

# Heterogeneous RBCs via Deep Multi-Agent Reinforcement Learning

Federico Gabriele\*  
Università La Sapienza di Roma  
Italy

Aldo Glielmo†  
Banca d'Italia‡  
Italy

Marco Taboga†  
Banca d'Italia‡  
Italy

## ABSTRACT

Current macroeconomic models with agent heterogeneity can be broadly divided into two main groups. Heterogeneous-agent general equilibrium (GE) models, such as those based on Heterogeneous Agent New Keynesian (HANK) or Krusell–Smith (KS) approaches, rely on GE and ‘rational expectations’, somewhat unrealistic assumptions that make the models very computationally cumbersome, which in turn limits the amount of heterogeneity that can be modelled. In contrast, agent-based models (ABMs) can flexibly encompass a large number of arbitrarily heterogeneous agents, but typically require the specification of explicit behavioural rules, which can lead to a lengthy trial-and-error model-development process. To address these limitations, we introduce MARL-BC, a framework that integrates deep multi-agent reinforcement learning (MARL) with real business cycle (RBC) models. We demonstrate that MARL-BC can: (1) recover textbook RBC results when using a single agent; (2) recover the results of the mean-field KS model using a large number of identical agents; and (3) effectively simulate rich heterogeneity among agents, a hard task for traditional GE approaches. Our framework can be thought of as an ABM if used with a variety of heterogeneous interacting agents, and can reproduce GE results in limit cases. As such, it is a step towards a synthesis of these often opposed modelling paradigms.

## KEYWORDS

multi-agent systems; reinforcement learning; real business cycle

### ACM Reference Format:

Federico Gabriele, Aldo Glielmo, and Marco Taboga. 2026. Heterogeneous RBCs via Deep Multi-Agent Reinforcement Learning. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026)*, Paphos, Cyprus, May 25 – 29, 2026, IFAAMAS, 10 pages. <https://doi.org/10.65109/VZHC1838>

## 1 INTRODUCTION

Macroeconomic modelling has traditionally relied on General Equilibrium (GE) models employing representative agents, such as the

real business cycle (RBC) and new Keynesian (NK) models. However, a well-known limitation of representative-agent models is their inability to account for agent heterogeneity [40, 52].

To capture heterogeneity, two main modelling approaches emerged in macroeconomics: (i) heterogeneous-agent GE models, and (ii) agent-based models (ABMs). The first approach embeds heterogeneity directly into a GE framework, typically extending a representative agent model through a ‘mean-field’ approach. In this setting, individuals act given aggregate quantities, which act as a ‘mean field’, in analogy with physics. Influential early examples aimed at introducing heterogeneous households into the RBC setting are those of Aiyagari [2] and of Krusell and Smith [41]. More recently, Heterogeneous Agent New Keynesian (HANK) models have incorporated similar extensions into NK models [10, 39]. These GE heterogeneous-agent models come at a high cost, as each agent is assumed to have ‘rational expectations’ and needs to keep track of the entire wealth or income distribution as state variable to make any decision, an assumption that can be thought of as both unrealistic and too computationally burdensome [46]. Furthermore, the high computational costs involved significantly restrict the degree of heterogeneity achievable with GE models. For example, often heterogeneity appears ‘ex-post’, meaning that all agents are initially identical and differentiate only as a result of the individual shocks they undergo during the simulation. The second major approach is agent-based modelling. ABMs abandon the representative-agent and rational expectations assumptions entirely. They simulate economies ‘from the bottom up’ with many interacting agents that are both ex-ante and ex-post heterogeneous, have no rational expectations, and instead are endowed with boundedly rational heuristic decision rules [11, 14, 22, 27, 28]. This flexibility of ABMs is also the reason why they face the strongest criticism. Because in ABMs the modeller typically needs to decide the agents’ behaviour directly, it can be difficult to correctly navigate the arbitrariness in their specification and pin down realistic rules [13, 59].

Only very recently, researchers have begun exploring reinforcement learning (RL) [53] – and especially multi-agent reinforcement learning (MARL) [3] – as a novel way to model heterogeneous agents in macroeconomics [8, 18]. RL agents learn optimal (or nearly optimal) behaviours through repeated interaction with an environment to maximise their reward signals over time. The RL learning paradigm seems to offer a natural synthesis between the extremes of GE models and ABMs: agents can be boundedly rational and diverse, yet their behaviours emerge endogenously from a principled optimisation process (learning to maximise rewards), so that the modeller can avoid searching directly for a behavioural rule [16]. Moreover, modern deep RL techniques can solve complex, high-dimensional and nonlinear optimisation problems, which would be difficult to tackle with traditional methods.

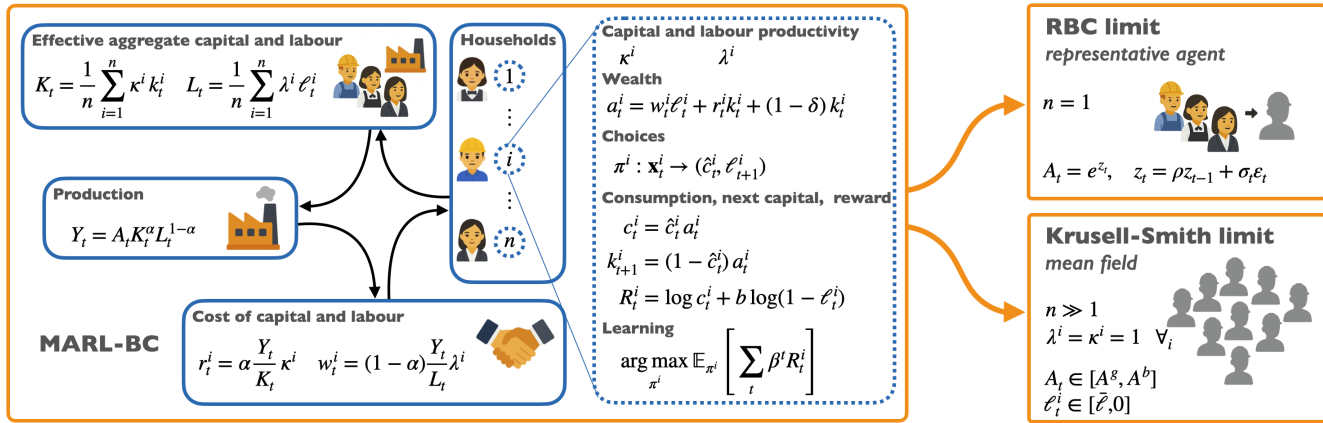
\*Work done during an internship at Banca d'Italia‡.

†aldo.glielmo@bancaditalia.it, marco.taboga@bancaditalia.it.

‡The views and opinions expressed in this paper are those of the authors and do not necessarily reflect the official policy or position of Banca d'Italia.



