

# Learning and Calibrating Heterogeneous Bounded Rational Market Behaviour with Multi-Agent Reinforcement Learning

Benjamin Patrick Evans  
JP Morgan AI Research  
London, United Kingdom  
benjamin.x.evans@jpmorgan.com

Sumitra Ganesh  
JP Morgan AI Research  
New York, USA  
sumitra.ganesh@jpmorgan.com

## ABSTRACT

Agent-based models (ABMs) have shown promise for modelling various real world phenomena incompatible with traditional equilibrium analysis. However, a critical concern is the manual definition of behavioural rules in ABMs. Recent developments in multi-agent reinforcement learning (MARL) offer a way to address this issue from an optimisation perspective, where agents strive to maximise their utility, eliminating the need for manual rule specification. This learning-focused approach aligns with established economic and financial models through the use of rational utility-maximising agents. However, this representation departs from the fundamental motivation for ABMs: that realistic dynamics emerging from bounded rationality and agent heterogeneity can be modelled. To resolve this apparent disparity between the two approaches, we propose a novel technique for representing heterogeneous processing-constrained agents within a MARL framework. The proposed approach treats agents as constrained optimisers with varying degrees of strategic skills, permitting departure from strict utility maximisation. Behaviour is learnt through repeated simulations with policy gradients to adjust action likelihoods. To allow efficient computation, we use parameterised shared policy learning with distributions of agent skill levels. Shared policy learning avoids the need for agents to learn individual policies yet still enables a spectrum of bounded rational behaviours. We validate our model’s effectiveness using real-world data on a range of canonical  $n$ -agent settings, demonstrating significantly improved predictive capability.

## KEYWORDS

Multi-agent reinforcement learning; Bounded rationality; Market simulation; Agent-based modelling; Skill heterogeneity

### ACM Reference Format:

Benjamin Patrick Evans and Sumitra Ganesh. 2024. Learning and Calibrating Heterogeneous Bounded Rational Market Behaviour with Multi-Agent Reinforcement Learning. In *Proc. of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2024)*, Auckland, New Zealand, May 6 – 10, 2024, IFAAMAS, 10 pages.

## 1 INTRODUCTION

Agent-based models (ABM) have achieved significant success in various domains, including business, epidemiology [33], economics,

and finance [5, 25]. However, despite these achievements, criticisms persist within different communities [2], notably in economics [37, 39, 61], primarily due to concerns regarding the decision-making rules in these systems. Frequently, these rules are manually specified heuristics, placing substantial reliance on the modeller’s judgement as the simulation results and validity depend upon the specific behavioural rules utilised [47, 61]. On the other hand, adaptive agents that optimise a utility function find greater acceptance across disciplines, as the agents’ behaviour is automatically derived in a principled manner.

Hence, introducing adaptive and learning agents into ABM, while allowing for heterogeneity and bounded rationality, could alleviate these concerns and improve the realism of the models. Recent progress in reinforcement learning (RL) helps to bring this closer to reality [1, 14, 59]. However, some significant challenges must be addressed before this can happen. In this work, we address the following essential question: *How can we learn heterogeneous bounded rational behaviours in an ABM?*

To address this question, we introduce a novel multi-agent RL (MARL) approach where agents exhibit skill heterogeneity [52], constrained by their strategic processing costs. In the limit, where these processing costs  $\rightarrow 0$ , perfectly rational mutually consistent equilibrium can be approximated. With uniform prior beliefs and homogenous processing costs among agents, quantal response type equilibrium can be approximated. With processing costs  $\rightarrow \infty$ , agents act based on their prior beliefs, e.g., driven by heuristics [27] or biases [17]. However, more generally, we can model a wide range of realistic behaviour with heterogenous agent bounds. This framework aims to enhance the simulation of complex social systems, which differs from many MARL methodologies focused on learning optimal behaviours. Instead, the work aligns with the literature on ABMs, focusing on understanding the resulting realistic dynamics emerging from human decision-making.

**Contributions:** We propose an approach to effectively learn heterogeneous agent skill levels (exhibiting diverse deviations from prior beliefs) within a MARL framework. We demonstrate the efficacy of the approach across several fundamental multi-agent economic environments. The proposed approach offers substantially enhanced accuracy in predicting human decisions (along with the subsequent dynamics) in controlled experiments compared to current state-of-the-art RL approaches and other equilibrium benchmarks. To improve efficiency, we utilise agent supertypes [63] and shared policy learning to learn diverse bounded rational behaviours. Heterogeneity is introduced by varying strategic processing costs of agents, measured through regularised divergences from their prior beliefs. This approach is general, expanding upon a new MARL framework for modelling complex systems, *Phantom* [3].



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

## 2 RELATED WORK

While MARL algorithms have made significant advances in approximating equilibria within complicated environments [49], a limitation is that most prevailing methodologies assume agents to be perfectly rational. This assumption is often overly stringent when simulating complex social systems [4], and may miss crucial real-world dynamics [7]. To broaden the applicability of MARL for modelling complex systems, our objective is to account for agent heterogeneity and bounded strategic abilities. In doing so, we make the connection with ABMs while automating some of the difficult modelling design (e.g., defining the agent behavioural rules) using RL. Recent surveys covering RL techniques in ABM have emphasized the usefulness of learning agent behaviours [50, 60, 67].

Behavioural economics has developed more realistic models of decision-making than the traditional *homo economicus* perfectly rational representative agent [40]. Instead, these models operate under the framework of bounded rationality. One prominent approach for relaxing the strict perfectly rational Nash equilibrium (NE) assumption and incorporating bounds in reasoning is the Quantal Response Equilibrium (QRE) [42, 43], which allows for deviations from optimal responses and the possibility of erroneous play. Such approaches generally feature consistent (and common) beliefs among agents, for example, by having the same processing costs across the population (mutual consistency). However, the need to consider agent heterogeneity has been stressed due to the ability to "bring about new outcomes not foreseeable from analysis of the homogeneous dynamics" [44]. Given that populations inherently encompass a spectrum of behaviours and beliefs, a "representative" agent often proves insufficient [30]. Recognizing the importance of strategic diversity, various extensions have been developed to accommodate a range of agent behaviours [51, 52], relaxing the mutual rationality and mutual consistency assumptions [12, 19]. Relaxing mutual consistency is beneficial for multi-agent settings, allowing for a population characterized by varying levels of strategic *boundedness* [38, 48, 65]. However, the computability of such game-theoretic models is often limited to relatively simple domains, prompting the use of approximation methods (such as MARL) within more complicated environments [68].

Despite the achievements in behavioural game theory and RL, a gap persists in combining these methods with ABMs, e.g. for learning the decision-making rules. While RL approaches often prioritize convergence towards rational equilibrium, this focus contrasts with the primary objective of the ABM community: understanding the properties emerging from heterogeneous boundedly rational agents. To the best of our knowledge, the seminal work on capturing realistic bounded behaviours in MARL is that of [45], based on rational inattention (RI). The authors of [45] underline that the existing body of work "fails to address the modeling of bounded rationality for more accurate [MARL] simulations," proposing an innovative framework to address this limitation. Our work differs in some important ways. Specifically, we remove the difficult processing cost estimation required in [45] (as discussed in Section 5) and allow for arbitrary prior beliefs for encoding behavioural biases. Additionally, we introduce agent skill heterogeneity, learnt through regularised policies, and propose an approach for efficiently calibrating these policies to real world dynamics, all features yet to be considered.

## 3 PROPOSED APPROACH

We introduce a novel MARL approach to effectively model a diverse range of bounded rational behaviours (or varying *skill*). This approach proves valuable for calibrating ABMs to real-world systems through learning regularised policies and, under limiting cases, establishing links to various existing equilibrium solution concepts.

