Reinforcement Learning and Consumption-Savings Behavior

Brandon Kaplowitz*

[Click here for latest version]

February 3, 2026

Abstract

This paper demonstrates how reinforcement learning can explain two puzzling empirical patterns in household consumption behavior during economic downturns. I develop a model where agents use Q-learning with neural network approximation to make consumption-savings decisions under income uncertainty, departing from standard rational expectations assumptions. The model replicates two key findings from recent literature: (1) unemployed households with previously low liquid assets exhibit substantially higher marginal propensities to consume (MPCs) out of stimulus transfers compared to high-asset households (0.50 vs 0.34), even when neither group faces borrowing constraints, consistent with Ganong et al. (2024); and (2) households with more past unemployment experiences maintain persistently lower consumption levels after controlling for current economic conditions, a "scarring" effect documented by Malmendier and Shen (2024). Unlike existing explanations based on belief updating about income risk or ex-ante heterogeneity, the reinforcement learning mechanism generates both higher MPCs and lower consumption levels simultaneously through value function approximation errors that evolve with experience. Simulation results closely match the empirical estimates, suggesting that adaptive learning through reinforcement learning provides a unifying framework for understanding how past experiences shape current consumption behavior beyond what current economic conditions would predict.

*New York University. Email: bgk258@stern.nyu.edu.

I am grateful to my committee, Jaroslav Borovicka, Thomas Sargent, Thomas Philippon, Andrew Caplin, Peter Ganong, Corina Boar, Jess Benhabib, Simon Gilchrist, and Gianluca Violante for helpful comments and suggestions. For valuable discussions and ideas, I thank Kevin Guo, Spencer Kwon, Sobhan Mohammadpour, Matthew Fellows, Samuel Sokota, Christian Schroeder De Witt, Jakob Foerster, Gabriele Farina, Chiara Gardenghi, Man Chon Iao, and Yatheesan Selvakumar. I also thank seminar participants at the NYU Macro Student Lunch, Simon Institute Theory of RL, AI4ABM Workshop, and NYU Stern Macro Workshop. All errors are my own.

1 Introduction

1.1 Motivation and Goals

The overarching goal of this paper is to understand the implications of sophisticated learning protocols for consumption behavior. Specifically, in this project, I focus on two empirical patterns documented in recent years and show that embedding reinforcement learning, a widely used protocol in computer science, in an otherwise standard model of consumption choice can generate these patterns.

In my model, a consumer is represented as an economically rational decision-maker with informational and computational constraints.[1] In particular, the model assumes that the consumer does not know their income process perfectly, nor are they able to solve a Bellman equation to obtain their true value function. Instead, they learn their continuation value over time using reinforcement learning. They form a parametric estimator of their continuation value and update the parameters of this estimator over time based on their experiences.

One crucial assumption that brings considerable tractability is that an agent uses a Markovian formulation of their decision problem, where continuation values depend only on current assets and income.[2] This is restrictive because it presupposes knowledge of the relevant states, but serves as a natural starting point; with this assumption, the problem becomes defined in terms of a Markov Decision Problem (MDP), where reinforcement learning is an appropriate solution procedure.[3] Two recent empirical findings motivate this analysis. First, Ganong et al. (2024) study household spending during the COVID-19 recession and document a puzzling pattern: unemployed households with previously low liquid assets exhibited much higher marginal propensities to consume (MPCs) out of stimulus transfers (0.53) than those with previously high assets (0.29), even when the transfers were large enough that neither group faced current borrowing constraints.

Second, Malmendier and Shen (2024) examine how past economic experiences affect current behavior using panel data from the Panel Study of Income Dynamics (PSID) and

1. I refer to this consumer as an *agent* throughout.

2. Other important assumptions made that are standard in an economics context, but may want to be questioned in this context, include: economic rationality of households (that is, that they are subjective expected utility maximizers); agents forming forward-looking expectations; no special memory, context, replay buffer, attention, or model of the environment; choice of learning algorithm among reinforcement learning algorithms; preferences admitting a unique utility representation and this utility representation being available to households; time-separability of preferences; no cost for computation time or other forms of complexity or access to multiple alternative solution schemes; and no preferences regarding the timing of the resolution of uncertainty, such as ambiguity aversion.

