy.

## 3. Research Method and Simulation Model Design

The simulation involves the preparation of a model in the NetLogo environment to investigate how the choice of different strategic targets shapes business performance. The simulation assumes 500 iterations, during which four possible operating strategies will be explored: maximising revenue, maximising the value of ROI, minimising the costs and maximising the value of social objectives. An agent simulation model will be created within NetLogo to reflect the dynamics of the strategic processes. A key role in this process will be played by ChatGPT, which will automatically recommend a strategic target for a given simulation cycle based on the analysis of financial and non-financial reports of the CD Projekt company. The decision to select a target will be based on an analysis of prevailing trends, such as the priority of increasing revenue, improving financial efficiency (ROI), reducing costs or increasing social impact (ESG, CSR). In the context of the social objective, general indicators have been selected to be introduced into the model, but which can influence the achievement of the others. These include donating a portion of turnover to social causes, removing provisions from contracts that are disadvantageous to customers but profitable for the organisation, and the organisation of free training by the company's human capital for local communities. These objectives may limit the possibility of achieving the sought-after higher or lower values of the other objectives, respectively.

During the 500 iterations of the simulation, it will be observed how the preference for one objective affects other areas. For example, choosing to maximise revenue may lead to discriminate SDGs. While an emphasis on maximising social value may reduce immediate revenue. The NetLogo software will be used to build result sets, visualise trends and benchmark analyses between alternative strategies, allowing the discrimination of other objectives to be accurately indicated when one objective is dominant. The final result of the simulation will therefore be a full characterisation of the interrelationship between, enabling the most sustainable strategies to be recommended based on the specific market and organisational conditions of the company in question.

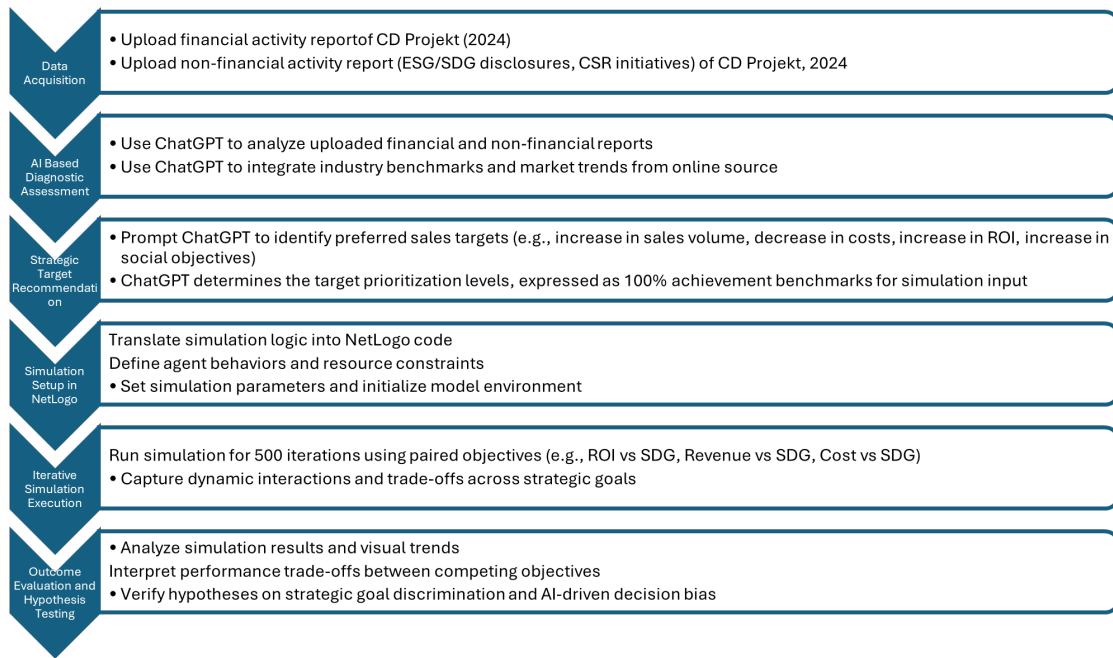
The cake sharing model reflects a situation in which agents compete for predetermined and limited resources or outcomes. Decisions about their allocation consequently affect the targets achieved. In the simulation adopted in the study, the strategic objectives related to the sales function relate to an increase in revenue, a reduction in costs, an increase in ROI and an increase in the social objectives index. These can be treated as a limited resource, or fragments of a whole cake to be shared with a resource constraint and influenced by the chosen strategy of action. Certain assumptions are made here as to the validity of the choice of model. It was therefore assumed that in a cake sharing model, agents would make decisions that account for what part of the cake related to the two goals set in every simulation would be achieved. This model seems to present well the conflict of interest between objectives and indicates the actual situations in which managers, executives and other executives seek a balance between sometimes conflicting demands such as increasing revenue and at the same time spending time on social responsibility tasks. The model thus presents a situation in which objectives compete with each other in a situation of choosing a strategic goal and direction under conditions of limited resources and conflicting interests between objectives.

