Enhancing Agent-Based Models with Inverse Reinforcement Learning for Macroeconomic Crisis Analysis

Elena Lickel

Supervised by Prof. Michael Wooldridge and Prof. Anisoara Calinescu

Department of Computer Science, University of Oxford

UNIVERSITY OF OXFORD

Key message

Crisis-policy simulations embed **value judgements** about what households are assumed to prioritise. These judgements are ethically consequential, yet are rarely made explicit. This project infers a household objective from microdata using **Inverse Reinforcement Learning (IRL)** and embeds it into a macro **agent-based model (ABM)**. The objective is then **held fixed** across crisis and policy counterfactuals, so differences in outcomes reflect policy and shocks.

Agent-based models in macroeconomics

Agent-based models (ABMs) simulate the macroeconomy *from the bottom up* as interacting households, firms, banks, and government operating under explicit constraints and market mechanisms. Their key advantage is that aggregate outcomes **emerge** from **heterogeneity** and **feedback loops**. These mechanisms matter in financial crises, where balance-sheet stress propagates through credit, labour, and goods markets.

State of the art. Recent macro ABMs increasingly replace hand-coded household rules with **(multi-agent) reinforcement learning (MARL)** to capture adaptation and strategic interaction in changing environments. This enables richer behavioural responses and endogenous adjustment to policy rules.

- **What this enables:** realistic, adaptive behaviour without specifying decision heuristics by hand.
- **Open challenge:** learning-based ABMs still rely on **hand-designed rewards** that are hard to justify, compare, or transfer across regimes, making conclusions sensitive to reward specification.
- **Proposed solution:** infer the household reward function through Inverse Reinforcement Learning (IRL).

Ethical relevance

Integrating **IRL** into macroeconomic **ABMs** supports ethically defensible policy simulation by grounding household behaviour in empirically inferred priorities rather than hand-designed assumptions.

- **Distributional justice:** moves beyond representative agents to quantify who bears costs and who benefits across heterogeneous households.
- **Safety and robustness:** reduces the simulation-to-reality gap by stress-testing policies against behavioural and shock variation, mitigating brittle recommendations.
- **Regime-change validity:** infers underlying motivations that can adapt under new policy regimes, addressing Lucas-style concerns.
- **Transparency and bias:** makes value judgements explicit and auditable; enables subgroup checks to mitigate biased inference.
- **Preference alignment:** supports welfare analysis reflecting revealed trade-offs under constraints, not only aggregate metrics.