3. Relaxing the agent "magically" knowing their true state variables represents a natural follow-up. See Section A.4 for further details.

NielsenIQ Homescan Consumer Panel. They find that households with more past unemployment experiences—either personal or through exposure to high regional unemployment—maintain persistently higher savings rates even after controlling for all current economic conditions including income, wealth, and employment status, and suggest it has a highly persistent effect that they capture via a weighted index. This "scarring" effect suggests that past unemployment experiences create lasting behavioral changes beyond their immediate economic impact.

These findings present challenges for standard full-information rational expectation models of consumption, even augmented with borrowing constraints.

The model assumes agents approximate their continuation value with a flexible parametric model with a neural network estimate at the rational solution and then proceed to update the parameters governing this model, the weights and biases of the neural network, over the course of their lifetime via gradient descent. To obtain their optimal consumption choice, agents fit a polynomial to the output of the neural network. For initialization, the model assumes each agent begins their life with an estimate of the continuation value that coincides with its rational expectations counterpart. This means that deviations from rational expectations in the model arise from experiences and learning dynamics rather than initial conditions. A key parameter is the learning rate, which controls the rate at which agents update the parameters in response to new observations. The baseline parameterization sets the learning rate to a high initial value and decays it at the rate $1/\sqrt{t}$ following the work of Robbins and Monro (1951) to satisfy the criteria of stochastic approximation for stochastic gradient descent and, hence, enable the possibility of convergence in the long run. I simulate the model with standard quarterly consumption-savings parameters heavily drawn from Ganong et al. (2024) and Krueger et al. (2016). I examine their consumption profiles over 50 quarters.

I use the calibrated model to replicate two results that connect to recently documented empirical facts. First, consistent with Ganong et al. (2024), agents who previously held low liquid assets exhibit substantially higher marginal propensities to consume (MPCs) out of stimulus transfers compared to previously high-asset agents, even when transfers of cash are large enough that neither group is currently borrowing constrained. Using data generated via simulation, I find estimates very similar to those found by Ganong et al. (2024). Second, following Malmendier and Shen (2024), agents with more past unemployment experiences maintain higher savings rates, significant even after controlling for current assets and income and dropping the two most recent observations—a *scarring* effect from past unemployment experiences.

Thus, this paper presents RL as an alternative unifying explanation for these facts. It

does not rely on ex-ante heterogeneity—the hypothesis in Ganong et al. (2024)—which cannot account for the experience effects observed in Malmendier and Shen (2024). RL is also distinct from the scarring mechanism in Malmendier and Shen (2024), as scarring focuses solely on changes in future expected unemployment probabilities. Empirically, my mechanism generates both lower consumption levels and higher MPCs, consistent with the data. In contrast, increased subjective unemployment risk (as learned through scarring) causes MPCs and consumption levels to fall together. This pattern is inconsistent with the higher MPCs that Ganong et al. (2024) find among agents with lower endogenous prior assets—consistent with agents who have experienced past unemployment and depleted their savings through consumption-smoothing.

2 Literature

This paper is related to several strands of the literature. The substantive question bears a direct connection to a large, long-standing body of work studying the effects of incomplete information and learning on consumption behavior.

