

# Multi-Agent Modeling of Food and Water Security Crises: From Historical Causation to Robust Policy Design

Doctoral Consortium

Daria Taniushkina

AI Center and IAI MSU, MIRAI, Moscow Institute of Steel and Alloys

Moscow, Russia

d.taniushkina@iai.msu.ru

## ABSTRACT

Food and water insecurity have repeatedly been believed to trigger political instability across history from Late Ming China (1630s famine), French (1780s) and Russian (1917–1918; both food scarcity) Revolutions to the Arab Spring (2010–2011 food price spikes) and 2021–2022 Kyrgyzstan–Tajikistan conflict (water scarcity). Yet these cases mostly lacking systematic data-driven analysis resulting in very limited predictive and preventive frameworks. Our research integrates agent-based modeling, game theory, and causal inference to understand root causes, event dynamics, and design robust policies for conflict prevention and mitigation. Thesis contributions are (1) agent-based model with heterogeneous strategic agent types (farmers, governments, suppliers, consumers, suppliers) capturing resource scarcity dynamics and agents response; (2) investigation of equilibrium robustness of agents behavior under adversarial influence; (3) interpretable policy design via LLM-augmented expert knowledge extraction and multi-agent reinforcement learning. Finally, I plan for retrospective analysis of historical conflict cases outlining augmented perspective on the real reasons behind them.

## KEYWORDS

Agriculture; Multi-Agent Systems; Social Conflicts

### ACM Reference Format:

Daria Taniushkina. 2026. Multi-Agent Modeling of Food and Water Security Crises: From Historical Causation to Robust Policy Design: Doctoral Consortium. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026)*, Paphos, Cyprus, May 25 – 29, 2026, IFAAMAS, 3 pages. <https://doi.org/10.65109/XKPC1340>

## 1 INTRODUCTION

Agriculture faces a dual challenge: increasing production to meet rising demand while adapting to much more volatile climate conditions that shift precipitation patterns, temperature extremes, and therefore water and food availability. As a result, food and water crises may occur more often across the globe, especially in regions with high population density and climates that are already unfavorable for agriculture. Historically, a strong correlation has been observed between major food crises and political instability that can lead to military conflict [15]. Notable historical cases include the collapse of the late Ming Dynasty in China (1630s–1644; famine),

the French Revolution (1770s–1780s; grain crisis/Flour War [39]), the Russian Revolution (1917; bread shortage [38]), and more recent events such as the Arab Spring (2010–2011; high food prices [10, 14]) and the Syria war (2006–2010; drought [7, 12]). Finally, despite food vs. crisis correlation, conflicts often have more complex causes where lack of food and water is not the only driver. For example, global food prices spiked in 2007–2008, but only about 40 countries experienced riots [4]. Furthermore, Tunisia, the Arab Spring epicenter, maintained low food price pass-through through subsidies [5]. Power disruptions might be another major factor [6, 9, 17, 29].

We employ multi-agent modeling coupled with assessment of water availability, land fertility, and crop yields to understand root causes, dynamics, and policies that can reduce the risk of conflict. With this framework, we aim to predict potential future conflict risks and develop prevention policies when possible. In our study, we use several classes of agents, including farmers (who may hoard grain when anticipating further price increases, amplifying scarcity), governments (export bans that increase global shortages; subsidies that deplete budgets; repression that escalates unrest), suppliers (speculation and market manipulation), and other actors, each parameterized by a small number of interpretable parameters. Each agent optimizes its own reward function, resulting in a (mean-field) saddle point optimization problem. A substantial part of the work focuses on designing policies based on LLM-augmented expert knowledge and mean-field multi-agent reinforcement learning (MAREL) [34, 35]. Preventive policies (subsidies, strategic reserves, water treaties) aim to avoid escalation and must be interpretable for stakeholders and robust, meaning effective even when agents adapt or when external actors attempt to manipulate behavior.

*Core thesis.* By integrating agent-based modeling of heterogeneous strategic behavior with learned behavioral distributions, robust optimization, and causal inference, we can build interpretable models that identify early warning signals and support robust preventive and corrective policies for food- and water-driven instability. In addition, the resulting framework can help clarify the drivers behind historical conflicts and inform ways to prevent future ones.