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



**Figure 1: The MARL-BC framework and two limit cases.** The left orange square contains a schematic illustration of the MARL-BC framework. A population of  $n$  heterogeneous RL households,  $i = 1, \dots, n$ , possesses distinct productivities of capital  $\kappa_i$  and of labour  $\lambda_i$ , and provides capital  $k_t^i$  and labour  $l_t^i$ . These individual inputs aggregate into total capital  $K_t$  and total labour  $L_t$ , which drive production  $Y_t$  via a Cobb-Douglas function. The individual interest rates  $r_t^i$  and the wages  $w_t^i$  are assumed to be proportional to the corresponding individual productivities. Wealth  $a_t^i$  is given by past depreciated capital  $(1 - \delta)k_t^i$ , earned wages  $w_t^i l_t^i$ , and earned interest  $r_t^i k_t^i$ . Each agent decides the fraction  $c_t^i$  of wealth to consume and the amount of labour to provide. These choices determine the instantaneous reward  $R_t^i$ , which the agent learns to maximise in the long run. The two right orange boxes represent classic macroeconomic models as limiting cases of our framework. The RBC model (top) is obtained using a representative agent ( $n = 1$ ) and an AR(1) process for the technology shocks  $A_t$ . The Krusell–Smith mean-field model (bottom) is obtained using many ex-ante identical agents and technology and labour switching between discrete values.

In this work, we follow this direction and show how deep RL techniques can bridge the gap between ABMs and heterogeneous-agent GE models. Our contributions are as follows:

- We develop the Multi-Agent Reinforcement Learning Business Cycle (MARL-BC), a MARL-based framework that extends the classic RBC model with multiple households characterised by rich and flexible heterogeneity.
- We demonstrate that the training of MARL-BC is computationally feasible, employing state-of-the-art RL algorithms such as Proximal Policy Optimisation (PPO) or Soft Actor Critic (SAC) or Deep Deterministic Policy Gradient (DDPG).
- We show that MARL-BC can recover classic textbook RBC results when using a single agent.
- We further show that MARL-BC can recover the mean-field Krusell–Smith model when using a large population of ex-ante identical agents.
- Finally, we illustrate the capability of the framework to simulate richer heterogeneity among agents, impossible to achieve with a GE approach.

Our results demonstrate the value and relevance of the MARL-BC framework for a growing community of researchers in both economics and computer science, thus paving the way for numerous applications and future investigations.

**Related work.** The idea of using RL as a model for realistic human decision making is attracting increasing research and development interest for economic and financial simulations [8, 56]. Influential early contributions that sparked interest in this direction were the studies in [38] and [64], where deep multi-agent RL was employed to simulate emergent economic behaviours within simplified toy

economies. Initially, its adoption occurred mostly within the financial sector, where RL was successfully applied to model trading strategies in different markets [5, 12, 25, 33, 43], and has led to the development of specialised open-source platforms [4, 6]. In macroeconomics, recent literature has begun exploring RL techniques to extend classic GE frameworks [18]. For instance, [17, 37] use RL for monetary policy, while [19, 20, 24, 45] employ MARL for tax policy optimisation. In [9] and [61], RL is used to solve representative-agent and heterogeneous GE models respectively, while in [36] it is leveraged to solve a hybrid model that combines features of both macro ABMs and GE models. Parallel studies focused on extending RL with ABMs. For instance, in [16], the authors incorporate RL firms within a standard ABM, while in [34] the authors use RL to enhance the calibration of ABMs. In [1], RL is used to design robust macroeconomic policies resilient to model misspecification and uncertainty. Similarly, [26] proposes a general framework combining adaptive agent behaviours through RL within ABMs. Our work is also related to another line of ongoing research which aims to tackle the curse of dimensionality of GE models with deep learning, but without using RL [29, 30]. Finally, our work can contribute to research studying the effects of bounded rationality, typically done without RL, in existing ABMs or GE models [7, 23].

**Significance.** Our work stands out clearly in comparison with the described literature for the following reasons. Existing literature at the *economics* end of the spectrum has mostly focused on single-agent RL, showing that it can solve representative-agent GE models. However, research in economics has *not* explored how to build more expressive models using a multi-agent RL approach. Conversely, the *computer-science* literature has experimented with multi-agent RL, showing that it can give rise to rich emergent

economic behaviours. However, research in computer science has mostly overlooked fundamental models from macroeconomics, and it has *not* convincingly shown how multi-agent RL can extend such models. In our view, this has created a clear knowledge gap between the two disciplines, which currently impedes communication and mutually beneficial exchanges of ideas. For example, this implied that, on the one hand, economic modelling could not leverage the full potential of multi-agent RL as a viable strategy for extending its traditional modelling frameworks and, on the other hand, that research and development in computer science could not progress rapidly towards the construction of frameworks of actual relevance for the economics discipline.

Our work represents a clear step forward in bridging this knowledge gap, as it provides a foundation for linking research in two disciplines. For this reason, we believe it has significant potential for stimulating fruitful mutual exchanges. Specifically, we bridge the two research lines by proposing a framework that (i) exactly reproduces the representative-agent RBC solution when  $n = 1$ , (ii) converges to the KS mean-field equilibrium as  $n \gg 1$ , and (iii) extends beyond both limits to easily accommodate a high degree of discrete heterogeneity that GE methods cannot easily handle.

The rest of this work is structured as follows. In Sec. 2 we describe the MARL-BC framework. In Sec. 3, we provide details on the training and testing procedure we followed for the numerical experiments. In Sec. 4, we describe the results of our numerical experiments, dividing the discussion into three parts: the replication of textbook RBC results, the replication of mean-field Krusell–Smith results, and the illustration of the full capabilities of our framework. Finally, in Sec. 5, we conclude.

## 2 THE MARL-BC FRAMEWORK

We describe here the Multi-Agent Reinforcement Learning Business Cycle (MARL-BC) framework, in its general form. This will allow us to analyse the traditional models, such as the textbook RBCs or the Krusell–Smith, simply as limiting cases. Our framework and the mentioned limiting cases are illustrated in Figure 1. Specifically, the orange box in the figure contains the general MARL-BC framework. The framework is composed of two pieces: the heterogeneous RBC environment, depicted as the four solid blue boxes, is described in Sec. 2.1, while the RL household problem, depicted as the dashed blue box, is described in Sec. 2.2.

### 2.1 The heterogeneous RBC environment

The model consists of  $n$  types of households  $i = 1, \dots, n$  and a single firm. At each time  $t$ , each type of household (which we will also refer to as a ‘household’) can own a different amount of capital  $k_t^i$  and can provide a different amount of labour  $\ell_t^i$ . Furthermore, each agent has a different productivity of capital  $\kappa^i$  and of labour  $\lambda^i$ . We assume here that capital and labour productivities are fixed in time, but the model can be easily extended to allow them to vary.

**Effective aggregate capital and labour.** Households’ capital and labour are assumed to give rise to the effective aggregate capital  $K_t$  and labour  $L_t$  as

$$K_t = \frac{1}{n} \sum_{i=1}^n \kappa^i k_t^i, \quad L_t = \frac{1}{n} \sum_{i=1}^n \lambda^i \ell_t^i, \quad (1)$$

which are averages of the households’ variables, weighted by the corresponding productivities. Although these are formally written averages, we take the population mass to be one (a convention that also ensures consistency with standard GE models), so averages coincide with aggregates, quantities are expressed on a per capita basis, and the firm’s conditions in Eq. (3) apply.