Our general formulation is as follows. We focus on  $N$ -agent systems, where each agent  $i \in N$  seeks to maximise their reward (or *utility*) function  $U_i$  by taking actions from their action space  $a \in A$ . Importantly, these agents may not act perfectly rationally. The system is characterised by a state space  $S$ , and agents possess a (potentially partial) observation of the current state  $s_i$  along with prior beliefs about their potential actions  $q_i$  (a probability distribution over the action space). The behaviour of each agent is governed by their policy  $\pi_i$ , which is a mapping from states to a distribution over actions. Agents act based on their policy  $a_i \sim \pi_i$ , receiving reward  $U_i(a_i, s_i)$ . Again, these actions may not be perfectly rational and instead may be *satisficing* [9, 56].

### 3.1 Components

**3.1.1 Reward.** In standard RL, agents aim to learn an optimal policy that maximises their expected cumulative discounted reward  $U$ :

$$\pi_i^*(a|s_i) = \max_{\pi_i} \mathbb{E}_{\pi_i} \left[ \sum_{t=0}^{\infty} \gamma^t U(a_t|s_{i,t}) \right] \quad (1)$$

However, often, in real world systems, the agents we seek to model are boundedly rational, driven by prior beliefs and a limited amount of processing power to improve upon these priors. To address these considerations, we incorporate generic limitations in agents' reasoning abilities by reformulating the maximisation problem into a constrained one<sup>1</sup>:

$$\max_{\pi} \mathbb{E}_{\pi_i} \left[ \sum_{t=0}^{\infty} \gamma^t U(a_t|s_{i,t}) \right] \quad \text{subject to} \quad I(\pi_i, s_{i,t}, q_i) < \bar{I} \quad (2)$$

where agents maximise  $U$  while adhering to a constraint  $\bar{I}$  on their processing costs  $I$ . The processing cost restricts how far the learnt policy  $\pi_i$  can diverge from the prior belief  $q_i$ . We can equivalently reformulate Eq. (2) as the maximisation of a modified reward:

$$\pi_i^\lambda(a|s_i) = \max_{\pi_i} \mathbb{E}_{\pi_i} \left[ \sum_{t=0}^{\infty} \gamma^t (U(a_t|s_{i,t}) - \lambda I(\pi_i, s_{i,t}, q_i)) \right] \quad (3)$$

where  $\lambda$  controls the strength of regularisation, modulating the *boundedness* (or *skill*) of the agent. As  $\lambda \rightarrow \infty$ , the agent is entirely driven by their prior beliefs (performing no strategic reasoning), whereas as  $\lambda \rightarrow 0$ , the agent is unbounded and approximates rational behaviour. Eq. (3) makes the formulation general and compatible with existing RL algorithms without requiring in-depth modification to the loss or optimisation process.

**3.1.2 Processing Costs.** Quantifying information processing costs in a generalised manner is desirable, as this enables compatibility

<sup>1</sup>We maintain exponential discounting here. However, hyperbolic discounting might provide a better alignment with human decision-making [53], although current RL approaches have not yet exhibited significant differences between the two [22].

with existing optimisation algorithms. Following recent achievements in constraining agent decision-making using information-theoretic costs [19], we adopt a similar approach. This information-theoretic treatment abstracts the underlying causes of such constraints, allowing a focus on learning behaviour without necessitating an in-depth understanding of the specific psychological factors at play. From an optimisation standpoint, this is advantageous, as the process remains independent of the particular details of how decisions are formulated [57].

One of the most common information-theoretic constraints is an entropy constraint, e.g.  $H(\pi_i) = -\sum p(\pi_i) * \log(p(\pi_i))$ , restricting deviations from uniform behaviour. For example, (the logit form of) QRE can be seen as maximisation under an entropy constraint. However, much research has shown the usefulness of incorporating arbitrary prior beliefs (not just uniform) [18], motivating extensions that measure the divergence from an arbitrary prior distribution based on the Kullback-Leibler (KL) divergence  $D_{KL}$  [46].

KL regularisation has shown success in relevant domains for formally capturing these costs. For example, [36] demonstrates the usefulness of a penalty for minimising  $D_{KL}$  from expert policies in tree search and [58] analyses  $D_{KL}$  in a RL context for improving convergence in two-player games. However, neither of these considers the domain we propose here for better capturing human-like play in complex multi-agent social systems.

Specifically, we propose using the following information processing costs

$$I(\pi, s, q) = D_{KL}(\pi \parallel q) = \sum_{a \in A} \pi(a|s) \log \frac{\pi(a|s)}{q(a|s)} \quad (4)$$

to constrain  $\pi_i$  from diverging too far from agents' prior beliefs  $q$  at each state, limiting their strategic abilities. Eq. (4) can also be seen as equivalent to enforcing an  $H$  constraint when assuming the prior beliefs are uniform (making connections with QRE, as discussed in Section 5). Additionally, with this representation, the contribution of a specific action  $a$  to the divergence can be identified, e.g.,

$$I_a(\pi, s, q) = \pi(a|s) \log \frac{\pi(a|s)}{q(a|s)} \quad (5)$$

meaning the adjustment to the reward function in Eq. (3) can be directly linked to  $a$  rather than  $\pi_i$ , i.e.,

$$\pi_i^\lambda(a|s) = \max_{\pi_i} \mathbb{E}_{\pi_i} \left[ \sum_{t=0}^{\infty} \gamma^t (U(a_t|s_{i,t}) - \lambda I_a(\pi_i, s_{i,t}, q_i)) \right] \quad (6)$$

which is advantageous for optimisation purposes. We use this formulation throughout. As we sample more actions from this policy, we would approximate  $D_{KL}$  as  $I(\pi, s, q) = \sum_{a \in A} I_a(\pi, s, q)$ .

In RL, information-theoretic regularisation is often employed to enhance the convergence or robustness of algorithms. For instance, Proximal Policy Optimization (PPO) utilises a  $D_{KL}$  penalty term to prevent excessively large changes in the policy during training steps and improve the convergence. Similarly, the Soft Actor-Critic (SAC) algorithm employs  $D_{KL}$  in its policy improvement step, limiting divergence from the previous Q-function [31]. Moreover, Maximum Entropy RL introduces an entropy term to enhance exploration, convergence, and robustness [20, 21]. In contrast, our approach restricts divergence from an arbitrary prior belief  $q$  to reflect the constraints in information processing present during human decision-making

rather than being aimed at improving the algorithm's convergence. These prior beliefs  $q$  (also called "magnets" [58] or "anchors" [36]) may change throughout training and inference (e.g. with updated information) and can take many forms, for example, demonstrating bias towards certain actions, encoding heuristics, averaging over past decisions, or preferring historically well-performing actions.

**3.1.3 Heterogeneous Behaviours.** Effectively learning behaviours that capture the diversity of the population's decision-making is crucial for integration into agent-based simulations.

Two initial approaches could be employed for capturing heterogeneous skills of agents: Firstly, learning optimal (homogeneous) behaviours  $\pi^*$  as in Eq. (1) and applying heterogeneous bounds at inference, e.g.,  $\pi_i = \pi^* + \eta_i$ , where  $\eta_i$  is a noise term. Secondly, individual learning with heterogeneous  $\lambda$ 's, i.e.,  $\pi_i = \max_{\pi} \mathbb{E}_{\pi} [\sum_{t=0}^{\infty} \gamma^t (U(a_t|s_{i,t}) - \lambda_i I(\pi, s_i, q_i))]$ . In the following section, we describe why these approaches are insufficient before proposing an alternative that overcomes these limitations.