## 2 BACKGROUND

*Related work.* In [25], the authors developed an agent-based model (ABM) for global food security and food riots, simulating production shocks and analyzing riot emergence. Their model was validated on the 2007–2008 crisis but lacked strategic agent optimization. In [1], ABMs were used for food systems with spatial interactions, demonstrating how township protests spread via agent dynamics. For water, [31] proposed coupled ABM–drought models



This work is licensed under a Creative Commons Attribution International 4.0 License.

*Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026)*, C. Amato, L. Dennis, V. Mascardi, J. Thangarajah (eds.), May 25 – 29, 2026, Paphos, Cyprus. © 2026 International Foundation for Autonomous Agents and Multiagent Systems ([www.ifaamas.org](http://www.ifaamas.org)). <https://doi.org/10.65109/XKPC1340>

analyzing human–drought feedback, but without game-theoretic strategic agents. Water allocation risks are studied in [3, 19, 20, 23].

Causal inference from historical data is an important part of this work. Structural causal models (SCMs) formalize causation via directed acyclic graphs and structural equations [26]. SCMs enable counterfactual reasoning (e.g., “What if France had grain reserves in the 1780s?”) and causal effect identification from observational data [27]. For political instability, [8] applied causal frameworks to revolutions, identifying necessary vs. sufficient conditions. In [39], climate conditions were used as an exogenous shock to estimate the causal effect of harvest failure on French Revolution peasant revolts, finding a significant positive effect [39]. Related causal evidence is discussed in [12]. Self-learning could help with limited data [2, 21].

Finally, several case studies support the view that water and food scarcity can contribute to social instability. In [39], the authors quantified the causal effect of droughts on social instability, focusing on the French Revolution. A strong correlation between food price index peaks and riot onset during the Arab Spring was reported in [14], with a notable exception of Tunisia [22]. Political conditions and unemployment were important co-factors as well [11].

*Gap.* Despite substantial interest in this topic, an end-to-end pipeline that jointly analyzes multiple correlated factors (water and food availability, social and macroeconomic factors), their roles in conflict risk and dynamics, and practical ways to mitigate risk is still largely missing. Furthermore, to the best of our knowledge, there is no standardized dataset that allows the researchers to compare methods on this problem. This thesis goal is to close these gaps.

### 3 COMPLETED AND ONGOING WORK

This thesis develops an integrated, dynamic model of food and water security in which environmental degradation, climate variability, and strategic behavior jointly shape the emergence of scarcity, migration, and political instability. The central premise is that crisis dynamics are rarely driven by climate shocks alone: they are amplified (or dampened) by responses from farmers, markets, and governments, as well as by slower-moving constraints such as soil degradation and declining land productivity.

The first component constructs a land fertility module that treats soil quality as a state variable evolving over time. It was addressed earlier [13, 24, 28, 30, 36, 37] by assessing future land fertility/risks under different climate change scenarios. The results support our concern that highly populated regions may face major agricultural risk. In addition, we can model not only future conditions but also past conditions, which allows for the assessment of major episodes of political instability. Important adjustments are still needed: supply chains, crop yields, and other factors must be calibrated to reflect technological, behavioral, and logistical changes over time.

The second component builds a multi-agent simulation of food and water systems with heterogeneous agents, including producers, consumers, intermediaries, and state actors. Agents differ in wealth, landholdings, risk preferences, geographic exposure, and access to information, and they operate under different levels of strategic complexity. In ongoing work, we experiment with the number and types of agents to produce a simple but powerful and interpretable framework. To scale strategic interaction to very large populations, the dissertation uses mean-field multi-agent reinforcement learning

to approximate local interactions through aggregate neighborhood effects, enabling analysis at the scale of millions of farmers while retaining adaptive decision-making [40]. A complementary work in progress focuses on learning behavioral responses to weather and price shocks so that behavioral distributions [33] can shift endogenously as conditions change.

The third component provides a theory and robustness layer that explains when decentralized behavior becomes collectively harmful and how policy can shift incentives. Resource scarcity is framed as a strategic setting in which farmer hoarding and government interventions (subsidies, export controls, reserve releases) can admit Nash equilibria that are privately rational yet socially suboptimal, producing amplification dynamics consistent with prisoner’s dilemma and commons-type failures. The dissertation also studies robustness under adversarial manipulation and distribution shifts in behavior by evaluating policies under worst-case stress tests (disinformation, logistical sabotage). This work is in progress and planned for the next 12 months.