**Production.** Given aggregate capital and labour, production  $Y_t$  is given by a Cobb–Douglas function

$$Y_t = A_t K_t^\alpha L_t^{1-\alpha}. \quad (2)$$

with output elasticity of capital  $\alpha$ , and technology  $A_t$ .

**Cost of capital and labour.** We assume that labour and capital markets are perfectly competitive. Therefore, wages  $w_t^i$  and interest rates  $r_t^i$  are proportional to the marginal productivities of capital and labour respectively,

$$r_t^i = \alpha \frac{Y_t}{K_t} \kappa_i, \quad w_t^i = (1 - \alpha) \frac{Y_t}{L_t} \lambda_i. \quad (3)$$

**Households’ wealth.** Each household receives wages and interest as compensation for its labour and capital contributions to aggregate output while it faces depreciation of its capital stock at a fixed rate  $\delta$ . Hence, the wealth  $a_t^i$  of household  $i$  at time  $t$  is

$$a_t^i = w_t^i \ell_t^i + r_t^i k_t^i + (1 - \delta) k_t^i, \quad (4)$$

which is the sum of labour earnings  $w_t^i \ell_t^i$ , returns on capital  $r_t^i k_t^i$ , and the depreciated stock of capital  $(1 - \delta) k_t^i$ .

### 2.2 The RL households

**Action space.** The RL households’ action at each time step  $t$  is a tuple  $(\hat{c}_t^i, \ell_{t+1}^i)$  of two continuous numbers, the consumption fraction  $\hat{c}_t^i$  and the labour  $\ell_{t+1}^i$ . While theoretically both of these fall in the range  $(0, 1)$ , for numerical stability we clip them in the range  $(0.01, 0.99)$ .

The consumption fraction  $\hat{c}_t^i$  determines  $c_t^i$ , which is the amount of the household’s wealth spent on consumption

$$c_t^i = \hat{c}_t^i a_t^i, \quad (5)$$

as well as the amount dedicated to investments in next-step capital


$$k_{t+1}^i = (1 - \hat{c}_t^i) a_t^i. \quad (6)$$

By choosing  $\hat{c}_t^i$  and  $\ell_{t+1}^i$ , the households need to balance current consumption and leisure with future earnings, as will be made clearer in the rest of this section.

**Observation space.** The observation space  $\mathbf{x}_t^i$  of the RL household  $i$  can flexibly contain any information on the state of the agent and on the state of the economy that is deemed important for the agent’s decisions. A rather general choice, coherent with standard economic models, would be

$$\mathbf{x}_t^i = (k_t^i, K_t, \ell_t^i, L_t, A_t, \kappa^i, \lambda^i), \quad (7)$$

containing the values of individual and aggregate capital, past individual and aggregate labour, and technology. In practice, not all of these variables are needed to obtain interesting dynamics or to reproduce traditional macroeconomic models, and, in fact, we will sometimes select a subset as specified in the results section.



**Figure 2: State-of-the-art RL schemes can solve the household optimisation problem.** Mean rewards obtained by different RL schemes as a function of the number of training steps. All schemes considered succeed in learning a policy that maximises cumulative rewards. SAC, DDPG and TD3 are seen to perform significantly better than PPO in terms of sample efficiency. However, PPO’s greater computational efficiency (not shown in this graph) can make also this method a good candidate for the task.

**Reward.** Each RL household learns a deterministic policy

$$\pi^i : \mathbf{x}_t^i \rightarrow (\hat{c}_t^i, \ell_{t+1}^i) \tag{8}$$

giving consumption fraction  $\hat{c}_t^i$  and labour  $\ell_{t+1}^i$  as a function of the state of the household and of the economy  $\mathbf{x}_t^i$ . The policy  $\pi^i$  is learned by maximising  $\mathcal{R}^i$ , the expected sum of discounted rewards

$$\mathcal{R}^i = \mathbb{E}_{\pi^i} \left[ \sum_t \beta^t R_t^i \right], \tag{9}$$

where  $\beta \in (0, 1)$  is the agent’s discount factor and the superscript in  $\beta^t$  here indicates an exponentiation to the power  $t$ , differently from anywhere else in this work, where it is just used as an index. The reward  $R_t^i$  of household  $i$  at time  $t$  is defined as

$$R_t^i = \log c_t^i + b \log(1 - \ell_t^i). \tag{10}$$

It increases with increasing consumption  $c_t^i$  and leisure  $(1 - \ell_t^i)$ , with  $b > 0$  controlling the trade-off between the two. Rewards are known as ‘utilities’ in the economic literature. Furthermore, the specific function in Eq. (10) is very commonly used since the logarithmic utility is, in fact, a special case of the well-known ‘Constant Relative Risk Aversion’ (CRRA) utility function when the coefficient of relative risk aversion equals one [49].

**Game-theoretic considerations.** From a game-theoretic perspective, the households can be thought of as participating in a stochastic game with  $n$  players  $i = 1, \dots, n$ , a common state, individual actions  $(\hat{c}_t^i, \ell_{t+1}^i)$ , a transition kernel induced by Eqs. (1–4), and per-period payoffs  $R_t^i$ . This, in turn, ensures the existence of optimal policies for the households assuming a solution concept such as Markov perfect equilibrium [44]. However, the corresponding problem is well known to be computationally hard [21]. Consequently, although the approach would be theoretically possible, it is impractical due to the high complexity and dimensionality of the

underlying model and its strategy space. In addition, relying on the full information set of the model is an unrealistic assumption, as real-world agents typically lack full observability of their environment. In contrast, we train independent learners each having access only to the partial information set  $\mathbf{x}_t^i$  and optimising for a policy  $\pi^i$  that approximates a best response to the evolving behaviour of the others. In our numerical experiments, we show that for sufficiently large  $n$ , the resulting MARL scheme reproduces the equilibrium behaviour of [41]. However, from a theoretical standpoint, establishing convergence of independent MARL dynamics to mean field equilibria in stochastic games remains a challenging and open problem [62, 63].

### 3 EXPERIMENTAL SETUP


In our numerical experiments, we first showcase the capabilities of our framework to reproduce the traditional representative agent RBC model and the mean-field Krusell–Smith (KS) model (Sec. 4.1 and 4.2), and we then show the capabilities of the framework in the general form (Sec. 4.3). In the rest of the paper, we refer to these three experimental settings as ‘RBC’, ‘KS’ and ‘General’.

**Parameters.** Table 1 reports the main parameters used in our numerical experiments, along with their names and symbols. The values reported are grouped depending on the type of experiment performed, either ‘RBC’, ‘KS’, or ‘General’. If a single value is reported, then the same value has been used for all experiments. We note here that most of the parameters used are standard for macroeconomic models of these types [49].

**Training and testing.** We implemented MARL-BC using well-known open-source libraries in Python. Specifically, we use the MARL interface of PettingZoo [54], and use SuperSuit [55] to leverage RL algorithms from Stable-Baselines3 [48] in an independent MARL fashion. We consider the following RL algorithms: Deep Deterministic Policy Gradient (DDPG) [42], Twin Delayed Deep Deterministic Policy Gradient (TD3) [31], Soft Actor Critic (SAC) [35], Proximal Policy Optimization (PPO) [51]. We use default hyperparameters as detailed in [32, Appendix].