*Post-hoc bounds at inference.* One potential approach involves optimising  $U$  without including information processing costs during training (eliminating boundedness) to learn  $\pi^*$ . Subsequently, this boundedness parameter is applied only to the learned policies during inference  $\pi_i = \pi^* + \eta_i$ , where  $\eta_i$  is the noise term. For example, [8] applies dropout during simulation, and noisy introspection applies noise into the decisions relaxing the equilibrium requirement [28]. Such methods introduce a range of skill levels in action execution, e.g. through introducing noise  $\eta_i$  into the action selection process. However, if heterogeneous bounds were implemented in this manner, agents would not learn how to adapt to the behaviour of other bounded agents, as the best response to the optimal policy is not necessarily the best response to a noisy policy  $BR(\pi^*) \neq BR(\pi_i)$ . To illustrate this point, consider a simple rock-paper-scissors (RPS) environment. In RPS, the perfectly rational equilibrium policy is  $\pi^* = \{p(R), p(P), p(S)\} = \{\frac{1}{3}, \frac{1}{3}, \frac{1}{3}\}$ . However, if agent 2s policy is instead fixed as  $\pi_2 = \{\frac{0}{3}, \frac{0}{3}, \frac{3}{3}\}$  (e.g. they are boundedly rational and biased towards playing S), the rational best response for agent 1 is  $\pi_1 = BR(\pi_2) = \{\frac{3}{3}, \frac{0}{3}, \frac{0}{3}\}$ . Had agent 1 not observed  $a_2 \sim \pi_2$  during training, they would not have learned to exploit  $\pi_2$ . While a simple example, this consideration becomes pivotal in demonstrating the emergence of auto-curricula [6]. To capture the interplay among heterogeneously skilled agents, the notion of boundedness must be present throughout the learning process rather than only during inference.

*Individual Learning.* An alternative approach involves assigning heterogeneous processing penalties  $\lambda_i$  to each agent  $\pi_i = \pi^{\lambda_i}$  in Eq. (3), e.g. with heterogeneous logit responders [30]. In this scenario, any standard MARL algorithm could be employed, wherein all agents, each governed by their unique constraint, strive to optimise their rewards, adjusting their behaviours in response to the observed outcomes. However, this method would prove inefficient due to the necessity of learning  $N$  individual policies and calibrating  $N$  different processing costs  $\lambda_i$  (one for each agent). This inefficiency becomes a crucial concern as ABMs often have a large  $N$ . Additionally, the learnt policies would not generalise across different  $\lambda_i$ , requiring retraining each time a new  $\lambda_i$  is introduced.

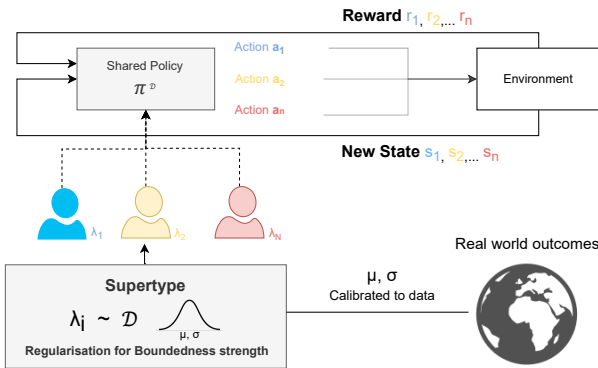
It becomes clear we need a scalable alternative that can deal with the heterogeneous boundedly rational behaviour of agents.

### 3.2 Shared Policy Learning

Rather than learning individual policies with heterogeneous  $\lambda_i$ , all aiming to solve Eq. (3), we wish to learn a generalised policy  $\pi(\dots | s_i, i, \lambda_i, q_i)$ . This representation treats agents' prior beliefs and processing resources as part of the observation space (and for simplicity of notation, we will use  $\pi_i(\dots | s_i) = \pi(\dots | s_i, i, \lambda_i, q_i)$ ), enabling generalised policy learning based on these state observations. This formulation provides a way of efficiently representing a diverse population of agents  $i \in N$  with varying bounds in strategic reasoning abilities  $\lambda_i$  through a single parameterised policy with an augmented observation space.

However, calibrating  $\lambda_i$  remains an important issue. While calibrating  $\lambda_i$  to each agent  $i$  may seem ideal, this is computationally impractical with larger  $N$ , and additionally, could lead to overfitting to specific behavioural parameters due to the large number of required parameters ( $N$ ). Furthermore, in practice, as  $\lambda_i$  are unobserved, assigning these values exactly is difficult.

We adopt an alternative approach to address this challenge by assigning individual strategic processing resources as samples from a probability distribution  $\lambda_i \sim \mathcal{D}$ , addressing the uncertainty of the agents' exact  $\lambda_i$  values and keeping the number of free parameters low. Any  $\mathcal{D}$  could be utilised (and the proposed approach is agnostic to the particular distribution used), but here, we employ the Gaussian distribution  $\mathcal{D} = \mathcal{N}(\mu, \sigma)$ , where  $\mu$  controls the mean processing costs, and  $\sigma$  the heterogeneity. This way, we only need to calibrate the parameters  $\mu$  and  $\sigma$  (rather than  $N$  separate parameters), which is typically  $\ll N$ . For instance, when dealing with  $N = 100$  agents, we are calibrating just 2 parameters instead of 100, helping to avoid overspecification. During learning, this sampling approach allows for interpolation across a range of  $\lambda_i$ , reducing the computational complexity and enforcing a "smoother" policy. This smoothness arises from observing many different behaviours during training, resulting in a more robust policy forced to interpolate across  $\lambda_i$  values, reducing the potential of overfitting.



**Figure 1: Proposed Approach: Shared policy learning with heterogeneous bounds through agent supertypes.**

To learn a generalised policy  $\pi^{\mathcal{D}}$  for  $\lambda_i \sim \mathcal{D}$ , we use agent supertypes [63], enabling efficient scaling through shared policy

learning, while still capturing a range of behaviours. Agent supertypes have exhibited promise, particularly in applications such as calibrating rational behaviour in over-the-counter markets [62]. However, the potential for incorporating heterogeneous strategic reasoning skills to better approximate human decision-making has yet to be explored. Here, we propose a novel approach to extend the use of supertypes to capture diverse bounded rational behaviours based on the regularised policies introduced in Eq. (3).

Under the proposed approach, a regularised policy for the supertype is established  $\pi^{\mathcal{D}}$ , which is exposed to different regularisation strengths  $\lambda_i \sim \mathcal{D}$  throughout training.  $\pi^{\mathcal{D}}$  learns to extrapolate over the regularisation strengths, reducing the number of policies to train while still enabling heterogeneous behaviour. Through this process, agents learn to adapt their behaviour in response to the varying processing resources across the agent population, accounting for potential auto curricula. The inputs to the supertype  $(\mu, \sigma)$  are calibrated such that the simulation outcome closely matches the real world dynamics (from the calibration data). An overview of this high-level process is depicted in Fig. 1.

The shared policy  $\pi^{\mathcal{D}}$  takes an agent's id  $i$ , processing resources  $\lambda_i$ , and prior beliefs  $q_i$  as inputs (as components of  $s_i$ ). Including  $i$  facilitates the learning of (potentially) competitive behaviour between agents of the same supertype. By adjusting the input  $\lambda_i$  values,  $\pi^{\mathcal{D}}$  can effectively demonstrate a spectrum of skill levels, while only needing to learn a single (generic) policy. Allowing for arbitrary  $q_i$  accounts for the effect of various prior beliefs.

The underlying assumption of this supertype approach is that all agents in the supertype have the same  $U$  function; however, they possess varying levels of skill in maximising  $U$ . Given the inherent uncertainty surrounding the precise nature of agent decision-making, a compelling case is made for capturing a spectrum of regularised behaviours. The regularisation offers dual advantages: firstly, it enables deviations from perfect rationality in agent behaviour; secondly, it effectively encompasses uncertainties from both the modeller's perspective and the agents being modelled (i.e., uncertainty in the model's formulation and the agents' decision processes) [18]. This motivation aligns with the use of bounded rationality in situations characterised by fundamental uncertainty of the agent [26] and also helps to address concerns regarding modeller judgement (e.g. model misspecification) by permitting a range of information-constrained behaviour [54].