Learning and Consumption One influential strand explores the implications of unknown income transition probabilities, where an agent learns over time about their income process but perfectly solves their dynamic optimization problem.[4] This approach includes Cogley and Sargent (2008), who introduced *anticipated utility* models to macroeconomics, where agents use a Bayesian approach to learn about an underlying exogenous stochastic process for income or some other state they have to forecast, but then solve their value function to a fixed point, assuming no further learning happens from tomorrow onward. While specifics of the learning process and economic settings differ, Friedman (1957), Muth (1960), Jovanovic (1979), Marcet and Sargent (1989), Evans et al. (1994), Guvenen (2007), and Kozlowski et al. (2020) all fall into this category. Malmendier and Shen (2024), who document the scarring fact mentioned earlier, use a similar model to explain their fact.

This paper differs from that work in that it does not assume that agents learn about their income process or solve a fixed-point problem every period. Instead, they directly estimate the value function associated with different choices each period. In particular, this paper never requires them to solve their value function forward in time via iteration, given a particular set of beliefs. However, agents remain forward-looking in that they always are forming estimates of their expected value function and making decisions based on estimated future expected utility realizations.

4. By iterating forward on the perceived income transition probabilities to a fixed point.

Reinforcement Learning The second large strand of literature is the reinforcement learning literature, which has been widely applied in computer science and to some extent in economics. This literature focuses on estimating the value function of a Markov decision process (MDP) and using this to make decisions about how to act in the future. Sutton and Barto (2018) provide a comprehensive overview and introduction to this literature. Additionally, there has been a large body of recent literature that suggests that humans learn in ways that resemble (deep) reinforcement learning. Daw et al. (2011), Dabney et al. (2020), Botvinick et al. (2020), Muller et al. (2024), Kasdin et al. (2025), Masset et al. (2025), and Sousa et al. (2025)

In economics, the most relevant papers focusing specifically on reinforcement learning as a learning mechanism are the works by Barberis and Jin (2023) and Ilut and Valchev (2024). Barberis and Jin (2023) study how some puzzles in finance may be explained by a combination of "model-free" learning (represented by a variant of Q-Learning) and "model-based" learning (represented by the correctly specified Q-function, given a learned perceived distribution over returns, akin to adaptive learning). Ilut and Valchev (2024) study a particular Gaussian process-based variant of Q-Learning as a theory for decision-making and explore how agents may trade off between collecting samples and using a linear model to update their value functions over time. They interpret their model via a unified Bayesian lens that determines how optimally agents should decide when to use the model-free versus model-based approach. They also look at how additional factors, such as rational inattention, influence how agents trade off between the model-based and model-free approach.

This paper is motivated by two empirical facts from the recent literature, which I will briefly summarize here. The first fact is from Ganong et al. (2024), who study how households' spending patterns during the COVID-19 recession varied based on past financial wealth, and the second fact is from Malmendier and Shen (2024), who study how past unemployment experiences affect current consumption behavior.

Marginal Propensities to Consume During the COVID-19 recession, between March 29, 2020 and July 31, 2020, Congress temporarily added $600 Federal Pandemic Unemployment Compensation (FPUC) top-ups to regular UI with the CARES Act.

The paper leverages a unique natural experiment created by administrative processing delays in state unemployment insurance systems. This quasi-random variation caused otherwise-similar unemployed workers to receive benefits at different times in different states. Ganong et al. use this to identify causal effects of benefit receipt on household spending patterns and to explore how these effects differ based on household liquid asset levels, particularly looking at differences in marginal propensities to consume (MPCs) across households with

different histories. The paper examines differences in consumption due to staggered stimulus receipt across matched otherwise similar households in different states and defines MPCs as the additional share consumed by households that received stimulus payments within a month of stimulus receipt relative to those that didn't, as a share of the total stimulus payment received. This is meant to estimate the slope of the consumption function—what fraction of a dollar of additional income is immediately consumed by households. This provides the first fact to target.

Fact #1: During COVID-19, the average MPC out of stimulus transfers for previously low-asset unemployed households is substantially higher (0.53) than that of previously high-asset unemployed households (0.29), in Ganong et al.'s preferred specification. This result holds even though, at present, neither type of agent is borrowing constrained (Ganong et al. 2024).