Symbol	Description	Value (RBC   KS   General)
$n$	Number of households	1   20   20
$T$	Episode length	500
$\kappa_i$	Capital productivity	1   1   {0, 0.8, 1, 1.2, 0.98, 1.02}
$\lambda_i$	Labour productivity	1   1   {0.98, 1, 1.02}
$A_t$	Technology	AR(1)   AR(1)   {KS, AR(1)}
$\alpha$	Output elasticity	0.36
$\delta$	Capital depreciation	{1, 0.025}   0.025   0.025
$\beta$	Discount factor	0.95
$b$	Leisure weight	5   0   {0, 5}

**Table 1: Parameters used in the experiments.** When three values are reported, separated by a vertical bar, they refer respectively to the three result categories (RBC, KS, and General) as specified in the main text. For the AR(1) technology process we use  $\rho = 0.9$  and  $\sigma = 0.01$ . The KS technology process is described in the main text.



**Figure 3: MARL-BC can reproduce representative-agent results with a single agent.** The **left** panel shows how the RL household recovers the optimal consumption fraction  $\hat{c}$  and labour supply  $\ell$  of Eq. (13), for a textbook, analytically solvable, RBC model with full capital depreciation ( $\delta = 1$ ). The **centre** panel shows a similar convergence, but for a more complex model with a partial depreciation ( $\delta = 0.025$ ). This model cannot be analytically solved, and the optimal choices are obtained here using numerical tools common in economics. Finally, the **right** panel shows that MARL-BC is also capable of reproducing the dynamical behaviour of the standard RBC models as captured by an impulse response function. The plot shows the impulse response for consumption following a productivity shock of one standard deviation, with the dashed orange line indicating the standard RBC result and the full blue line indicating the equivalent MARL-BC behaviour.

To facilitate quick convergence and scalability, we adopt the standard MARL paradigm of parameter sharing [3]. In practice, this means a single neural network represents all the agents. Importantly, this paradigm does *not* force an identical behaviour for all agents, since the different agents’ characteristics (such as the individual productivities of labour and capital) appear as observations for the neural network, which can thus encode completely different behaviours. Parameter sharing significantly reduces the computational burden, improves sample efficiency and scalability, and can help to stabilise training. Furthermore, it can also be thought of as a stylised form of ‘social learning’ [47, 58] since information discovered by an agent is immediately made available to similar agents via shared parameters.

We train all algorithms for  $10^6$  steps in the single-agent environment. For the multi-agent runs, we perform  $10^5$  per-agent updates, corresponding to  $n \cdot 10^5$  total steps, sometimes more if needed for convergence (as detailed in [32, Appendix]). Unless otherwise stated, we report mean and standard deviation over five different training runs. The code for the MARL-BC environment used to perform our experiments is available in open source at <https://github.com/fedegabriele/MARL-BC>.

## 4 RESULTS

### 4.1 Representative agent RBC limit

The classic RBC model entails a single representative household, assumed to internalise all aggregate variables. As illustrated in the top right panel of Figure 1, our framework can recover exactly the RBC model as a special case by just choosing a single agent and unitary productivity of labour and capital, i.e.

$$n = 1, \quad \kappa = \lambda = 1, \tag{11}$$

so that individual and aggregate quantities coincide:  $K_t = k_t$  and  $L_t = \ell_t$ . Furthermore, as for all standard RBCs, we choose the technology to move according to the exponential of an AR(1) process

$$A_t = e^{z_t}, \quad z_t = \rho z_{t-1} + \sigma \varepsilon_t, \tag{12}$$

where  $\rho$  and  $\sigma$  are the persistence and volatility of the stochastic process, and  $\varepsilon_t$  is a random number sampled independently from a standard normal distribution.


**Algorithmic performance.** Figure 2 shows the four RL schemes considered (PPO, SAC, TD3 and DDPG) trained on the textbook RBC problem. All RL algorithms considered for this problem found successful policies to optimise cumulative rewards. More specifically, SAC, TD3 and DDPG appear to strongly outperform PPO in terms of speed of convergence towards the optimal solution, with SAC being the most stable learner among all. As we will show later (see Figure 6), SAC is the strongest candidate for the multi-agent experiments. DDPG is another valid option for the single-agent experiments since, although it takes slightly more steps to converge, it is also faster than SAC in terms of wall-clock time. For these reasons, we choose to leverage the speed of DDPG in the RBC limit experiments shown here and the efficacy and stability of learning of SAC for the multi-agent experiments.

**Reproducing textbook RBCs ( $\delta = 1$ ).** Here, we reproduce the results of the simplest type of RBC model as covered in most macroeconomic textbooks, one with full depreciation ( $\delta = 1$ ). In this setting, capital cannot be transferred to the next step and closed-form solutions can be obtained [15]. The optimal policy  $(\hat{c}_t^*, \ell_t^*)$  reads

$$\hat{c}_t^* = (1 - \alpha\beta), \quad \ell_t^* = \frac{\alpha}{b(1 - (1 - \alpha)\beta) + \alpha}. \tag{13}$$

The left panel of Figure 3 shows how the RL household learns to recover the optimal actions of the textbook RBC. Specifically, after around  $10^4$  training steps, the consumption fraction  $\hat{c}_t$  and labour  $\ell_t$  (shown in blue and orange) approximately converge to the optimal values shown as dashed lines, and then remain stably around those values.

**Reproducing typical RBCs ( $\delta = 0.025$ ).** More general RBCs set up a partial depreciation, often of  $\delta = 0.025$  as done in our experiments. Relaxing the full-depreciation assumption makes the model



**Figure 4: MARL-BC can reproduce mean-field Krusell–Smith results with a large number of identical agents.** The left panels are scatter plots of the aggregate capital at two consecutive times  $K_t$  and  $K_{t+1}$ , at four training snapshots, with progressively more training steps (namely  $2^7 \approx 10^2$ ,  $2^{11} \approx 2 \cdot 10^3$ ,  $2^{19} \approx 5 \cdot 10^4$  and  $2^{21} \approx 2 \cdot 10^6$ ). The characteristic linear time evolution of aggregate capital (cf. the first figure of the original paper [41]) is here recovered with a sufficient number of training steps. The centre panels depict the wealth distributions of the agents before and after training. While the Gini index for the untrained model is 0.08, it moves to 0.18 for the trained model distribution, a number which is in line with the original work. Finally, the right panels illustrate the curves of the marginal propensities to consume learned by the MARL-BC agents. Analogously to the KS agent policy, the consumption policy learned by the MARL-BC agents is almost completely flat when wealth is high, while it increases sharply when wealth is low.

impossible to solve analytically. In these cases, economic modellers typically resort to approximate solutions relying on first-order conditions and linearisation as implemented in widespread software such as ‘Dynare’ [57]. The centre panel of Figure 3 shows that, after approximately  $10^4 - 10^5$  training steps, the RL household of our framework learns optimal consumption and labour choices that are coherent with those computed with the Dynare software.