## 4 EMPIRICAL RESULTS: $n$ -AGENT SETTINGS

To verify that the proposed approach can capture a range of interesting behaviour not predicted by the analytically derived equilibrium or standard state-of-the-art MARL approaches, we compare the predictions from the proposed model against these approaches on a range of canonical  $n$ -agent economic environments involving human participants.


### 4.1 Process Overview

We assess the performance of the proposed approach in three well-established multi-agent economic environments: supply chains,

<sup>2</sup>The Cournot competition environments each carry  $\frac{1}{2}$  weight when computing the average to address the interdependence and avoid biasing the average ranking (although in this case, the rankings would not change without such a weighting)

**Table 1: 5x2-fold validation results for each environment. Each cell displays the root mean squared errors as mean  $\pm$  standard deviation, along with (rankings) for between-environment comparison [15]. The last row presents the average rank<sup>2</sup>. Lower rankings indicate better performance.**

	Rational	MARL	Proposed
Supply Chain	0.33 $\pm$ 0.004 (2.5)	0.33 $\pm$ 0.004 (2.5)	0.02 $\pm$ 0.005 (1)
Cournot			
- Duopoly	0.16 $\pm$ 0.001 (3)	0.13 $\pm$ 0.001 (2)	0.04 $\pm$ 0.001 (1)
- Triopoly	0.16 $\pm$ 0.002 (3)	0.15 $\pm$ 0.002 (2)	0.03 $\pm$ 0.001 (1)
Cobweb	0.02 $\pm$ < 0.001 (2)	0.03 $\pm$ < 0.001 (3)	0.01 $\pm$ < 0.001 (1)
<b>Rank</b>	2.5	2.5	<b>1</b>



**Figure 2: Triopoly calibration results for values of the boundedness parameter  $\mu$  and heterogeneity parameter  $\sigma^*$**

oligopolies, and cobweb markets. To validate our approach, we leverage laboratory experiments conducted in each setting, comparing the predictions with actual human behaviour. We compare the proposed approach with analytically derived solutions and a state-of-the-art MARL algorithm (PPO). In each case, we perform repeated 5x2cross-fold validation [15] to estimate the generalisation ability, ensuring the models do not overfit to the calibration data. The squared  $L_2$ -loss function (the mean squared error) is used as our performance metric [16]. In the presented tables, we report the root mean squared error for interpretability. To facilitate comparisons across environments, we use rankings based on resulting errors (where the lowest error receives rank=1, and ties are split by using the average rank had there been no ties) [15]. To calibrate our model, we perform a grid search over  $\mu, \sigma$  values, choosing  $\mu, \sigma$  with the lowest training error for use on the unseen test set. The optimisation process never sees the test data. Additional experiment details are given in the online appendix<sup>3</sup>.

**4.1.1 Calibration.** The calibration results for one environment (triopolies), displaying the values of the heterogeneity ( $\sigma^*$ ) and boundedness ( $\mu$ ) parameters, are shown in Fig. 2. Similar plots are available for all environments in the online appendix. The proposed approach offers the flexibility to incorporate perfect rationality or homogeneity by setting  $\mu = 0$  (removing bounds) or  $\sigma^* = 0$  (removing heterogeneity). However, it’s noteworthy that the optimal values never align with  $\mu = 0$  or  $\sigma^* = 0$ , highlighting the usefulness of both heterogeneity and processing costs across all three environments. Furthermore, since homogeneity and unbounded reasoning can be considered special cases of our proposed approach,

<sup>3</sup>Online appendix available at: <https://arxiv.org/abs/2402.00787>

this eliminates the need to determine such assumptions *a priori* (as often required in many existing methods). Instead, our approach enables the calibration of these properties based on the specific environments of interest.

## 4.2 Results

The out-of-sample performance of each algorithm is compared in Table 1. The proposed approach achieves the highest accuracy across all three environments, resulting in the best overall rank. The state-of-the-art (standard) MARL approach and the analytically derived rational case generally perform equivalently, indicating that the MARL algorithm approximated the true rational equilibrium well. However, both the alternatives performed poorly in capturing the experimental data, demonstrating that rationality and homogeneity are too strict of an assumption even in these relatively simple multi-agent settings. These results motivate the relaxation of perfect rationality and the introduction of skill heterogeneity when using MARL to model complex systems.

To better understand the results and the reason for the improved capabilities of the proposed approach, we analyse each environment in more detail. For each environment, we begin with a brief description, before presenting the results.

### 4.2.1 Supply Chains.

*Description.* The supply chain environment is a capacity allocation problem with a single good with cost  $c$  and price  $p$ . There is one supplier with a limited capacity  $K$ , and  $I$  retailers. Each retailer  $i$  makes a request  $0 < x_i \leq X, x_i \in \mathbb{Z}$ , and the supplier responds by offering  $y_i$ . Retailers are allocated goods proportionate to their request  $y_i \propto x_i$ :

$$y_i = K \times \frac{x_i}{\sum_{j \in I} x_j} \quad (7)$$

inducing the potential for (rationally) inflated order sizes to ensure the required quantities are met. Each retailer receives a fixed demand  $D > \frac{K}{I}$ , i.e., resources are limited. The reward is given by


$$U(x_i) = D \times (p - c) - \omega \times \max(y_i - D, 0) - s \times \max(D - y_i, 0) \quad (8)$$

where  $\omega$  is the wastage cost, and  $s$  is the shortage cost ( $s = p - c$ ). The rational (Nash) solution to this task, irrespective of  $\omega, s$  (when in the limited capacity case of  $K < I * D$ ), is for retailers to submit their maximum request  $X$ , each retailer then receiving  $y_i = \frac{K}{I}$  units due to the proportional allocation. Any lower of a request would result in the retailer receiving  $y_i < \frac{K}{I} < D$ .

We utilise the experimental results of [11], with  $I = 2, \omega = 2$ , and  $s = 5$ . There were 30 subjects, composed of university students, randomly paired in 30 repeated decision rounds to make a game with two retailers in each round. The capacity is  $K = 90$ , with each retailer receiving demand  $D = 50$  and able to make a maximum request of  $X = 100$ . As  $2D > K$ , we are faced with limited capacity.

*Results.* The experimental results are displayed in Fig. 3, showing substantial deviations from purely rational (Nash) play. The NE is to request the maximum  $x_i = X = 100$ . The standard MARL approach learns the NE here; however, this is a poor predictor of what happens experimentally (Fig. 3). Experimentally, the most commonly occurring requests are in the 60 – 80 range, far lower than the NE. The proposed approach is a very good fit for the





**Figure 3: Supply Chain.** Experimental data from [11] are shown as grey bars. The proposed approach is shown in orange. The standard MARL approach is shown as the dashed purple line, and the NE is denoted by the black bar.

experimentally observed behaviour, capturing the overall trend, demonstrating that subjects have varying strategic bounds, giving rise to a range of outcomes not predicted by a rational representative agent, and helping to motivate the bounded rationality assumptions.