The final component translates model insights into interpretable prevention policies and validates the framework against historical cases. Policy design emphasizes transparency through rule-based decision structures (e.g., threshold-triggered reserve releases) augmented by explainable machine learning when forecasting is required; post-hoc explanations use SHAP to communicate which factors drive predicted instability and to generate actionable counterfactuals [18]. An LLM-assisted workflow is explored to extract candidate policy rules from institutional guidance and historical precedent, followed by expert vetting and formalization into interpretable decision trees [32]. Validation proceeds through historical replay and calibration on the French Revolution, the Arab Spring, and the Syria drought, comparing simulated timing, intensity, and spatial patterns of unrest and migration to historical records [7, 10, 12, 14, 16, 39]. The primary contribution will be a unified, empirically anchored modeling and policy framework that links environmental dynamics and strategic adaptation to instability risk, enabling early warning and robust, interpretable intervention design. This work is planned for completion in the next 24 months.

In the next year, I plan to finalize the historical dataset assembly and complete preliminary validation of the ABM that models agent responses to food and water scarcity. I plan to iteratively expand the dataset, focusing initial experiments on two separate datasets: one representing an earlier historical era and the other representing the recent century, motivated by data availability and also by a step change in technology and infrastructure. The remaining time will be devoted to causal inference and root cause analysis.

*Feedback request.* I would appreciate feedback from mentors on (i) the right balance between model realism and interpretability when selecting agent types, state variables, and policy levers; (ii) whether the mean-field MARL formulation and the chosen population-level state representation are appropriate; (iii) how to set up a defensible calibration and validation protocol, including what metrics and ablation studies would most convincingly separate climate-driven shocks from strategic amplification effects. I would also value guidance on how to make the work more impactful to both the multi-agent systems community and domain stakeholders.

## ACKNOWLEDGEMENT

The work is supported by the Ministry of Economic Development of the Russian Federation (agreement No. 139-15-2025-013, dated June 20, 2025, subsidy identifier 000000C313925P4B0002)