**Reproducing impulse response functions.** We further demonstrate that our framework can reproduce standard impulse response functions. To compute them, we initialise the model at its deterministic steady state (computed with  $\sigma_z = 0$ ), hit the economy with a one-standard-deviation productivity shock, and observe the relaxation to the steady state value. The right panel of Figure 3 overlays the impulse–response functions for consumption from our framework (blue lines with the shaded area) and Dynare (black line); the two lines are statistically consistent, confirming that our framework correctly reproduces not only the stationary choices but also the dynamic response to external shocks.

### 4.2 Mean-field Krusell–Smith limit

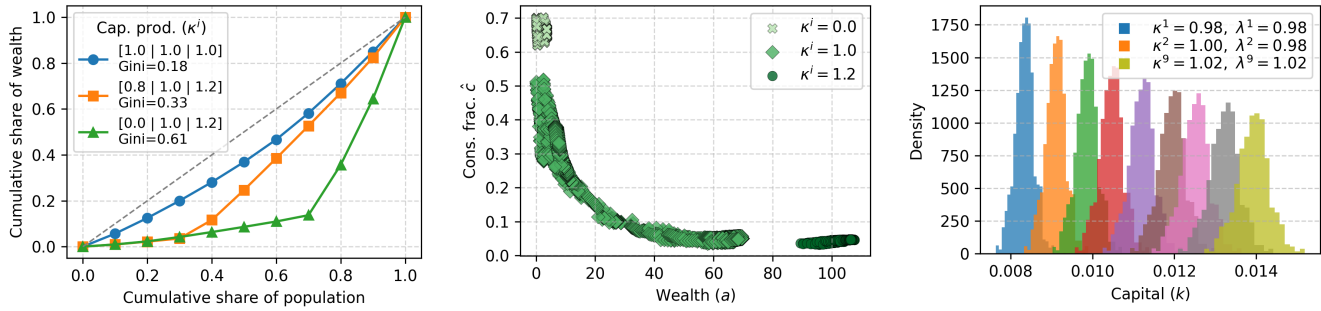
The celebrated ‘Krusell–Smith’ (KS) model [41] is a mean-field model with ex-ante identical households that face discrete shocks of individual unemployment and aggregate technology. As illustrated in the bottom right panel of Figure 1, our framework can recover the mean-field KS limit with a large number of identical agents

$$n \gg 1, \quad \kappa^i = \lambda^i = 1 \quad \forall i. \quad (14)$$

In practice, in our experiments we use  $n = 20$  agents as we expect this to be sufficient to reproduce trends and distributions of the mean field limit. Furthermore, to recover the precise KS

model behaviour, we set up the aggregate and individual shocks in close agreement with the original model. Specifically, we let the aggregate technology factor  $A_t$  follow a two-state Markov chain  $A_t \in [A_t^g, A_t^b]$ . In the ‘good’ ( $g$ ) or ‘bad’ ( $b$ ) state, technology is respectively higher  $A_t^g = 1.02$  or lower  $A_t^b = 0.98$ . Furthermore, individual households also follow a two-state Markov chain, being either employed and supplying  $\ell_t^i = \bar{\ell} = 1.11$  units of labour, or being unemployed and supplying no labour ( $\ell_t^i = 0$ ). The transition probabilities between aggregate and individual states are taken from [41]. They are calibrated to US data to make the expected duration of each regime eight periods, and in such a way that the unemployment rate is around 4% in good states and 10% in bad states. Since labour is fixed, we let the RL household choose only the consumption fraction  $\hat{c}_t^i$  and set  $b = 0$  in the reward function. Finally, in agreement with [41], we set the observation space to  $\mathbf{x}_t^i = (k_t^i, \ell_t^i, K_t, A_t)$ . The results of these experiments are given in Figure 4 through a single illustrative training run, and discussed below. Statistics over different seeds are available in [32, Appendix].

**Reproducing the classic KS law of motion.** The left panel of Figure 4 illustrates the convergence properties of our MARL-BC framework relative to the classic ‘law of motion’ for aggregate capital postulated in the original paper [41]. The four small scatter plots display the pairs  $(K_t, K_{t+1})$  after approximately  $10^2$ ,  $2 \cdot 10^3$ ,  $5 \cdot 10^4$ , and  $2 \cdot 10^6$  SAC updates. In the KS procedure, a linear map is postulated whose coefficients are estimated until they converge in a self-consistent procedure. By contrast, using MARL-BC a perfectly linear relationship ( $R^2 > 0.99$ ) between consecutive aggregate capitals emerges endogenously without any a priori assumption.



**Figure 5: MARL-BC allows for modelling rich ex-ante heterogeneity and for the emergence of heterogeneous behaviour.** The left panel reports the Lorenz curves, and the corresponding Gini indices, of the wealth distributions computed using MARL-BC, either in the standard KS setting (blue dots) or in the extended KS with heterogeneous capital productivities (orange circles and green triangles). The graph illustrates how MARL-BC allows one to model a wider range of wealth inequalities. The centre panel reports the marginal propensities to consume of the three classes of agents (of low, medium and high returns) for the KS model with Gini index equal to 0.61. The graph illustrates how the three household groups learn markedly different policies. The right panel illustrates the kind of wealth distribution that one can obtain from an RBC model with 9 agents with different combinations of labour and capital productivity. The legend shows only 3 colours, but the pattern is clear: higher productivities imply higher wages or capital returns that, in turn, give rise to higher capital accumulation.

**Reproducing KS distributional properties.** The two histograms in the centre panel of Figure 4 report the stationary distribution of individual capital before and after convergence. Before convergence, the scale of the capital distribution is very small, ranging only from 0 to 1.2. This is a sign that the agents have not yet learned consumption policies capable of accumulating capital. Not surprisingly, the Gini index is very low. After convergence, the Gini index increases to 0.18, a value relatively close to the 0.25 value from the original KS calculation.

**Reproducing KS marginal propensity to consume.** The two graphs in the right panel of Figure 4 report the agents’ marginal propensity to consume as a function of their wealth, before and after training (top and bottom graphs). The graphs show that the learning process converges to curves that are flat for large values of wealth ( $a > 20$ ) and increase rapidly for low values of wealth (especially after  $a < 5$ ). Interestingly, there is very little difference in the consumption fractions if the agents are employed (orange dots) or unemployed (blue dots). These shapes for the curves of the marginal propensity to consume are another key result of the original KS paper that the MARL-BC can easily recover.

### 4.3 Greater heterogeneity with MARL-BC

We illustrate here how MARL-BC can seamlessly extend classic models by leveraging its flexibility to represent a greater level of agent heterogeneity. Specifically, we extend the KS model with heterogeneous capital productivities and then extend the RBC model with heterogeneity in both capital and labour productivities.