This is surprising because much previous work argued that high MPCs observed during unemployment resulted from agents being borrowing-constrained and lacking sufficient liquid assets to smooth consumption.

In fact, benefits were large enough that unemployed households temporarily moved up the asset distribution relative to before they were unemployed, with the median household moving from the 38th percentile to the 63rd percentile of the liquidity distribution, an adjusted measure of bank account balance sheets. Because MPCs fall as asset levels increase in a full-information rational expectations model with borrowing constraints, households with significant asset buffers should not exhibit large differences in MPCs even if there remain persistent differences in their asset levels, raising a puzzle of what causes this difference. Additionally, despite high liquidity, these households also exhibited high marginal propensities to consume relative to past estimates from the literature such as Kaplan and Violante (2022).[5]

Furthermore, if agents are rational and fully informed, previous asset conditions from several years ago should have little effect on today's decision-making since assets fully capture the impact of the history on states today.

5. Borrowing constraints alone result in lower MPCs in calibrated simulation, between 5% and 30% for quarterly MPCs out of an unexpected windfall of $500, with the lower end being estimates for average quarterly MPCs in models without illiquid assets, 25% as the MPC for the poor hand-to-mouth household, and 30% for the wealthy hand-to-mouth household in two-asset variants with a liquid and illiquid asset. According to Kaplan and Violante (2022), in turn citing Jappelli and Pistaferri (2010) and Havranek and Sokolova (2020), a large body of literature argues in favor of MPCs out of transitory income changes of $500–$1,000 of 15–25%. (Kaplan and Violante 2022).

Scarring Malmendier and Shen (2024) investigate whether personal experiences of economic downturns create persistent behavioral changes that extend beyond the immediate economic impact. Using panel data from the PSID (1993–2013) and NielsenIQ Homescan Consumer Panel (2004–2013), they construct an experience-based index that captures households' exposure to unemployment over their lifetime from both personal and macroeconomic sources, such as regional or national unemployment rates, with more recent experiences receiving greater weight through a weighting scheme that places less weight the further ago an unemployment experience occurred, using a weighted averaging scheme. The core question of the paper examines whether past unemployment experiences affect current consumption behavior even after controlling for all observable economic characteristics, including current income, wealth, employment status, and demographic factors. That is, among two otherwise identical households *today*, does one household's behavior change when they had unemployment experiences in their past or lived through periods of higher national unemployment? By regressing household consumption on this unemployment experience index while holding constant current economic conditions, they can identify whether past hardships leave lasting behavioral "scars" that manifest as persistently lower consumption and higher savings. This provides the second fact to target.

> **Fact #2:** Comparing households by their previous unemployment experiences, using a weighted moving average index of unemployment that decays in how long ago the unemployment experiences occurred, the regression coefficient for consumption on past unemployment experiences is negative after controlling for all observables (including current assets of the household) (Malmendier and Shen 2024).

These two facts might appear to contradict each other, but they do not for several reasons. One, I will especially emphasize is the difference in focus on marginals versus averages for the consumption function—that is the difference between the slope of the consumption function and its level.

Marginals versus Averages Ganong et al. (2024) is about differences in changes in consumption between two groups (the change in consumption before and after a shock occurs). In contrast, Malmendier and Shen (2024) is about differences in average savings levels (without differences over time). The first deals with the slope of a line that connects consumption policy points, whereas the other deals with the policy level itself, conditioned on different historical experiences.[6]

6. Other differences across the two papers include:

I propose a mechanism, reinforcement learning, that is consistent with both facts. It will generate a decrease in average consumption levels and an increase in MPCs (a shift downward in the consumption policy and a steepening) after periods of unemployment.

The project aims to explain each fact based on a single learning-based mechanism. This approach ensures that agents do not require ex-ante heterogeneity to explain each result. Instead, experiences are sufficient to generate the ex-post heterogeneity observed in the data.