#### 4.2.2 Cournot Oligopoly.

*Description.* The Cournot competition is an environment modelling oligopolies in a market. In a Cournot market,  $K$  firms must simultaneously choose what quantities  $q_i \in \mathbb{Z}$  of a homogenous good to produce. The reward for firm  $i$  depends on the market price  $p$  of the good and the individual  $q_i$ , i.e.:

$$U_i = p \times q_i \quad (9)$$

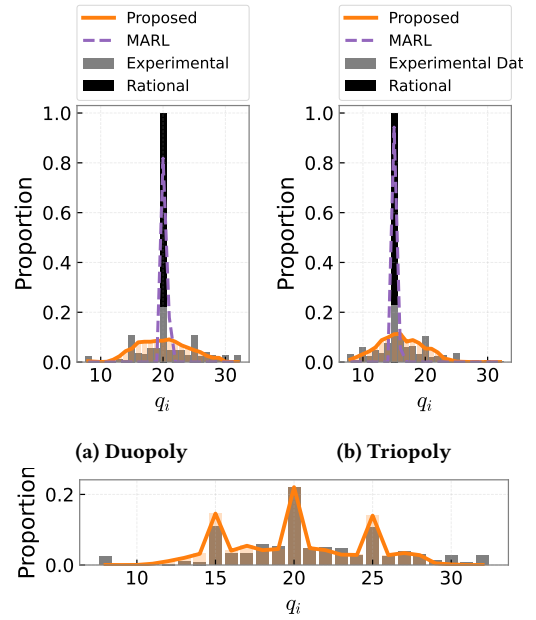
where  $p$  is determined by the total production of all goods:

$$p = A - B \times \sum_{k=1}^K q_k \quad (10)$$

We use the experimental data from [23, 32], for duopolies and triopolies (experiments 7,8,9,10 from [23]). Following [32], we group experiments 7, 10 (duopoly) together and experiments 8, 9 (triopoly) together. In these experiments,  $A = 2.4$ ,  $B = 0.04$ , and  $8 \leq q_i \leq 32$ . There were 64 participants for the duopoly experiments, and 66 for the triopoly, composed of university students.

*Results.* The results for duopolies (triopolies) are presented in Fig. 4a (Fig. 4b). With the experimental data, we see a significant deviation from both the rational behaviour and the standard MARL predictions. The unique NE for duopolies and triopolies is 20 and 15 respectively. While these actions are the most common in each case, these occurrences comprise  $\leq 20\%$  of the total decisions, and the remaining  $\approx 80\%$  are sub-optimal decisions (under the assumption of mutual rationality). Again, the proposed model is a good fit for the experimental data in both duopolies and triopolies, capturing this significant deviation from the optimal choice while still capturing the maximal peak from the experimental data.

Under the processing cost constraint, agents choose actions proportionate to the expected reward and the level of regularisation in their decision function (their skill level). This means there are specific over representative peaks in the experimental data, for example, at 15 and 25 in the duopoly case and 20 in the triopoly case, which the model with uniform priors can not capture. These peaks



**(c) Duopoly with a priori preference towards prominent numbers (10, 15, 20, 25, 30), reflecting a cognitive bias.**

**Figure 4: Cournot Oligopoly.** Experimental data from [23] is shown as grey bars. The proposed approach is shown in orange. The standard MARL approach is shown as the dashed purple line, and the NE is denoted by the black bar.

can not be explained from expected reward alone, as there is no particular reason that 15 would have such high preference. Instead, these demonstrate an *a priori* preference of the agents towards particular prominent numbers (0, 5, 10, ...), a known cognitive bias [10, 13]. Owing to the model’s flexibility in allowing for arbitrary prior beliefs, this can be modelled with  $q_i$  with higher weightings on these prominent numbers. An example of the resulting decisions when using such priors is shown in Fig. 4c, providing a significantly improved fit, capturing all of the experimental peaks. We do not use such a model when comparing results in Table 1, as this modification was made post-hoc (after seeing the experimental data), but it shows the usefulness of incorporating prior beliefs when known, demonstrating an additional strength of the model.

#### 4.2.3 Cobweb Market.

*Description.* In a cobweb market [35], there are  $K$  producers who must estimate the price  $\hat{p}_{i,t}$  of a good at the next timestep  $t$ . The reward for a producer  $i$  is based on the accuracy of their prediction compared to the market price  $p_t$ :

$$U_{i,t} = \max(0, 1300 - 260(p_t - \hat{p}_{i,t})^2) \quad (11)$$

which is lower bound by 0, e.g., the producers cannot receive negative utilities. Producers have no contact with others, but at the end of each round, producers observe the realised market price  $p_t$ .

The market price depends on the demand  $D$  and supply curves  $S$ .  $D$  is linear with price and is subject to small normally distributed

demand shocks  $\eta_t$ , and  $S$  non-linearly increases with the producer’s expected price, i.e.,


$$D(p_t) = a - bp_t + \eta_t, \quad S(\hat{p}_{i,t}) = \tanh(\psi(\hat{p}_{i,t} - K)) + 1 \quad (12)$$

where  $\psi$  controls the non-linearity and stability of the market. The realised market price is given by

$$p_t = \frac{a - \sum_{k \in K} S(\hat{p}_{k,t})}{b} + \epsilon_t \quad (13)$$

where  $\epsilon_t \propto \eta_t$ . Under rational expectations, producers all predict the price to be the intersection of  $S$  and  $D$ ,  $p^*$ , e.g.,  $\hat{p}_{i,t} = p^* + \epsilon_t, \forall i \in K$ , meaning the rational predictions will, on average, fall in line with the equilibrium price with fluctuations  $\propto \epsilon_t$ .

We utilise the experimental data of [35], with  $a = 13.8$ ,  $b = 1.5$ ,  $\epsilon_t \sim \mathcal{N}(0, 0.5)$  and  $\psi = 2$ . There were 36 participants, generally undergraduate economics, psychology, and science students.



**Figure 5: Cobweb Markets. Experimental data from [35] is shown as grey bars. The proposed approach is shown in orange. The standard MARL approach is shown as the dashed purple line, and the black line denotes the rational expectations solution.**

*Results.* The cobweb market results are visualised in Fig. 5, displaying the distribution of realised prices  $p_t$ . Both the standard MARL approach and the rational expectations are poor predictors of the observed phenomena from the experimental data. While the mean of the experimental data often aligns with the rational and MARL case, the distribution spread is far broader, indicating persistent excess volatility, with much larger standard deviations than those expected under the rational expectations hypothesis or MARL approach. This is a noteworthy stylised fact of markets incompatible with the rationality assumption of all agents [34, 35]. As excess volatility is known to occur in many markets [64], understanding the causes and being able to model this volatility is an important use of ABM. The proposed approach offers a much better fit, capturing the mean of the data and the overall distribution of price fluctuations, reproducing the observed excess price volatility (Fig. 5), providing an explanation of the endogenous formation of excess volatility based on bounded rationality.

### 4.3 Key Takeaways

The proposed approach demonstrated strong out-of-sample performance across these three economic and financial environments, outperforming the comparisons and validating the model in controlled environments. Specifically, we showcased the value of:

- **Boundedness:** Incorporating bounded rationality resulted in substantially improved predictive accuracy (Table 1).
- **Heterogeneity:** Allowing for heterogeneous processing costs improved upon assuming mutual consistency (Fig. 2b)
- **Non uniform priors:** Arbitrary prior beliefs explained phenomena incompatible with deviations from expected utility alone (Fig. 4c)

Additionally, while the proposed approach relaxes these three assumptions, if desired, these can still be recovered as special (limit) cases as discussed in Section 5. The benefit of the proposed approach is that these assumptions do not need to be established *a priori*, rather they are calibrated to the environment of interest.

## 5 DISCUSSION AND RELATION TO EQUILIBRIUM SOLUTIONS

Flexibility is one of the model’s strengths. However, this flexibility comes at the expense of exact analytical tractability, and generally, we are limited by the theoretical guarantees of the underlying RL algorithm (here, PPO). Despite this, in this section we show the relation to the decision functions of other equilibrium solution concepts and provide discussions on the equilibrium approximations.