#### **4. Simulation Procedure for Strategic Goal Setting and Discussion**

The simulation process is conducted within the NetLogo modeling environment. It begins with the initialization of the simulation system and the definition of the operational environment in which agents will act. The next stage involves the integration of a generative AI system, in this case ChatGPT, which partially processes source data to support strategic goal selection. This source data consists of the CD Projekt company's financial and non-financial annual report for 2024, obtained from its official website. Following data integration, ChatGPT is prompted to analyze the input in the context of current market trends and industry forecasts. The AI system is given a standard instruction: *"Based on the company's financial and non-financial reports, as well as market trends and industry forecasts, select the most appropriate strategic target from the indicated options. In doing so, take into account the company's declared Social Development Goals (SDGs) and determine which objective should be prioritized in the current context to ensure both business performance and social value creation."* Based on this analysis, ChatGPT selects one strategic goal expressed as a percentage level of target achievement which is then adopted for simulation purposes. Subsequently, the simulation model in NetLogo is parameterized in accordance with the AI selected strategic objective. The first test phase consists of a single simulation iteration, during which agents act in alignment with the selected strategic priority. Following this, the model proceeds through 500 iterations, in which the evolution of key performance indicators (KPIs) is tracked for all four strategic goals: revenue, costs, ROI, and social impact SDGs related. Each objective is scaled in percentage terms, where 100% represents full attainment of a given goal within a single cycle. For objectives aimed at increase (e.g., revenue, ROI, social impact), upward movement is considered positive, whereas for cost reduction, downward movement is preferred. Importantly, due to the resource-constrained design of the model each agent can only reach 100% progress in one objective per iteration, thus requiring trade-offs between two competing goals in every simulation

The parameters remain constant throughout all 500 iterations to ensure consistency in outcome tracking and to allow the accumulation of a robust data set. This structure facilitates the systematic analysis of how the prioritization of one dominant objective influences the achievement levels of other goal. By comparing performance outcomes under each scenario, the simulation enables the identification of potential bias or "goal discrimination" resulting from AI-guided decision making.

Finally, the results of the simulation are presented visually through a series of graphs, illustrating the dynamic interdependencies between objectives and highlighting how the system behaves under different prioritization strategies. This provides empirical grounding for verifying the research hypotheses and evaluating the long-term consequences of AI-influenced strategic goal-setting.



Source: Author’s own work

**Figure 1: Structure of the Simulation Process**

The simulation model is illustrated in Figure 1. It comprises six main stages to explore how AI guided decisions affect strategic goal setting within a resource constrained environment. It is specifically applied to CD Projekt, using their financial and non-financial data for 2024. The model integrates financial performance metrics with sustainability oriented goals aligned with SDGs to evaluate the trade-offs inherent in AI-driven strategy selection.