3 Model

Agents are infinitely-lived and discount future utility at rate $\beta < 1$. Each agent decides the fraction of their cash-on-hand to consume or save in a risk-free asset with a fixed return R per period. Each agent is fully rational each period in all variants of the model, exactly making the optimal choice, given current estimates.

An agent's cash-on-hand comprises income from two sources: a stochastic labor income process that evolves independently over time and is independent of the agent's choices and a deterministic and perfectly controllable financial income deriving from the agent's asset holdings this period.

Notation Denote cash-on-hand x, share of cash-on-hand consumed ψ, savings choices a', labor income y, risk-free assets a, and fixed return on assets per period $R > 1$. Throughout the paper, prime notation ($'$) will denote next-period values. Current-period variables appear without primes, next-period with primes. Time subscripts are avoided except where explicitly tracking multiple periods is necessary. I denote savings as a' because savings chosen today become assets tomorrow.

Timescale: The timescale for each claim is very different. The first claim spans from a few months to two years. The second covers several years to as much as forty years.

Asset Controls: The first claim does not control current assets, whereas the second aims to do so. Both approaches have implications that might introduce possible confounders or selection effects. For example, in the first model, high state-dependent consumption could be concentrated among certain households, leading to the low asset holdings observed. However, as emphasized already and noted by Ganong et al. (2024), due to the focus on marginals, this would not be sufficient to explain the gap of 0.24 in MPCs, since both households should have minimal differences in MPCs being far away from the borrowing constraint under a full-information rational expectations model without behavioral effects.

Group Definitions: The first claim concerns households that previously held low liquid assets versus those that held high liquid assets. The second pertains to previous unemployment experiences. Although correlated, these two definitions are not identical. So far, my analysis has treated both claims as referring to past unemployment. Other explanations might introduce further degrees of freedom (e.g., intergenerational wealth), but these may be difficult to justify or may complicate the model substantially.

Definitions and Relationships Cash-on-hand is defined as $x \coloneqq y + Ra$, consumption is the share of cash-on-hand consumed $c = \psi x$, and savings is the residual, $a' = (1 - \psi)x$.[7] Lastly, I note that $a'_t(a_t, y_t) = a_{t+1}$, that is the agent's a' choice today, given a realization of assets and income, becomes the agent's realized assets for tomorrow. Savings choices are made before knowing the realization of tomorrow's income state y', which, conditional on today's income, is the realization of a Bernoulli random variable.

Income Process Income transitions are captured by a first-order Markov transition matrix P, such that $y' \sim P_{yy'}$. Where relevant, I denote the transition probability $p(y' = y_i | y = y_j)$ as p_{ij}, $i, j \in \{e, u\}$, where y_e is the high-income state (employment) and y_u is the low-income state (unemployment). The model assumes $\beta R < 1$ so that the rational information expectations agent has no motive to accumulate assets beyond smoothing consumption in response to risky shocks and ensuring sufficient assets to avoid hitting their borrowing constraint in the future (precautionary savings). I explicitly consider two types of agents: an agent referred to as the *rational benchmark* with complete and perfect information about their decision problem and that solves their decision problem optimally, and a learning agent, who I refer to as the *reinforcement learner*, that lacks this information, instead using reinforcement learning to estimate their expected-value function each period.

3.1 Decision Problems

Rational Benchmark In the rational benchmark, the agent's consumption choice can be represented via the Bellman equation.

$$\begin{aligned} V(a, y) = \max_{c} \; & u(c) + \beta \, \mathbb{E}_{y'|y} \Big[V(a', y') \Big], \\ & a' \geqslant \underline{a} \\ & c = Ra + y - a' \\ & c > 0 \\ & y' \sim P_{yy'}, \end{aligned} \tag{1}$$

where $u(c)$ is the agent's utility function for a strictly positive consumption level c, $V(a, y)$ is the value function at the agent's state variables, the asset, income pair a, y, and $\mathbb{E}_{y'|y}$ is the expectation over the distribution of the income tomorrow given today's income state. I set $\underline{a} = 0$ to capture the agent's no-borrowing constraint.