**KS with heterogeneous capital returns.** We extend the KS model by dividing the  $n = 20$  agents into three groups, respectively with low, middle and high capital productivity. Specifically, we assign 14 agents (or 70% of the total) to the middle productivity group, and 6 agents (or 30% of the total) equally between the low and high productivity groups. Although this assignment is largely arbitrary, the described proportions were selected to resemble those of

a normal distribution, where roughly 70% of the data falls within one standard deviation of the mean and 30% lies in the two tails. We perform two experiments, with either mild or marked differences in capital productivity among the three groups. To be more precise, in the ‘mild’ experiment we let the three productivities  $\kappa^i$  be  $\{0.8, 1.0, 1.2\}$ , while in the ‘marked’ experiment we let them be  $\{0.0, 1.0, 1.2\}$ . As capital productivities have an immediate effect on capital returns via Eq. (3), this KS extension can neatly represent the heterogeneous returns on capital of real economies [60].

The left panel of Figure 5 reports the Lorenz curves, and the corresponding Gini indices, of the original KS experiment (blue dots) as well as of the mildly heterogeneous experiment (orange squares) and markedly heterogeneous experiment (green triangles). The graph nicely illustrates how the introduction of heterogeneous capital productivity, and hence of heterogeneous capital returns, allows for the modelling of more unequal economies, with Gini indices increasing to 0.33 and 0.61 in the mildly and markedly heterogeneous returns models. Different capital returns also give rise to different marginal propensities to consume as reported in the centre panel of Figure 5. Specifically, the graph shows a scatter plot of the consumption fraction as a function of wealth for the markedly heterogeneous experiment.

The graph shows that the households with lower capital returns (green crosses) learn a policy known as ‘hand-to-mouth’, in which they consume almost all of their wealth each step without any capital accumulation. This type of policy is often encoded in GE models as a heuristic, while it emerges endogenously with our framework as a result of learning. On the other extreme, households with higher capital returns learn to consume a very low portion of their wealth since this policy leads to wealth accumulation and thus to higher future consumption.

**RBC with heterogeneous returns and wages.** To further illustrate the flexibility and scope of MARL-BC, we build a heterogeneous version of the standard RBC described in Sec. 4.1 with nine


agents, each with capital and labour productivity taken from a  $3 \times 3$  grid containing values in the set  $\{0.98, 1, 1.02\}$ . Thus, the agent with the lowest productivity has  $\kappa^i = \lambda^i = 0.98$ , the one with the highest productivity has  $\kappa^i = \lambda^i = 1.02$ , and the other agents have all combinations in the middle. In turn, these different productivities give rise to heterogeneity in both capital returns and wages, which influence consumption and labour decisions as well as wealth levels. In the right panel of Figure 5, we illustrate the results of this setup by showing how the nine agents stabilise with overlapping yet distinct levels of wealth after a sufficient number of training steps. Although we do not show this in the figure, the different productivities also give rise to significantly different choices for consumption fraction and labour supply.

**Scalability to hundreds of agents.** To test the scalability of MARL-BC to simulations with a large number of heterogeneous agents, we expand on the experiment discussed in the previous paragraph and perform training runs for progressively denser, equispaced, grids of capital and labour productivities ( $\kappa^i$  and  $\lambda^i$ ), from  $3 \times 3 = 9$  agents to  $23 \times 23 = 529$  agents. The results of these runs are shown in Figure 6. The top panel in the figure reports the mean and two standard deviations of the best episodic return attained during training runs of  $10^5$  per-agent updates. The bottom panel reports the corresponding wall-clock time taken for training on a single CPU. We find that the performance of SAC remains stably high across all model sizes. On the contrary, we find that PPO and DDPG strongly underperform respectively in regimes of low and high numbers of agents. Finally, we find that the computational cost rises steadily with the number of agents, and yet remains within practical limits for SAC (it takes roughly two hours to train the largest model for roughly  $50 \cdot 10^6$  steps). These results indicate that MARL-BC can be successfully trained with hundreds of households, even using very modest hardware. Further details on the scalability of the framework are available in [32, Appendix].

## 5 CONCLUSIONS

In this work, we present the Multi-Agent Reinforcement Learning Business Cycle (MARL-BC) framework, a macroeconomic modelling framework that integrates deep MARL within a real business cycle (RBC) environment. We show that MARL-BC can reproduce the canonical representative-agent RBCs and the Krusell–Smith mean-field models, and it can extend them by simulating multiple heterogeneous agents.

One obvious limitation of our framework is the computational cost, with accurate training of RL agents requiring up to a few hours to terminate successfully. However, the use of vectorised GPU computation and of more modern MARL schemes is likely to greatly mitigate this shortcoming in the near future. Specifically, by implementing the MARL environment described in a GPU-compatible form, for instance using specialised software [50], would allow one to fully leverage GPU acceleration (which can easily exceed a factor of 20). Furthermore, while the independent learning scheme used here has the advantage of being simple to understand and implement, more sophisticated MARL approaches could improve learning stability and speed [3]. This would allow much faster and more robust training, hence greatly facilitating usability, prototyping and extension of MARL-BC.



**Figure 6: MARL-BC can scale up to hundreds of agents.** The **top** panel reports the best rewards achieved by MARL-BC agents during training runs of  $10^5$  per-agent updates, meaning  $n \cdot 10^5$  steps in total. The **bottom** panel reports the time needed for the corresponding runs, on a single CPU machine. Both quantities are shown as a function of the number of agents ( $n$ ) in the model. PPO is seen to underperform with respect to its competitors. On the contrary, SAC is observed to obtain comparatively high rewards, in reasonable time limits, even for the largest models considered.

While this study focused primarily on the description of MARL-BC, its limit cases, and its capabilities, future work could investigate the use of the framework to study specific economic problems. Natural issues could be related to economic inequality, but also to the economic and financial consequences of asymmetric changes in labour productivity, such as those potentially deriving from the spreading of AI tools in the workplace. By showing how a model of multiple interacting agents (akin to an ABM) can give rise to heterogeneous GE results, our work marks a step towards a synthesis of these often opposed methodologies. Furthermore, by showing how a MARL approach can recover and extend classic results from the economics literature, our work has significant potential to foster communication and exchanges between communities in computer science and economics that are interested in RL modelling, but that have not previously found common ground.

## ACKNOWLEDGMENTS

We would like to thank Alessandro Moro, Valerio Astuti, Valerio Nispi Landi, Sergio Santoro (Banca d'Italia), Doyne Farmer and José Morán (Oxford University), Marta Grześkiewicz (Cambridge University) and Sara Casella (LUISS) for useful feedback, and Michele Colombi (Scuola Normale Superiore) for double-checking some of the results.