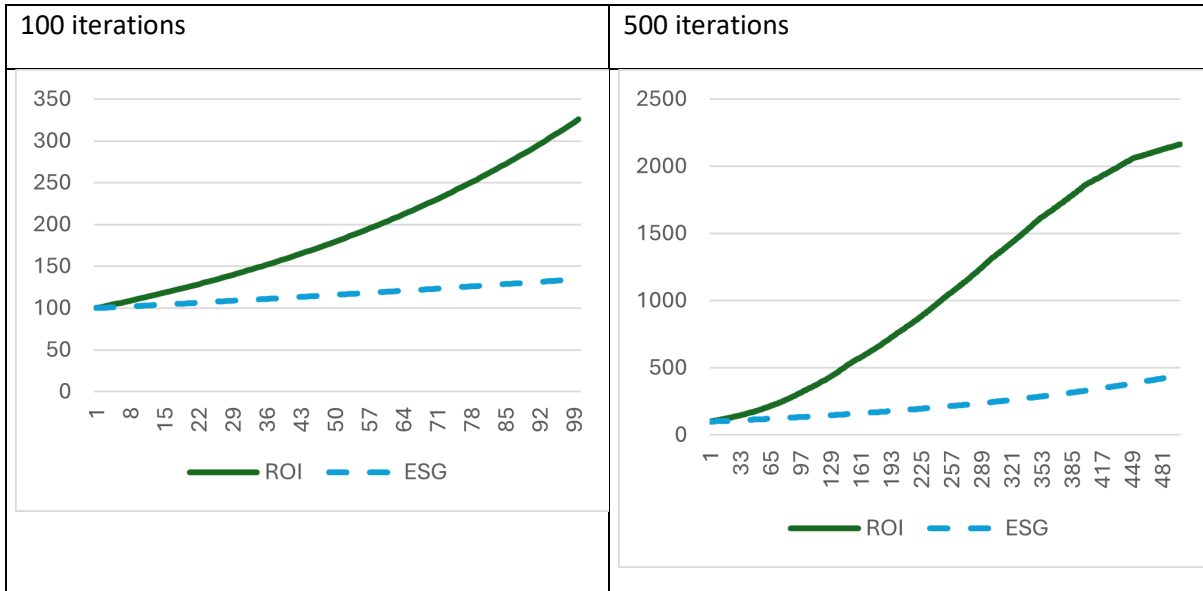
The simulations produced four distinct scenarios, each reflecting the system’s preference for a specific strategic target. In every case, the prioritization of one objective resulted in the relative marginalization or discrimination of the remaining goal. Though the intensity of this effect varied across scenarios. Notably, even when the simulation was interrupted after 100 iterations and the intermediate results were fed back into ChatGPT, the AI did not recommend substantial adjustments to the chosen target and their relation. This suggests that, under stable input conditions, the generative AI consistently maintains its original prioritization strategy, reinforcing the imbalance among competing objectives.