7. Note that in this definition, I implicitly assume that agents want to consume as much as possible; mathematically, utility is strictly monotone in consumption and economically, agents' preferences satisfy local non-satiation.

Simulated Agent Under reinforcement learning, the model assumes the agent cannot exactly solve for their value function via iteration to a fixed point. This could be due to some combination of the fact that the agent does not know their income process $P_{yy'}$, or even the family of distributions their income process comes from, or cannot solve this fixed-point problem due to computational complexity. However, they can still evaluate their utility function $u(c)$. This is intended to reflect a real-world setting where (1) a household has considerable incomplete information about their income process, and (2) the household has not been able to get income data from others.

Because V is defined as the fixed point of the Bellman operator—which depends on knowing $P_{yy'}$ to evaluate the right side of, which the agent is unable to evaluate exactly, and $c^*(a, y)$ as the optimal consumption policy that satisfies this functional equation—the agent can no longer solve exactly for V. The model assumes that agents do not even know the form of $P_{yy'}$, which rules out anticipated utility-based learning, correctly specified regression, etc. This is a strong assumption in that it assumes agents have no knowledge about the underlying income process (rather than just misspecified knowledge or some form of sampling-based estimation) and that agents cannot solve this fixed problem. I further assume that agents learn solely from their own observations, ruling out the possibility of learning from other agents.

Agents have a perceived expected value function (EV), which I will denote as $\widehat{EV}(y, a'; \phi)$. The parameter ϕ denotes a set of parameters associated with a functional approximator. For example, in the case of a neural network, which I will use, these will represent the hidden weights of each layer, while for a polynomial these represent the coefficients on each term.

Given this perceived EV, I can define $Q(\cdot)$ associated with any a' as follows

$$Q(a, y, a'; \phi) = u(Ra + y - a') + \beta \widehat{EV}(y, a'; \phi) \tag{2}$$

Q gives the expected value of taking arbitrary *feasible* savings choice a' given the estimate $\widehat{EV}(\cdot)$. The agent then forms their consumption policy as

$$a'^*(a, y; \phi) \coloneqq \arg\max_{a'} Q(a, y, a'; \phi) \tag{3}$$

$$c^*(a, y; \phi) \coloneqq Ra + y - a'^*(a, y; \phi) \tag{4}$$

The agent learns by updating its conditional expected value $\mathbb{E}_{y'|y}[V]$ and using this, in turn, to update their Q.

3.2 Learning

This section describes how actions unfold and learning occurs in each period. The agent enters the current period with an asset level a and an income level y. For expositional clarity, I subdivide each period into two stages, day and night. During the day, the agent chooses a'.During the night, the agent learns the realization of y'. At that point, he also observes flow utility u and computes their updated $\widehat{EV}$, which in turn updates Q and the agent's policy. After updating their policy, at the beginning of the following day, the agent executes the updated policy.

Day During the day, they choose Eq. (3).

Night At night, the agent experiences their flow utility $u(Ra + y - a'^*(a, y; \phi))$ and learns their following period's income realization y'. Using this, they forms the following *empirical* estimate of the Q-value.

$$Q^{empirical}(a, y, a', y'; \phi) \coloneqq u(Ra + y - a') + \beta \max_{a''} Q(a', y', a''; \phi). \tag{5}$$

The temporal difference (TD) error is defined as:

$$\epsilon(a, y, a', y'; \phi) \coloneqq Q^{empirical}(a, y, a', y'; \phi) - Q(a, y, a'; \phi) \tag{6}$$