*Quantal Response Equilibrium.* With homogeneous processing costs  $\lambda_i = \lambda$  and uniform prior beliefs  $q_i(a) = q$ , the approach can be seen as approximating QRE (as QRE converges to NE with  $\lambda \rightarrow 0$  [29], approximation of NE too). Using a similar formulation to Section 3.1.1, with QRE, each agent chooses  $a$  to maximise  $U$ , subject to an entropy  $H$  constraint:

$$\max \pi_i(a)U(a|\pi_{-i}) \quad \text{subject to} \quad H(\pi_i) \geq H_{\min} \quad (14)$$

where  $\pi_{-i}$  gives the action profile of the other agents. To derive the quantal response decision function  $QR_i$ , we use the method of Lagrange multipliers and the principle of maximum entropy to convert Eq. (14) into an unconstrained optimisation problem. Given the usual constraints on the probability function (that  $QR_i(a) \geq 0, \forall a$  and  $\sum_{a \in A} QR_i(a) = 1$ ), we get the following Lagrangian [18]:

$$\mathcal{L} = - \sum_{a \in A} QR_i(a)U(a|QR_{-i}) - \zeta \left( \sum_{a \in A} QR_i(a) - 1 \right) + \lambda \left( H(QR_i) - H_{\min} \right) \quad (15)$$

where taking the first order conditions and solving for  $QR_i$  yields

$$QR_i(a) = \frac{e^{U(a|QR_{-i})/\lambda}}{\sum_{a' \in A} e^{U(a'|QR_{-i})/\lambda}} \quad (16)$$

To demonstrate that the decision function implied by the  $D_{KL}$  constraint in Eq. (4) (with uniform priors and homogenous  $\lambda$ ) reduces to the same functional form as Eq. (16), we get:

$$I(\pi, s, q) = \sum_{a \in A} \pi(a|s) \log \frac{\pi(a|s)}{q(a|s)} = \sum_{a \in A} \pi(a|s) (\log(\pi(a|s)) - C) \quad (17)$$

plugging into  $\mathcal{L}$

$$\mathcal{L} = - \sum_{a \in A} \pi_i(a|s_i)U(a|\pi_{-i}) - \zeta \left( \sum_{a \in A} \pi_i(a|s_i) - 1 \right) + \lambda \left( \sum_{a \in A} \pi(a|s) (\log(\pi(a|s)) - C) - \bar{I} \right) \quad (18)$$

and the decision function reduces to:

$$\pi_i(a|s_i) = \frac{C e^{U(a|s_i)/\lambda}}{\sum_{a' \in A} C e^{U(a'|s_i)/\lambda}} = \frac{e^{U(a|s_i)/\lambda}}{\sum_{a' \in A} e^{U(a'|s_i)/\lambda}} \quad (19)$$

confirming equivalent functional forms to Eq. (16) under uniformity and homogeneity. The key difference is  $QR_i$  depends directly on the policies of other agents  $QR_{-i}$ , whereas  $\pi_i$  captures this indirectly via the state  $s_i$ .

The QRE corresponds to a fixed point of these QR functions [29], assuming that  $\lambda$  is homogeneous and common knowledge among the agents. In contrast, under the proposed approach, rather than explicitly attempting to find the fixed point solution, gradient descent and simulation are used to find  $\pi_i$  that maximises  $U$ , with a neural network  $f$  (with inputs  $s_i$  including  $q_i, \lambda_i$ ), and no common knowledge of  $\lambda_j, q_j, j \neq i$ . The outputs of  $f$  are  $|A|$  logits (one for each  $a \in A$ ), which are passed through a softmax function, giving learnt policies of the form:

$$\hat{\pi}_i(a) = \frac{e^{f_i(a|s_i)}}{\sum_{a' \in A} e^{f_i(a'|s_i)}} \quad (20)$$

where each agent is continually attempting to learn  $f_i$  that maximises their expected reward from using  $\hat{\pi}_i$  (here using PPO with Generalized Advantage Estimation [55]). Of course, except in very specific settings [66], we do not have general convergence guarantees, so we say the proposed approach *approximates* these equilibria.

The benefit of the proposed approach is the flexibility of  $f$  in allowing for various behaviours from heterogeneous agents (e.g., varying  $\lambda_i$  and  $q_i$ ) and computability when deriving the equilibria would otherwise be intractable, such as when  $\lambda_i$  and  $q_i$  are not common knowledge. When allowing heterogeneous  $\lambda_i$  and  $q_i$ , we approximate a Subjective Heterogeneous QRE [52], a type of Bayesian equilibrium [24], where agents may have different (potentially incorrect) subjective beliefs about the type distributions of the other agents (in this case, the values of  $\lambda_i$  and  $q_i$  in the population).

*Relation to Rational Inattention.* As mentioned in Section 2, the key relevant work in this area is [45]. While we share a similar goal, our work differs in some important ways. [45] requires estimating the mutual information (MI) for processing costs using a separate estimation engine. MI is defined over the joint probabilities as:

$$MI = - \sum_{a \in A} p(a, s_i) \log \frac{p(a, s_i)}{p(s_i)p(a)} \quad (21)$$

which has a dependence on the unconditional  $p(a)$  which must be solved with approximation techniques [18]. As the divergence we utilise in Eq. (4) does not have this same dependence, such estimation is not required, providing an alternative formulation allowing for arbitrary prior beliefs  $q_i$ , useful for representing cognitive biases (as demonstrated in Fig. 4c) or encoding behavioural heuristics. Furthermore, as discussed, we allow for a range of heterogeneous agent skills learnt through regularised policies, and propose an approach for efficiently calibrating these policies with agent supertypes and shared policy learning, both yet to be considered.

## 6 CONCLUSIONS

Agent-based models have much promise for explaining complex phenomena in a broad range of disciplines. However, a key criticism is how the behavioural rules are defined. Learning realistic

behavioural rules calibrated to real-world systems is essential to improve the models and promote continued uptake. In this work, we proposed an efficient MARL approach for inferring these decisions by calibrating heterogeneously skilled learning agents to real-world systems through shared policy learning and agent supertypes.

Under the proposed approach, agents possess diverse strategic processing abilities, represented through regularisation in their decision function. This regularisation is in the form of information processing costs, leading to varying levels of boundedly rational strategic behaviour, depending on the strength of regularisation. This agent skill heterogeneity is a critical aspect of many systems and is a departure from traditional equilibrium definitions. However, we demonstrate that this heterogeneity better captures many phenomena, as demonstrated under the various laboratory settings here and observed in many other real-world situations. For example, in market settings, institutional investors may have higher access to information and more extensive processing abilities than retail investors, altering the resulting market dynamics and potentially giving rise to behaviour deviating from the mutually consistent equilibrium. Relaxing this strict notion of equilibrium allows modelling a much broader range of dynamics.

The proposed approach does not impose strict assumptions on rationality, mutual consistency, or homogeneity but instead simulates the emergent outcomes through learning among the interacting agents. While these assumptions are not imposed, they can be recovered as special cases of the proposed approach, eliminating the requirement of determining which features are relevant *a priori*. We evaluated the proposed approach in various economic environments, demonstrating improved out-of-sample predictive accuracy compared to existing state-of-the-art MARL methods (PPO) and analytically derived equilibrium solutions. This work provides a valuable tool for modelling complex social systems and calibrating these models to real-world dynamics, particularly when analytical approaches become intractable, setting the foundation for more advanced simulations, e.g. limit order books [41].

## DISCLAIMER

This paper was prepared for informational purposes by the Artificial Intelligence Research group of JPMorgan Chase & Co and its affiliates (“J.P. Morgan”) and is not a product of the Research Department of J.P. Morgan. J.P. Morgan makes no representation and warranty whatsoever and disclaims all liability, for the completeness, accuracy or reliability of the information contained herein. This document is not intended as investment research or investment advice, or a recommendation, offer or solicitation for the purchase or sale of any security, financial instrument, financial product or service, or to be used in any way for evaluating the merits of participating in any transaction, and shall not constitute a solicitation under any jurisdiction or to any person, if such solicitation under such jurisdiction or to such person would be unlawful. © 2024 JPMorgan Chase & Co. All rights reserved.

## REFERENCES

- [1] Li An, Volker Grimm, Yu Bai, Abigail Sullivan, BL Turner II, Nicolas Malleon, Alison Heppenstall, Christian Vincenot, Derek Robinson, Xinyue Ye, et al. 2023. Modeling agent decision and behavior in the light of data science and artificial intelligence. *Environmental Modelling & Software* (2023), 105713.