Approach in one figure

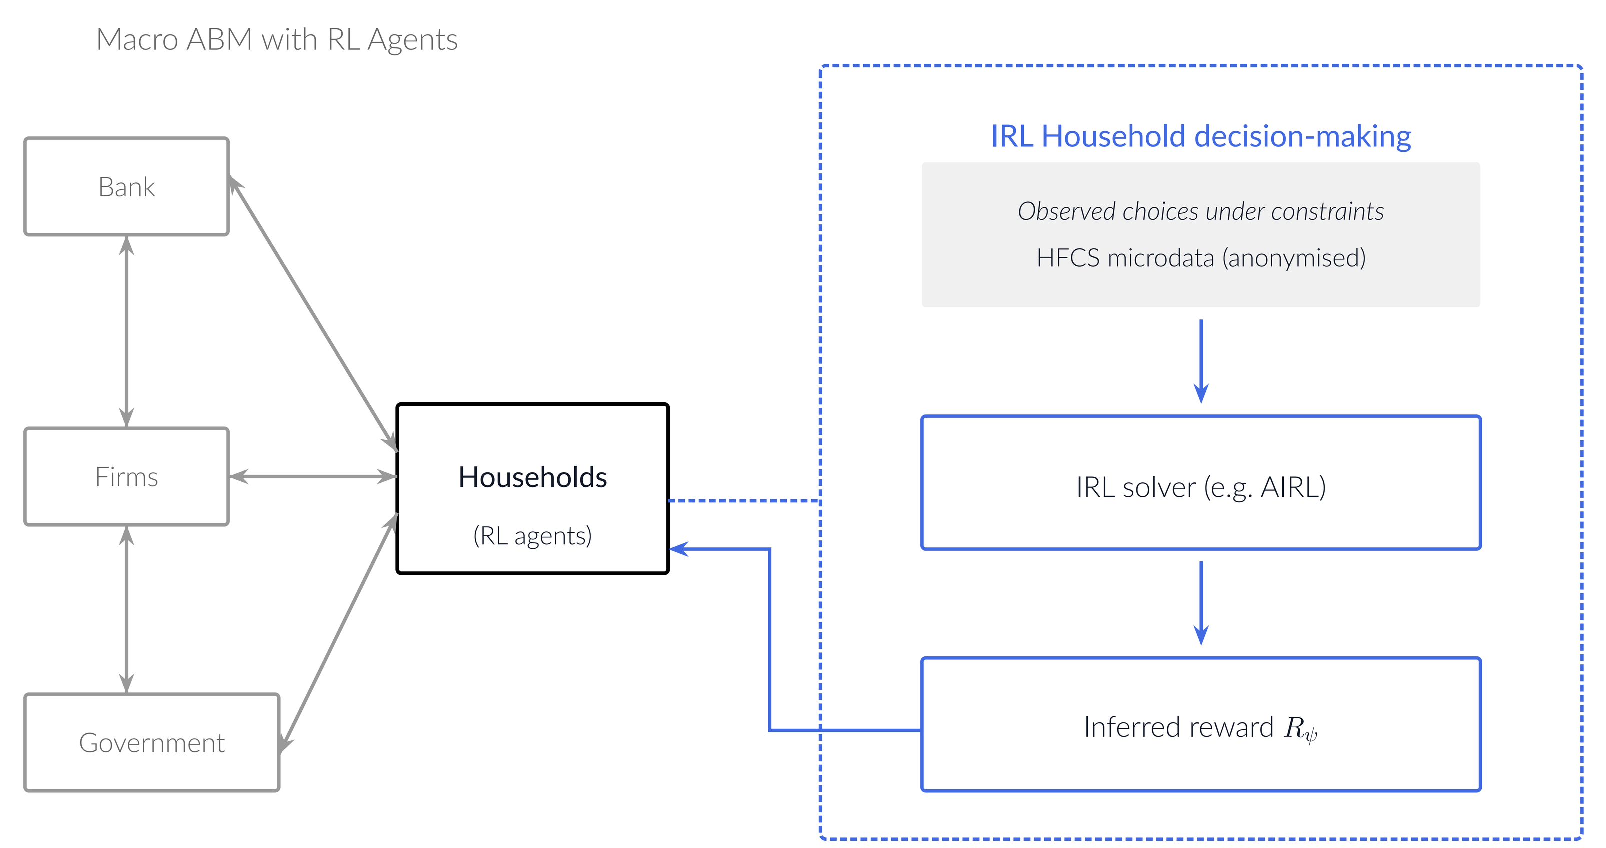

Methodological approach

Baseline crisis ABM. Build on an established macro ABM with interacting households, firms, banks, and government (e.g. [1]) to retain validated crisis mechanisms and comparability with prior work.

Household decision model. Represent households as adaptive agents with state s (income, wealth, debt, housing, prices, constraints) and actions a (consumption–saving, borrowing/repayment, portfolio/housing adjustments), i.e. a policy $\pi(a \mid s)$ under explicit constraints.

Objective inference (AIRL). Use **Adversarial IRL** to infer a reward $R_\psi(s, a)$ that rationalises observed choices under constraints.

- **Transferable rewards:** targets reward structure disentangled from environment dynamics, supporting regime and policy transfer.
- **Multi-agent fit:** developed for interacting-agent settings with non-stationarity, matching ABM feedback loops (firms, banks, policy).
- **Practical at scale:** avoids requiring a fully specified transition model, enabling inference in high-dimensional macro state spaces.

Counterfactual protocol. Infer R_ψ once, embed it in the ABM, and **hold it fixed.** Counterfactuals change only shocks, prices, constraints, and policy rules; households adapt actions to new conditions, not their objective, enabling clean attribution to policy and shocks.

Data and data ethics

Data. The project uses the **ECB Household Finance and Consumption Survey (HFCS)**: harmonised, **anonymised** microdata on household balance sheets and constraints across euro-area countries. Key inputs include **income, liquid assets, debt and debt service, housing tenure/value, employment status**, and core demographics used for heterogeneity.

Governance and safeguards.

- access via formal application and approved secure-use conditions,
- **data minimisation**: only variables required for modelling household decisions,
- **purpose limitation**: objective inference and subgroup-level policy evaluation (no targeting),
- no re-identification attempts; results reported only in aggregate/subgroups.

Contribution

- **Objective inference for macro ABMs:** integrate IRL/AIRL to estimate a household reward $R_\psi(s, a)$ from microdata rather than hand-crafting preferences.
- **Counterfactual discipline:** fix R_ψ across crises and policy regimes to separate policy effects from behavioural re-specification.
- **Distributional crisis analysis:** enable subgroup/decile evaluation of interventions within a heterogeneous ABM grounded in micro evidence.
- **Interdisciplinary bridge:** combine macroeconomics, multi-agent learning, and ethical governance of data-driven policy models (with INET Oxford).

Takeaway

Inferring and fixing household objectives in crisis ABMs makes policy simulations more **transparent**, more **robust across regimes**, and more **accountable** in their distributional conclusions.

References

[1] S. Brusatin, M. Tedeschi, M. Gallegati, and D. Delli Gatti, "Simulating the economic impact of rationality through reinforcement learning and agent-based modelling," in *Proceedings of the ACM International Conference on AI in Finance (ICAIF)*, 2024.

[2] J. Fu, K. Luo, and S. Levine, "Learning robust rewards with adversarial inverse reinforcement learning," *arXiv preprint arXiv:1710.11248*, 2018.

[3] L. Yu, J. Song, and S. Ermon, "Multi-agent adversarial inverse reinforcement learning," in *Proceedings of the International Conference on Machine Learning (ICML)*, 2019.

elena.lickel@cs.ox.ac.uk