Source: Author’s own work

**Figure 2: Simulation for the revenue growth target, 100 and 500 iterations**

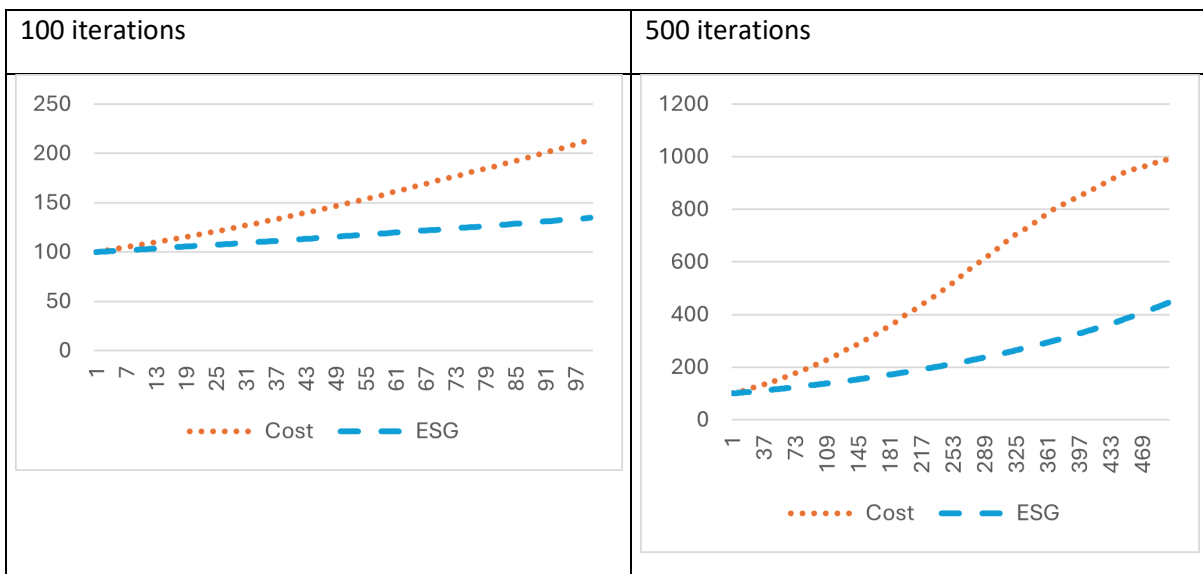
Revenue exhibits (Figure 2) a steep and accelerating upward trajectory, especially after iteration 100, ultimately exceeding 2000 units by the final iterations. This suggests that when revenue is prioritized as the strategic target, it dominates the simulation, compounding strongly over time. ESG performance, in contrast, increases at a much slower, linear-to-moderate exponential pace. Its value remains significantly lower throughout the simulation, never exceeding 500 units, and showing limited responsiveness despite being tracked in parallel. The divergence between the two lines becomes increasingly pronounced as the number of iterations progresses, visually illustrating the discriminatory effect: when revenue is prioritized as the dominant strategic goal, ESG-related performance is not proportionally developed and remains marginalized.



Source: Author’s own work

**Figure 2: Simulation for the ROI growth target, 100 and 500 iterations**

ROI begins to rise steadily from its baseline value (~100) and follows a consistent upward trajectory, ending above 300. ESG performance also increases but at a far slower and almost linear rate, ending just above 120. The gap between ROI and ESG widens gradually over these 100 iterations. In the short term, prioritizing ROI leads to significantly stronger growth in financial return compared to ESG progress. ESG performance, although not flat, shows only marginal improvements, suggesting a subordinate role when ROI is the dominant goal.



Source: author’s own work

**Figure 2: Simulation for the Cost reduce target, 100 and 500 iterations**

Costs increase exponentially, surpassing 1000 by the final iterations. This likely reflects sustained emphasis on cost as a strategic variable. Meanwhile ESG performance improves steadily, reaching over 400, but the growth remains significantly slower compared to the escalation of costs. The disparity becomes more visible over time, indicating that the longer the AI prioritizes cost as the dominant goal, the greater the lag in ESG performance. Over the full simulation period, the emphasis on decreasing cost leads to an increasingly disproportionate outcome. While ESG objectives show some progress, they are clearly deprioritized.

Across all simulation scenarios, ESG performance consistently remains the least developed metric when paired with any dominant financial goal. Despite minor improvements, ESG never surpasses 25% of the dominant goal's value by the end of the simulation. This confirms both hypotheses:

*H1: There is a clear trade-off—maximizing a financial goal leads to underperformance in ESG.*

*H2: ChatGPT-based AI decision-making exhibits consistent bias toward financial metrics, systematically discriminating against social objectives in strategic goal-setting.*

**Table 1: Regression for variables Dependent Variable: ESG, Independent Variables: Revenue, Cost, ROI**

Variable	Coefficient	Std. Error	t-Statistic	P-value
const	101.12	0.17	608.74	0.000
Revenue	-0.0676	0.0019	-36.39	0.000
Cost	+0.0953	0.0023	41.68	0.000
ROI	+0.0125	0.0004	32.58	0.000

Source: Author’s own work

R<sup>2</sup> equals 0.966 what means that the model explains 96.6% of the variance in ESG values, which is very high result. All variables are highly statistically significant (p < 0.001). Revenue has a negative and significant impact on ESG. This confirms that prioritizing revenue is strongly associated with lower ESG gains a key sign of ESG discrimination when profit dominates strategy. Cost decrease as a strategic goal has a strong effect on ESG. This suggests that some cost goal realization may result with worse ESG performance. ROI is positively associated with ESG, but its magnitude is small. While statistically significant, its practical influence is limited.