- [2] Li An, Volker Grimm, Abigail Sullivan, BL Turner II, Nicolas Malleson, Alison Heppenstall, Christian Vincenot, Derek Robinson, Xinyue Ye, Jianguo Liu, et al. 2021. Challenges, tasks, and opportunities in modeling agent-based complex systems. *Ecological Modelling* 457 (2021), 109685.
- [3] Leo Ardon, Jared Vann, Deepika Garg, Thomas Spooner, and Sumitra Ganesh. 2023. Phantom-A RL-driven Multi-Agent Framework to Model Complex Systems. In *Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems*. 2742–2744.
- [4] Jasmina Arifovic and John Duffy. 2018. Heterogeneous agent modeling: experimental evidence. In *Handbook of Computational Economics*. Vol. 4. Elsevier, 491–540.
- [5] Robert L Axtell and J Doyne Farmer. 2022. Agent-based modeling in economics and finance: Past, present, and future. *Journal of Economic Literature* (2022).
- [6] Bowen Baker, Ingmar Kanitscheider, Todor Markov, Yi Wu, Glenn Powell, Bob McGrew, and Igor Mordatch. 2020. Emergent Tool Use From Multi-Agent Autocurricula. In *International Conference on Learning Representations*. <https://openreview.net/forum?id=SkxpxJBKwS>
- [7] Jean-Philippe Bouchaud. 2008. Economics needs a scientific revolution. *Nature* 455, 7217 (2008), 1181–1181.
- [8] Luigi Campanaro, Daniele De Martini, Siddhant Gangapurwala, Wolfgang Merkt, and Ioannis Havoutis. 2023. Roll-Drop: accounting for observation noise with a single parameter. In *Learning for Dynamics and Control Conference*. PMLR, 718–730.
- [9] Andrew Caplin, Mark Dean, and Daniel Martin. 2011. Search and satisficing. *American Economic Review* 101, 7 (2011), 2899–2922.
- [10] Tao Chen. 2018. Round-number biases and informed trading in global markets. *Journal of Business Research* 92 (2018), 105–117.
- [11] Yefen Chen, Xuanming Su, and Xiaobo Zhao. 2012. Modeling Bounded Rationality in Capacity Allocation Games with the Quantal Response Equilibrium. *Management Science* 58, 10 (2012), 1952–1962. <https://doi.org/10.1287/mnsc.1120.1531> arXiv:<https://doi.org/10.1287/mnsc.1120.1531>
- [12] Juin-Kuan Chong, Teck-Hua Ho, and Colin Camerer. 2016. A generalized cognitive hierarchy model of games. *Games and Economic Behavior* 99 (2016), 257–274.
- [13] Benjamin A Converse and Patrick J Dennis. 2018. The role of “Prominent Numbers” in open numerical judgment: Strained decision makers choose from a limited set of accessible numbers. *Organizational Behavior and Human Decision Processes* 147 (2018), 94–107.
- [14] Molood Ale Ebrahim Dehkordi, Jonas Lechner, Amineh Ghorbani, Igor Nikolic, Emile Chappin, and Paulien Herder. 2023. Using Machine Learning for Agent Specifications in Agent-Based Models and Simulations: A Critical Review and Guidelines. *Journal of Artificial Societies and Social Simulation* 26, 1 (2023).
- [15] Janez Demšar. 2006. Statistical comparisons of classifiers over multiple data sets. *The Journal of Machine Learning Research* 7 (2006), 1–30.
- [16] Greg d’Eon, Sophie Greenwood, Kevin Leyton-Brown, and James Wright. 2023. Loss Functions for Behavioral Game Theory. *arXiv preprint arXiv:2306.04778* (2023).
- [17] Benjamin Enke, Uri Gneezy, Brian Hall, David Martin, Vadim Nelidov, Theo Offerman, and Jeroen Van De Ven. 2023. Cognitive biases: Mistakes or missing stakes? *The Review of Economics and Statistics* 105, 4 (2023), 818–832.
- [18] Benjamin Patrick Evans and Mikhail Prokopenko. 2021. A maximum entropy model of bounded rational decision-making with prior beliefs and market feedback. *Entropy* 23, 6 (2021), 669.
- [19] Benjamin Patrick Evans and Mikhail Prokopenko. 2023. Bounded rationality for relaxing best response and mutual consistency: the quantal hierarchy model of decision making. *Theory and Decision* (17 May 2023). <https://doi.org/10.1007/s11238-023-09941-z>
- [20] Benjamin Eysenbach and Sergey Levine. 2019. If maxent rl is the answer, what is the question? *arXiv preprint arXiv:1910.01913* (2019).
- [21] Benjamin Eysenbach and Sergey Levine. 2022. Maximum Entropy RL (Provably) Solves Some Robust RL Problems. In *International Conference on Learning Representations*. <https://openreview.net/forum?id=PiSAD3caaA2>
- [22] William Fedus, Carles Gelada, Yoshua Bengio, Marc G Bellemare, and Hugo Larochelle. 2019. Hyperbolic discounting and learning over multiple horizons. *arXiv preprint arXiv:1902.06865* (2019).
- [23] L.E. Fouraker and S. Siegel. 1963. *Bargaining Behavior*. McGraw-Hill.
- [24] John Geanakoplos. 1994. Common knowledge. *Handbook of game theory with economic applications* 2 (1994), 1437–1496.
- [25] John Geanakoplos, Robert Axtell, Doyne J Farmer, Peter Howitt, Benjamin Conlee, Jonathan Goldstein, Matthew Hendrey, Nathan M Palmer, and Chun-Yi Yang. 2012. Getting at systemic risk via an agent-based model of the housing market. *American Economic Review* 102, 3 (2012), 53–58.
- [26] Gerd Gigerenzer. 2020. What is bounded rationality? In *Routledge handbook of bounded rationality*. Routledge, 55–69.
- [27] Gerd Gigerenzer and Wolfgang Gaissmaier. 2011. Heuristic decision making. *Annual review of psychology* 62 (2011), 451–482.
- [28] Jacob K Goeree and Charles A Holt. 2004. A model of noisy introspection. *Games and Economic Behavior* 46, 2 (2004), 365–382.
- [29] Jacob K Goeree, Charles A Holt, and Thomas R Palfrey. 2020. Stochastic game theory for social science: A primer on quantal response equilibrium. *Handbook of Experimental Game Theory* (2020), 8–47.
- [30] Russell Golman. 2011. Quantal response equilibria with heterogeneous agents. *Journal of Economic Theory* 146, 5 (2011), 2013–2028.
- [31] Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. 2018. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In *International conference on machine learning*. PMLR, 1861–1870.
- [32] Teck-Hua Ho, So-Eun Park, and Xuanming Su. 2021. A bayesian level-k model in n-person games. *Management Science* 67, 3 (2021), 1622–1638.
- [33] Nicolas Hoertel, Martin Blachier, Carlos Blanco, Mark Olfson, Marc Massetti, Marina Sánchez Rico, Frédéric Limosin, and Henri Leleu. 2020. A stochastic agent-based model of the SARS-CoV-2 epidemic in France. *Nature medicine* 26, 9 (2020), 1417–1421.
- [34] Cars Hommes and Thomas Lux. 2013. INDIVIDUAL EXPECTATIONS AND AGGREGATE BEHAVIOR IN LEARNING-TO-FORECAST EXPERIMENTS. *Macroeconomic Dynamics* 17, 2 (2013), 373–401. <https://doi.org/10.1017/S1365100511000162>
- [35] Cars Hommes, Joep Sonnemans, Jan Tuinstra, and Henk Van De Velden. 2007. Learning in cobweb experiments. *Macroeconomic Dynamics* 11, S1 (2007), 8–33.
- [36] Athul Paul Jacob, David J Wu, Gabriele Farina, Adam Lerer, Hengyuan Hu, Anton Bakhtin, Jacob Andreas, and Noam Brown. 2022. Modeling strong and human-like gameplay with KL-regularized search. In *International Conference on Machine Learning*. PMLR, 9695–9728.
- [37] Arnold Kling. 2018. Agent-based modeling: Promises and pitfalls. [https://www.econlib.org/archives/2011/02/agent-based\\_mod.html](https://www.econlib.org/archives/2011/02/agent-based_mod.html)
- [38] Maciej Latek, RL Axtell, and Bogumil Kaminski. 2009. Bounded rationality via recursion. In *Proceedings of Eighth International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2009)*. 457–464.
- [39] Roberto Leombruni and Matteo Richiardi. 2005. Why are economists sceptical about agent-based simulations? *Physica A: Statistical Mechanics and its Applications* 355, 1 (2005), 103–109.
- [40] Steven D Levitt and John A List. 2008. Homo economicus evolves. *Science* 319, 5865 (2008), 909–910.
- [41] Penghang Liu, Kshama Dwarakanath, and Svitlana S Vyetenko. 2022. Biased or limited: Modeling sub-rational human investors in financial markets. *arXiv preprint arXiv:2210.08569* (2022).
- [42] Richard D McKelvey and Thomas R Palfrey. 1995. Quantal response equilibria for normal form games. *Games and economic behavior* 10, 1 (1995), 6–38.
- [43] Richard D McKelvey and Thomas R Palfrey. 1998. Quantal response equilibria for extensive form games. *Experimental economics* 1 (1998), 9–41.
- [44] Richard D McKelvey, Thomas R Palfrey, and Roberto A Weber. 2000. The effects of payoff magnitude and heterogeneity on behavior in  $2 \times 2$  games with unique mixed strategy equilibria. *Journal of Economic Behavior & Organization* 42, 4 (2000), 523–548.
- [45] Tong Mu, Stephan Zheng, and Alexander R Trott. 2022. Modeling Bounded Rationality in Multi-Agent Simulations Using Rationally Inattentive Reinforcement Learning. *Transactions on Machine Learning Research* (2022). <https://openreview.net/forum?id=DY1pMrmDkm>
- [46] Pedro A Ortega and Daniel A Braun. 2013. Thermodynamics as a theory of decision-making with information-processing costs. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* 469, 2153 (2013), 20120683.
- [47] Osonde A Osoba, Raffaele Vardavas, Justin Grana, Rushil Zutshi, and Amber Jaycocks. 2020. Modeling agent behaviors for policy analysis via reinforcement learning. In *2020 19th IEEE International Conference on Machine Learning and Applications (ICMLA)*. IEEE, 213–219.
- [48] Benjamin Patrick Evans and Mikhail Prokopenko. 2023. Bounded strategic reasoning explains crisis emergence in multi-agent market games. *Royal Society Open Science* 10, 2 (2023), 221164.
- [49] Julien Perolat, Bart De Vylder, Daniel Hennes, Eugene Tarassov, Florian Strub, Vincent de Boer, Paul Muller, Jerome T Connor, Neil Burch, Thomas Anthony, et al. 2022. Mastering the game of Stratego with model-free multiagent reinforcement learning. *Science* 378, 6623 (2022), 990–996.
- [50] Ashreeta Prasanna, Sascha Holzhauser, and Friedrich Krebs. 2019. Overview of machine learning and data-driven methods in agent-based modeling of energy markets. *INFORMATIK 2019: 50 Jahre Gesellschaft für Informatik–Informatik für Gesellschaft* (2019).
- [51] Jeevnt Rampal and Fernando Stragliotto. 2023. Heterogeneous Agent Quantal Response Equilibrium. (2023).
- [52] Brian W Rogers, Thomas R Palfrey, and Colin F Camerer. 2009. Heterogeneous quantal response equilibrium and cognitive hierarchies. *Journal of Economic Theory* 144, 4 (2009), 1440–1467.
- [53] Ariel Rubinstein. 2003. “Economics and psychology”? The case of hyperbolic discounting. *International Economic Review* 44, 4 (2003), 1207–1216.
- [54] Ellis Scharfenaker. 2020. Implications of quantal response statistical equilibrium. *Journal of Economic Dynamics and Control* 119 (2020), 103990.