## REFERENCES

- [1] Akash Agrawal, Joel Dyer, Aldo Glielmo, and Michael J Wooldridge. 2025. Robust policy design in agent-based simulators using adversarial reinforcement learning. In *The First MARW: Multi-Agent AI in the Real World Workshop at AAI 2025*.
- [2] S. Rao Aiyagari. 1994. Uninsured Idiosyncratic Risk and Aggregate Saving. *Quarterly Journal of Economics* 109, 3 (1994), 659–684.
- [3] Stefano V Albrecht, Filippos Christianos, and Lukas Schäfer. 2024. *Multi-agent reinforcement learning: Foundations and modern approaches*. MIT Press.
- [4] Selim Amrouni, Aymeric Moulin, Jared Vann, Svitlana Vyetenko, Tucker Balch, and Manuela Veloso. 2021. ABIDES-gym: gym environments for multi-agent discrete event simulation and application to financial markets. In *Proceedings of the Second ACM International Conference on AI in Finance*. 1–9.
- [5] Leo Ardon, Nelson Vadori, Thomas Spooner, Mengda Xu, Jared Vann, and Sumittra Ganesh. 2021. Towards a fully rl-based market simulator. In *Proceedings of the Second ACM International Conference on AI in Finance*. 1–9.
- [6] Leo Ardon, Jared Vann, Deepika Garg, Thomas Spooner, and Sumittra 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* (London, United Kingdom) (AAMAS '23). International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 2742–2744.
- [7] Yuki M Asano, Jakob J Kolb, Jobst Heitzig, and J Doyne Farmer. 2021. Emergent inequality and business cycles in a simple behavioral macroeconomic model. *Proceedings of the National Academy of Sciences* 118, 27 (2021), e2025721118.
- [8] Tohid Atashbar and Rui Aruhan Shi. 2022. Deep Reinforcement Learning: Emerging Trends in Macroeconomics and Future Prospects. *IMF Working Papers* 2022, 259 (12 2022), 1. <https://doi.org/10.5089/9798400224713.001>
- [9] Tohid Atashbar and Rui Aruhan Shi. 2023. AI and Macroeconomic Modeling: Deep Reinforcement Learning in an RBC model. *IMF Working Papers* 2023, 040 (2 2023), 1. <https://doi.org/10.5089/9798400235252.001>
- [10] Adrien Auclert, Matthew Rognlie, and Ludwig Straub. 2025. Fiscal and monetary policy with heterogeneous agents. *Annual Review of Economics* 17 (2025).
- [11] Robert L Axtell and J Doyne Farmer. 2025. Agent-based modeling in economics and finance: Past, present, and future. *Journal of Economic Literature* 63, 1 (2025), 197–287.
- [12] Tiziano Bacaloni, Aldo Glielmo, and Marco Taboga. 2025. Natural-gas storage modelling by deep reinforcement learning. In *Proceedings of the 6th ACM International Conference on AI in Finance (ICAIIF '25)*. Association for Computing Machinery, New York, NY, USA, 829–837. <https://doi.org/10.1145/3768292.3770348>
- [13] Marco Benedetti, Gennaro Catapano, Francesco De Scavias, Marco Favorito, Aldo Glielmo, Davide Magnanini, and Antonio Muci. 2022. Black-it: A ready-to-use and easy-to-extend calibration kit for agent-based models. *Journal of Open Source Software* 7, 79 (2022), 4622.
- [14] András Borsos, Adrian Carro, Aldo Glielmo, Marc Hinterschweiger, Jagoda Kaszowska-Mojas, and Arzu Uluc. 2025. Agent-Based Modelling at Central Banks: Recent Developments and New Challenges. *Bank of Italy Occasional Papers (Questioni di Economia e Finanza)* 956 (2025).
- [15] William A Brock and Leonard J Mirman. 1972. Optimal economic growth and uncertainty: the discounted case. *Journal of Economic Theory* 4, 3 (1972), 479–513.
- [16] Simone Brusatin, Tommaso Padoan, Andrea Coletta, Domenico Delli Gatti, and Aldo Glielmo. 2024. Simulating the economic impact of rationality through reinforcement learning and agent-based modelling. In *Proceedings of the 5th ACM International Conference on AI in Finance*. 159–167.
- [17] Mingli Chen, Rama Cont, Andreas Joseph, Michael Kumhof, Xinlei Pan, Xiong Wei, and Xuan Zhou. 2025. Deep reinforcement learning in a monetary model. *Bank of England Staff Working Paper* 1142 (2025).
- [18] Thomas R Cook and Nathan M Palmer. 2025. Reinforcement Learning in Macroeconomics. In *Oxford Research Encyclopedia of Economics and Finance*.
- [19] Michael Curry, Alexander Trott, Soham Phade, Yu Bai, and Stephan Zheng. 2022. Analyzing micro-founded general equilibrium models with many agents using deep reinforcement learning. *arXiv preprint arXiv:2201.01163* (2022).
- [20] Michael Curry, Alexander Trott, Soham Phade, Yu Bai, and Stephan Zheng. 2023. Learning Solutions in Large Economic Networks using Deep Multi-Agent Reinforcement Learning. In *Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems* (London, United Kingdom) (AAMAS '23). International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 2760–2762.
- [21] Constantinos Daskalakis, Paul W Goldberg, and Christos H Papadimitriou. 2009. The complexity of computing a Nash equilibrium. *Commun. ACM* 52, 2 (2009), 89–97.
- [22] Herbert Dawid and Domenico Delli Gatti. 2018. Agent-based macroeconomics. *Handbook of computational economics* 4 (2018), 63–156.
- [23] Giovanni Dosi, Mauro Napoletano, Andrea Roventini, Joseph E Stiglitz, and Tania Treibich. 2020. Rational heuristics? Expectations and behaviors in evolving economies with heterogeneous interacting agents. *Economic Inquiry* 58, 3 (2020), 1487–1516.
- [24] Kshama Dwarakanath, Jialin Dong, and Svitlana Vyetenko. 2024. Tax Credits and Household Behavior: The Roles of Myopic Decision-Making and Liquidity in a Simulated Economy. In *Proceedings of the 5th ACM International Conference on AI in Finance*. 168–176.
- [25] Benjamin Patrick Evans and Sumittra Ganesh. 2024. Learning and Calibrating Heterogeneous Bounded Rational Market Behaviour with Multi-agent Reinforcement Learning. In *Proceedings of the 23rd International Conference on Autonomous Agents and Multiagent Systems* (Auckland, New Zealand) (AAMAS '24). International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 534–543.
- [26] Benjamin Patrick Evans, Sihan Zeng, Sumittra Ganesh, and Leo Ardon. 2025. ADAGE: A Generic Two-layer Framework for Adaptive Agent based Modelling. In *Proceedings of the 24th International Conference on Autonomous Agents and Multiagent Systems* (Detroit, MI, USA) (AAMAS '25). International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 2510–2513.
- [27] J Doyne Farmer. 2024. *Making Sense of Chaos: A Better Economics for a Better World*. Yale University Press.
- [28] J. Doyne Farmer and Duncan Foley. 2009. The Economy Needs Agent-Based Modelling. *Nature* 460, 7256 (2009), 685–686.
- [29] Jesús Fernández-Villaverde. 2025. *Deep Learning for Solving Economic Models*. Technical Report. National Bureau of Economic Research.
- [30] Jesús Fernández-Villaverde, Galo Nuño, and Jesse Perla. 2024. *Taming the curse of dimensionality: quantitative economics with deep learning*. Technical Report. National Bureau of Economic Research.
- [31] Scott Fujimoto, Herke Hoof, and David Meger. 2018. Addressing function approximation error in actor-critic methods. In *International conference on machine learning*. PMLR, 1587–1596.
- [32] Federico Gabriele, Aldo Glielmo, and Marco Taboga. 2025. Heterogeneous RBCs via deep multi-agent reinforcement learning. *arXiv preprint arXiv:2510.12272* (2025). <https://arxiv.org/abs/2510.12272>
- [33] Kang Gao, Stephen Weston, Perukrishnen Vytelingum, Namid Stillman, Wayne Luk, and Ce Guo. 2023. Deeper hedging: A new agent-based model for effective deep hedging. In *Proceedings of the Fourth ACM International Conference on AI in Finance*. 270–278.
- [34] Aldo Glielmo, Marco Favorito, Debmallya Chanda, and Domenico Delli Gatti. 2023. Reinforcement Learning for Combining Search Methods in the Calibration of Economic ABMs. In *Proceedings of the Fourth ACM International Conference on AI in Finance*. 305–313.
- [35] 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.
- [36] Edward Hill, Marco Bardoscia, and Arthur Turrell. 2021. Solving heterogeneous general equilibrium economic models with deep reinforcement learning. *arXiv preprint arXiv:2103.16977* (2021).
- [37] Natascha Hinterlang and Alina Tänzer. 2021. Optimal monetary policy using reinforcement learning. (2021).
- [38] Michael Bradley Johanson, Edward Hughes, Finbarr Timbers, and Joel Z Leibo. 2022. Emergent bartering behaviour in multi-agent reinforcement learning. *arXiv preprint arXiv:2205.06760* (2022).
- [39] Greg Kaplan, Benjamin Moll, and Giovanni L. Violante. 2018. Monetary Policy According to HANK. *American Economic Review* 108, 3 (2018), 697–743.
- [40] Alan P. Kirman. 1992. Whom or What Does the Representative Individual Represent? *Journal of Economic Perspectives* 6, 2 (1992), 117–136.
- [41] Per Krusell and Anthony A. Smith. 1998. Income and Wealth Heterogeneity in the Macroeconomy. *Journal of Political Economy* 106, 5 (1998), 867–896.
- [42] Timothy P. Lillicrap, Jonathan J. Hunt, Alexander Pritzel, Nicolas Heess, Tom Erez, Yuval Tassa, David Silver, and Daan Wierstra. 2016. Continuous control with deep reinforcement learning. In *ICLR*, Yoshua Bengio and Yann LeCun (Eds.).
- [43] Chris Mascioli, Anri Gu, Yongzhao Wang, Mithun Chakraborty, and Michael Wellman. 2024. A Financial Market Simulation Environment for Trading Agents Using Deep Reinforcement Learning. In *Proceedings of the 5th ACM International Conference on AI in Finance*. 117–125.
- [44] Eric Maskin and Jean Tirole. 2001. Markov perfect equilibrium: I. Observable actions. *Journal of Economic Theory* 100, 2 (2001), 191–219.
- [45] Qirui Mi, Siyu Xia, Yan Song, Haifeng Zhang, Shenghao Zhu, and Jun Wang. 2024. TaxAI: A Dynamic Economic Simulator and Benchmark for Multi-agent Reinforcement Learning. In *Proceedings of the 23rd International Conference on Autonomous Agents and Multiagent Systems* (Auckland, New Zealand) (AAMAS '24). International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 1390–1399.
- [46] Benjamin Moll. 2025. The trouble with rational expectations in heterogeneous agent models: A challenge for macroeconomics. *arXiv preprint arXiv:2508.20571* (2025).
- [47] Kamal K Ndousse, Douglas Eck, Sergey Levine, and Natasha Jaques. 2021. Emergent social learning via multi-agent reinforcement learning. In *International conference on machine learning*. PMLR, 7991–8004.
- [48] Antonin Raffin, Ashley Hill, Adam Gleave, Anssi Kanervisto, Maximilian Ernestus, and Noah Dormann. 2021. Stable-baselines3: Reliable reinforcement learning implementations. *Journal of machine learning research* 22, 268 (2021), 1–8.