## REFERENCES

- [1] Anup Adhikari, Leen-Kiat Soh, Deepti Joshi, Ashok Samal, and Regina Werum. 2023. Agent Based Modeling of the Spread of Social Unrest Using Infectious Disease Models. *ACM Transactions on Spatial Algorithms and Systems* 9, 3, Article 17 (2023), 31 pages. <https://doi.org/10.1145/3587463>
- [2] Massih-Reza Amini, Vasilii Feofanov, Loic Pauletto, Lies Hadjadj, Emilie Devijver, and Yury Maximov. 2025. Self-training: A survey. *Neurocomputing* 616 (2025), 128904.
- [3] Erik J. H. Ansink. 2009. *Game-theoretic models of water allocation in transboundary river basins*. Ph.D. Dissertation. Wageningen University. <https://doi.org/10.18174/10993>
- [4] Rabah Arezki and Markus Brückner. 2011. *Food prices and political instability*. IMF Working Paper WP/11/62. International Monetary Fund. <https://www.imf.org/external/pubs/ft/wp/2011/wp1162.pdf>
- [5] Clemens Breisinger, Teunis van Rheenen, Claudia Ringler, Alejandro Nin Pratt, Nicholas Minot, Catherine Aragon, Bingxin Yu, Olivier Ecker, and Tingju Zhu. 2010. *Food security and economic development in the Middle East and North Africa: Current state and future perspectives*. IFPRI Discussion Paper 00985. International Food Policy Research Institute (IFPRI).
- [6] Kairui Feng, Min Ouyang, and Ning Lin. 2022. Tropical cyclone-blackout-heatwave compound hazard resilience in a changing climate. *Nature communications* 13, 1 (2022), 4421.
- [7] Peter H. Gleick. 2014. Water, drought, climate change, and conflict in Syria. *Weather, Climate, and Society* 6, 3 (2014), 331–340.
- [8] Jack A. Goldstone. 2014. *Revolutions: A Very Short Introduction*. Oxford University Press.
- [9] Joe Gorka, Tim Hsu, Wenting Li, Yury Maximov, and Line Roald. 2024. Cascading blackout severity prediction with statistically-augmented graph neural networks. *Electric power systems research* 234 (2024), 110738.
- [10] Cullen S. Hendrix and Henk-Jan Brinkman. 2013. Food insecurity and conflict dynamics: Causal linkages and complex feedbacks. *Stability: International Journal of Security and Development* 2, 2, Article 26 (2013), 18 pages. <https://doi.org/10.5334/sta.bm>
- [11] Sarah Johnstone and Jeffrey Mazo. 2011. Global Warming and the Arab Spring. *Survival: Global Politics and Strategy* 53, 2 (2011), 11–17. <https://doi.org/10.1080/00396338.2011.571006>
- [12] Colin P Kelley, Shahrzad Mohtadi, Mark A Cane, Richard Seager, and Yochanan Kushnir. 2015. Climate change in the Fertile Crescent and implications of the recent Syrian drought. *Proceedings of the National Academy of Sciences* 112, 11 (2015), 3241–3246. <https://doi.org/10.1073/pnas.1421533112>
- [13] Danil Kononykhin, Mikhail Mozikov, Kirill Mishtal, Pavel Kuznetsov, Dmitrii Abramov, Nazar Sotiriadi, Yury Maximov, Andrey V Savchenko, and Ilya Makarov. 2024. From data to decisions: Streamlining geospatial operations with multimodal globeflowgpt. In *Proceedings of the 32nd ACM International Conference on Advances in Geographic Information Systems*. 649–652.
- [14] Marco Lagi, Karla Z. Bertrand, and Yaneer Bar-Yam. 2011. The food crises and political instability in North Africa and the Middle East. *arXiv preprint arXiv:1108.2455* (2011). <https://doi.org/10.48550/arXiv.1108.2455>
- [15] Tim Lang and Michael Heasman. 2015. *Food Wars: The Global Battle for Mouths, Minds and Markets*. Routledge.
- [16] Georges Lefebvre. 1973. *The Coming of the French Revolution*. Princeton University Press. Translated by R. R. Palmer.
- [17] Aleksander Lukashovich, Vyacheslav Gorchakov, Petr Vorobev, Deepjyoti Deka, and Yury Maximov. 2023. Importance sampling approach to chance-constrained dc optimal power flow. *IEEE transactions on control of network systems* 11, 2 (2023), 928–937.
- [18] Scott M Lundberg and Su-In Lee. 2017. A unified approach to interpreting model predictions. In *Advances in Neural Information Processing Systems*, Vol. 30.
- [19] Kaveh Madani. 2010. Game theory and water resources. *Journal of Hydrology* 381, 3-4 (2010), 225–238. <https://doi.org/10.1016/j.jhydrol.2009.11.045>
- [20] Alexander Marusov, Vsevolod Grabar, Yury Maximov, Nazar Sotiriadi, Alexander Bulkin, and Alexey Zaytsev. 2024. Long-term drought prediction using deep neural networks based on geospatial weather data. *Environmental Modelling & Software* 179 (2024), 106127.
- [21] Yury Maximov, Massih-Reza Amini, and Zaid Harchaoui. 2018. Rademacher complexity bounds for a penalized multi-class semi-supervised algorithm. *Journal of Artificial Intelligence Research* 61 (2018), 761–786.
- [22] Jean-François Maystadt and Olivier Ecker. 2014. Extreme weather and civil war: Does drought fuel conflict in Somalia through livestock price shocks? *American Journal of Agricultural Economics* 96, 4 (2014), 1157–1182.