- [55] John Schulman, Philipp Moritz, Sergey Levine, Michael I. Jordan, and Pieter Abbeel. 2016. High-Dimensional Continuous Control Using Generalized Advantage Estimation. In *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2–4, 2016, Conference Track Proceedings*, Yoshua Bengio and Yann LeCun (Eds.). <http://arxiv.org/abs/1506.02438>
- [56] Herbert A. Simon. 1979. Rational decision making in business organizations. *The American economic review* 69, 4 (1979), 493–513.
- [57] Christopher A. Sims. 2003. Implications of rational inattention. *Journal of Monetary Economics* 50, 3 (2003), 665–690.
- [58] Samuel Sokota, Ryan D’Orazio, J Zico Kolter, Nicolas Loizou, Marc Lanctot, Ioannis Mitliagkas, Noam Brown, and Christian Kroer. 2023. A unified approach to reinforcement learning, quantal response equilibria, and two-player zero-sum games. *ICLR* (2023).
- [59] Callum Rhys Tilbury. 2023. *Reinforcement Learning for Economic Policy: A New Frontier?* Technical Report.
- [60] Yakup Turgut and Cafer Erhan Bozdog. 2023. A framework proposal for machine learning-driven agent-based models through a case study analysis. *Simulation Modelling Practice and Theory* 123 (2023), 102707.
- [61] Arthur Turrell. 2016. Agent-based models: understanding the economy from the bottom up. *Bank of England Quarterly Bulletin* (2016), Q4.
- [62] Nelson Vadori, Leo Ardon, Sumitra Ganesh, Thomas Spooner, Selim Amrouni, Jared Vann, Mengda Xu, Zeyu Zheng, Tucker Balch, and Manuela Veloso. [n.d.]. Towards multi-agent reinforcement learning-driven over-the-counter market simulations. *Mathematical Finance* n/a, n/a ([n. d.]). <https://doi.org/10.1111/mafi.12416> arXiv:<https://onlinelibrary.wiley.com/doi/pdf/10.1111/mafi.12416>
- [63] Nelson Vadori, Sumitra Ganesh, Prashant Reddy, and Manuela Veloso. 2020. Calibration of shared equilibria in general sum partially observable Markov games. *Advances in Neural Information Processing Systems* 33 (2020), 14118–14128.
- [64] Svitlana Vyetenko, David Byrd, Nick Petosa, Mahmoud Mahfouz, Danial Derovic, Manuela Veloso, and Tucker Balch. 2020. Get real: Realism metrics for robust limit order book market simulations. In *Proceedings of the First ACM International Conference on AI in Finance*. 1–8.
- [65] Ying Wen, Yaodong Yang, and Jun Wang. 2021. Modelling Bounded Rationality in Multi-Agent Interactions by Generalized Recursive Reasoning. In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence (Yokohama, Yokohama, Japan) (IJCAI’20)*. Article 58, 8 pages.
- [66] Kaiqing Zhang, Zhuoran Yang, and Tamer Basar. 2019. Policy optimization provably converges to Nash equilibria in zero-sum linear quadratic games. *Advances in Neural Information Processing Systems* 32 (2019).
- [67] Wei Zhang, Andrea Valencia, and Ni-Bin Chang. 2021. Synergistic integration between machine learning and agent-based modeling: A multidisciplinary review. *IEEE Transactions on Neural Networks and Learning Systems* (2021).
- [68] Stephan Zheng, Alexander Trott, Sunil Srinivasa, David C Parkes, and Richard Socher. 2022. The AI Economist: Taxation policy design via two-level deep multiagent reinforcement learning. *Science advances* 8, 18 (2022), eabk2607.