*H1 confirmed: When one financial goal dominates, ESG suffers or advances less rapidly.*

*H2 confirmed: AI prioritization leads to biased decision-making, especially against ESG, notably when revenue is favored.*

## 5. Summary

The study used ChatGPT as a generative artificial intelligence system to simulate the selection and prioritization of strategic targets. While financial and non-financial reports from a Warsaw Stock Exchange-listed company CD Projekt serve as data sources, the recommendation of target preferences proposed by OpenAI ChatGPT based on proprietary algorithms and market trend analysis remains partly unverifiable. Even asked, the LLM listed just three internet sources that seemed to be inefficient in explanations of the internet based market analysis made by ChatGPT. These were explaining only a part of its conclusions and final recommendations and were rooted in main www site of the company and stocks exchange official reports and publications. The OpenAI was not creative in that field as it was not looking for additional or unobvious sources. This introduces an inherent opacity into the decision-making process, particularly concerning how AI balances competing strategic objectives.

Using the NetLogo environment and the cake-sharing model, the study simulated the evolution of AI-recommended goals over 500 iterations, examining four distinct prioritization strategies: maximizing revenue, minimizing costs, increasing return on investment (ROI), and maximizing social objectives that were supposed as aligned with SDG principles. The simulations revealed consistent patterns of strategic discrimination, where the target selected by AI was advanced disproportionately at the expense of the other. Notably, even after pausing the simulation at iteration 100 and reintroducing the outcomes into ChatGPT for reassessment, no substantial adjustment in prioritization was proposed. This indicates a limited capacity of the AI to self-correct its decision trajectory, especially under stable input conditions.

The findings confirm Hypothesis H1: prioritizing a single goal—such as ROI or revenue—leads to a significant reduction in the advancement of non-prioritized goals, particularly ESG. Furthermore, the regression analysis based on all 500 iterations substantiates Hypothesis H2 by demonstrating that while ESG values do rise over

time, their rate of growth is significantly lower compared to financial indicators. Specifically, revenue was found to have a statistically significant negative impact on ESG performance, whereas cost and ROI had positive but asymmetric effects, with cost exerting a stronger influence.

These results underscore the risk that AI based decision tools can entrench financial performance at the expense of social objectives, particularly when ESG metrics are less quantifiable or slower to manifest in observable outcomes. In such contexts, AI may reinforce a narrow optimization logic, producing decisions that conflict with long-term sustainability goals.

Thus, it must be concluded that in AI-supported strategic planning, the prioritization of one objective inherently constrains the development of others, especially those linked to ESG. Managers and decision-makers should approach AI recommendations critically, equipped with a strong understanding of the algorithms and trade-off mechanisms involved. Only through algorithmic literacy and transparent model governance can organizations ensure that AI contributes to balanced, equitable, and responsible goal-setting.

**Ethics declaration:** Due to the fact that all the reports used in the research were available online on the official website of the company and the company is obliged to publish them it is no need to provide any permission to use them for the simulation. As it is also stated in the text that it is a simulation based on the reports and not an official company's strategic goals.

**AI declaration:** Selected sections of the paper such as the Abstract, Introduction, and Summary were reviewed with ChatGPT 5.0 to refine the language, improve clarity, and enhance readability through recommended word choices and adjustments of sentence structure.