- [23] Fahimeh Mirzaei-Nodoushan, Omid Bozorg-Haddad, and Hugo A. Loáiciga. 2022. Evaluation of cooperative and non-cooperative game theoretic approaches for water allocation of transboundary rivers. *Scientific Reports* 12 (2022), 3991. <https://doi.org/10.1038/s41598-022-07971-1>
- [24] Mikhail Mozikov, Daria Taniushkina, Aleksandr Bulkin, Yuanhao Liu, Nazar Sotiriadi, Andrey Osiptsov, Roman Sultimov, Ilya Makarov, and Yury Maximov. 2026. Neural Network-Based Approach to Evaluate Convective Hazards Frequency Shift under Climate Change. *Frontiers in Climate* 8 (2026).
- [25] Davide Natalini, Giangiacomo Bravo, and Aled Wynne Jones. 2019. Global food security and food riots – an agent-based modelling approach. *Food Security* 11, 5 (2019), 1153–1173. <https://doi.org/10.1007/s12571-017-0693-z>
- [26] Judea Pearl. 2009. *Causality: Models, Reasoning, and Inference* (2 ed.). Cambridge University Press.
- [27] Judea Pearl. 2010. An introduction to causal inference. *The International Journal of Biostatistics* 6, 2 (2010). <https://doi.org/10.2202/1557-4679.1203>
- [28] Polina Pilyugina, Timofey Chernikov, Maria Smirnova, Alexey Zaytsev, Alexander Bulkin, Evgeny Burnaev, Ilya S Belalov, Nazar Sotiriadi, Albert Efimov, Yury Maximov, et al. 2025. A physics-informed machine learning framework for permafrost stability assessment. *IEEE Access* (2025).
- [29] Jan Rydzak. 2019. Of blackouts and bandhs: The strategy and structure of disconnected protest in India. *Available at SSRN 3330413* (2019).
- [30] Valeriy Shevchenko, Aleksandr Lukashovich, Daria Taniushkina, Alexander Bulkin, Roland Grinis, Kirill Kovalev, Veronika Narozhnaia, Nazar Sotiriadi, Alexander Krenke, and Yury Maximov. 2024. Climate change impact on agricultural land suitability: An interpretable machine learning-based Eurasia case study. *IEEE access* 12 (2024), 15748–15763.
- [31] Ileen N. Streefkerk, Jens de Bruijn, Toon Haer, Anne F. Van Loon, Edisson A. Quichimbo, Marthe Wens, Khalid Hassaballah, and Jeroen C. J. H. Aerts. 2023. A coupled agent-based model to analyse human-drought feedbacks for agropastoralists in dryland regions. *Frontiers in Water* 4 (2023), 1037971. <https://doi.org/10.3389/frwa.2022.1037971>
- [32] Daniil Sukhorukov, Andrei Zakharov, Nikita Glazkov, Katsiaryna Yanchanka, Vladimir Kirilin, Maxim Dubovitsky, Roman Sultimov, Yury Maximov, and Ilya Makarov. 2026. Hierarchical AI-Meteorologist: LLM-Agent System for Multi-Scale and Explainable Weather Forecast Reporting. In *Proceedings of the AAAI 2026 Workshop on AI for Earth Sciences (AI4ES), Poster Track*.
- [33] Roman Sultimov. 2026. Auditable Institutional Coordination with Episodic Memory for Robust Evacuation under Multi-Source Disruption. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. AAMAS 2026.
- [34] Roman Sultimov, Mikhail Mozikov, Dmitrii Abramov, Mariia Kovalchuk, Maksim Malykh, Ilya Makarov, Andrei Osiptsov, Aleksandr Volkov, and Yury Maximov. 2026. RESPOND: Realistic Environment Simulation of Population and Natural Disasters with LLM-Driven Agents. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI 2026)*.
- [35] Roman Sultimov, Aleksandr Volkov, Mile Mitrovic, and Yury Maximov. 2026. LLM-Guided Multi-Agent Evacuation Coordination via Episodic Memory and Cognitive Task Analysis. In *Proceedings of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026)*.
- [36] Daria Taniushkina, Aleksander Lukashovich, Valeriy Shevchenko, Ilya S Belalov, Nazar Sotiriadi, Veronika Narozhnaia, Kirill Kovalev, Alexander Krenke, Nikita Lazarichev, Alexander Bulkin, et al. 2024. Case study on climate change effects and food security in Southeast Asia. *Scientific reports* 14, 1 (2024), 16150.
- [37] Daria Taniushkina, Aleksander Lukashovich, Valeriy Shevchenko, Daniil Sukhurokov, Yulia Melkozerova, Vladislav Pedashenko, Vsevolod Morozov, Roman Sultimov, Ilya Makarov, Aleksandr Volkov, and Yury Maximov. 2026. Bridging Global Climate Solutions and Local Realities: Evaluating Neural Networks for High-Resolution Downscaling. *IEEE Access* (2026). Accepted for publication.
- [38] Rex A. Wade. 2005. *The Russian Revolution, 1917*. Cambridge University Press.
- [39] Maria Waldinger. 2024. “Let them eat cake”: drought, peasant uprisings, and demand for institutional change in the French Revolution. *Journal of Economic Growth* 29, 1 (2024), 41–77. <https://doi.org/10.1007/s10887-023-09230-y> Earlier working paper versions circulated in 2021–2023.
- [40] Yaodong Yang, Rui Luo, Minne Li, Ming Zhou, Weinan Zhang, and Jun Wang. 2018. Mean Field Multi-Agent Reinforcement Learning. In *Proceedings of the 35th International Conference on Machine Learning (ICML) (Proceedings of Machine Learning Research, Vol. 80)*. PMLR, 5571–5580. <https://proceedings.mlr.press/v80/yang18d.html>