- [49] David H. Romer. 2018. *Advanced Macroeconomics* (5 ed.). McGraw-Hill Education, New York, NY. 800 pages.
- [50] Alexander Rutherford, Benjamin Ellis, Matteo Gallici, Jonathan Cook, Andrei Lupu, Garoar Ingvarsson, Timon Willi, Ravi Hammond, Akbir Khan, Christian Schroeder de Witt, Alexandra Souly, Saptarashmi Bandyopadhyay, Mikayel Samvelyan, Minqi Jiang, Robert Tjarko Lange, Shimon Whiteson, Bruno Lacerda, Nick Hawes, Tim Rocktaschel, Chris Lu, and Jakob Nicolaus Foerster. 2024. JaxMARL: Multi-Agent RL Environments and Algorithms in JAX. In *The Thirty-eight Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.
- [51] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347* (2017).
- [52] Thomas M Stoker. 1993. Empirical approaches to the problem of aggregation over individuals. *Journal of Economic literature* 31, 4 (1993), 1827–1874.
- [53] Richard S Sutton, Andrew G Barto, et al. 1998. *Reinforcement learning: An introduction*. Vol. 1. MIT press Cambridge.
- [54] J Terry, Benjamin Black, Nathaniel Grammel, Mario Jayakumar, Ananth Hari, Ryan Sullivan, Luis S Santos, Clemens Dieffendahl, Caroline Horsch, Rodrigo Perez-Vicente, et al. 2021. Pettingzoo: Gym for multi-agent reinforcement learning. *Advances in Neural Information Processing Systems* 34 (2021), 15032–15043.
- [55] Justin K Terry, Benjamin Black, and Ananth Hari. 2020. Supersuit: Simple microwrappers for reinforcement learning environments. *arXiv preprint arXiv:2008.08932* (2020).
- [56] Callum Rhys Tilbury. 2022. Reinforcement Learning for Economic Policy: A New Frontier? *arXiv preprint arXiv:2206.08781* (2022).
- [57] Sébastien Villemot et al. 2011. Solving rational expectations models at first order: what Dynare does. *url: https://citeseerx.ist.psu.edu/document* (2011).
- [58] Andrew Whalen, Thomas L Griffiths, and Daphna Buchsbaum. 2018. Sensitivity to shared information in social learning. *Cognitive science* 42, 1 (2018), 168–187.
- [59] Paul Windrum, Giorgio Fagiolo, and Alessio Moneta. 2007. Empirical validation of agent-based models: Alternatives and prospects. *Journal of Artificial Societies and Social Simulation* 10, 2 (2007), 8.
- [60] Inês Xavier. 2021. Wealth inequality in the US: the role of heterogeneous returns. *Available at SSRN 3915439* (2021).
- [61] Yucheng Yang, Chiyuan Wang, Andreas Schaab, and Benjamin Moll. 2025. Structural Reinforcement Learning for Heterogeneous Agent Macroeconomics. *arXiv preprint arXiv:2512.18892* (2025).
- [62] Batuhan Yardim, Semih Cayci, Matthieu Geist, and Niao He. 2023. Policy mirror ascent for efficient and independent learning in mean field games. In *International Conference on Machine Learning*. PMLR, 39722–39754.
- [63] Bora Yongacoglu, Gürdal Arslan, and Serdar Yüksel. 2024. Mean-field games with finitely many players: independent learning and subjectivity. *Journal of Machine Learning Research* 25, 419 (2024), 1–69.
- [64] 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.