## References

- Agrawal, A., Gans, J., Goldfarb A. (2019) *Economic policy for artificial intelligence*, Innovation policy and the economy, Vol.19, Iss.1, pp.139-159.
- Bandara, K., Shi P., Bergmeir, C., Hewamalage H., Tran, Q., Seaman B. (2019) *Sales Demand Forecast in E-commerce Using a Long Short-Term Memory Neural Network Methodology*, Neural Information Processing, pp.462 - 474
- Chen J., Lu X., Du Y., Rejtig M., Bagley R., Horn M., Wilensky U. (2024) Learning Agent-based Modeling with LLM Companions: Experiences of Novices and Experts Using ChatGPT & NetLogo Chat, International Conference on Human Factors in Computing Systems, DOI: <https://doi.org/10.1145/3613904.3642377>
- Fischer, H., Seidenstricker, S., Berger, T., Holopainen, T. (2022) *Artificial intelligence in B2B sales: Impact on the sales process*, Artificial Intelligence and Social Computing, Vol.28, Iss.28, pp.135-142.
- Frank D., Autor D., Bessen J.E., Brynjolfsson E., Cebrián M., Deming D., Feldman M., Groh M., Lobo J., Moro E., Wang D., Youn H., Rahwan I., 2019, *Toward understanding the impact of artificial intelligence on labor*, Proceedings of the National Academy of Sciences of the United States of America, DOI: <https://doi.org/10.1073/pnas.1900949116>
- Frey, C.B., Osborne M.A. (2017) *The Future of Employment: How Susceptible Are Jobs to Computerisation?*, Technological Forecasting and Social Change vol.114, pp.254–280
- Frey, C.B., Osborne, M.A. (2013) *The Future of Employment: How Susceptible Are Jobs To Computerisation*, Oxford Martin School Working Paper, [/Users/michaelosborne/Documents/Research/future\\_of\\_employment\\_paper/original/.texpadtmp/future\\_of\\_employment\\_original.dvi](/Users/michaelosborne/Documents/Research/future_of_employment_paper/original/.texpadtmp/future_of_employment_original.dvi)
- Heydari Fard, S. (2022) *Strategic injustice, dynamic network formation, and social movements*, Synthese, Vol.200, Iss.5 art. No.392
- Hildebrand, C., Bergner, A. (2019) *AI-driven sales automation: Using chatbots to boost sales*. NIM Marketing Intelligence Review, Vol.11, Iss.2, pp.36-41.
- [Home - CD PROJEKT https://www.cdprojekt.com/en/](https://www.cdprojekt.com/en/) (2025.05.30)
- <https://ccl.northwestern.edu/netlogo/> (2025.05.20)
- Huang, Ming-Hui and Roland T. Rust. (2018) *Artificial Intelligence in Service*, Journal of Service Research Vol.21 (2018): pp.155 - 172.
- Khalimonchuk I.V., Pozovna I. (2024), *The role of machine learning and artificial intelligence in optimizing costs and increasing revenues of technological companies*, Economic sustainability and business practices, DOI: DOI:10.21272/esbp.2024.3-04
- Kirchdorfer, L., Blümel, R., Kampik, T. et al., 2025, Discovering multi-agent systems for resource-centric business process simulation. *Process Sci* 2, 4 . DOI: <https://doi.org/10.1007/s44311-025-00009-5>
- Martínez, J., Llinas, B., Botello, J.G., Padilla, J.J., Frydenlund, E., (2024) *Enhancing GPT-3.5's Proficiency in Netlogo Through Few-Shot Prompting and Retrieval-Augmented Generation*. 2024 Winter Simulation Conference, WSC, IEEE, pp.666-677
- Mehendale, A., HR, N.S. (2018) *Application of Artificial Intelligence (AI) for Effective and Adaptive Sales Forecasting*, Journal of Contemporary Management Research, Vol.12, Iss.2.

- Paschen, J., Paschen, U., Pala, E., Kietzmann, J. (2021) *Artificial intelligence (AI) and value co-creation in B2B sales: Activities, actors and resources*. Australasian Marketing Journal, Vol. 29, Iss.3, pp.243-251.
- Sulis E., Taveter K, 2022, *Agent – Based Business Process Simulation. A Primer with Applications and Examples*, Springer,
- Tissue, S. Wilensky, U. (2004) *Netlogo: A simple environment for modeling complexity*, International conference on complex systems, May, Vol. 21, pp.16-21
- Trivedi, S., Patel, N. (2020) *The Role of Automation and Artificial Intelligence in Increasing the Sales Volume: Evidence from M, S, and, MM Regressions*, June 20, 2020. DOI: <http://dx.doi.org/10.2139/ssrn.4180379>