Using the definition of $Q(a, y, a'; \phi)$ and plugging in for utility, the flow utilities for the present period $u(Ra + y - a')$ cancel, yielding

$$\epsilon(a, y, a', y'; \phi) = \beta \left[\max_{a''} \Big(u(Ra' + y' - a'') + \beta \, \widehat{EV}(y', a''; \phi) \Big) - \widehat{EV}(y, a'; \phi) \right]. \tag{7}$$

I use this ϵ to adjust ϕ by gradient descent on the squared TD error. The temporal difference error ϵ captures the difference between agents' realized value from taking an action and agents' prior expectation of the value of the choice. It therefore represents the felicity "surprise" the agent has received from a particular choice. The agent uses ϵ to adjust their parameters. They then update their parameters ϕ governing the parametric estimator as

$$\phi' = \phi - \lambda \nabla_\phi \widehat{EV}(y, a'; \phi) \, \epsilon(a, a', a', y', a''; \phi), \tag{8}$$

and revise the policy as

$$a'^*(a, y; \phi') = \arg\max_{a' \in [\underline{a}, \overline{a}]} Q(a, y, a'; \phi'), \tag{9}$$

$$c^*(a, y; \phi') = Ra + y - a'^*(a, y; \phi'). \tag{10}$$

In simulation, unemployment experiences cause consumption to decrease and savings to increase, while the estimated MPC increases. For the remainder of this paper, I use the following definition for MPC.

Definition 3.1. The marginal propensity to consume out of a transfer τ is

$$MPC(a, y; \phi) \coloneqq \frac{c^*(a + \tau, y; \phi) - c^*(a, y; \phi)}{\tau}. \tag{11}$$

3.3 Parameterization of EV

The model assumes the agent approximates $\widehat{EV}$ with a two-layer neural network with rectified linear unit activations, abbreviated as ReLU. I use this due to its desirable properties in approximating complex functions efficiently, because it is the most popular existing activation function (Rasamoelina et al. 2020), and because it has established guarantees both on convergence to a global optimal fit in the mean-squared error sense in the regression case, and on convergence to an optimal policy in the Q learning case (Xu and Gu 2020). The ReLU network is trained in a supervised fashion to imitate the rational benchmark at time 0 to achieve a precise fit to its value function. This specification abstracts from experimentation motives. This is in the spirit of the anticipated utility formulation adopted by much of the learning literature. Importantly, by using a parameterized estimator of $\widehat{EV}$, I partially capture the benefits of future exploration on policies to the extent these can be approximated in the $\widehat{EV}$ estimate. However, these only approximately do so, and the requirement to always use exactly optimal policies rules out the possibility of explicit exploration heuristics, designed to represent the option value from how observations will affect the future evolution of uncertainty given policy choices.

3.4 Smoothed Neural Network

A neural network with N layers can be represented as follows. Given an input X^0, each layer applies the recursion:

$$X^{i+1} = \sigma^i\Big(A^i X^i + b^i\Big), \quad \forall i \in \{0, \ldots, N-1\}, \tag{12}$$

where $X^0 = [y, a']^T$ serves as the argument for $\widehat{EV}(y, a')$. A^0 is an $h \times 2$ matrix of hidden weights, and X^1 is an $h \times 1$ vector of hidden biases. Each A^i for $i = 1, \ldots, N-1$ is an $h \times h$ matrix, while A^N is $1 \times h$, producing the agent's final estimate of $\widehat{EV}(y, a')$ for an input pair of (y, a'). The parameter vector ϕ is the concatenation of all these weights and biases as

$$\phi = \begin{bmatrix} \text{vec}(A^0) \\ \text{vec}(A^1) \\ \vdots \\ \text{vec}(A^{N-1}) \\ b^0 \\ b^1 \\ \vdots \\ b^{N-1} \end{bmatrix} \tag{13}$$

σ^i is an activation function applied element-wise. For all layers, I use the ReLU function,

$$\text{ReLU}(x) = \max(x, 0).$$

Because ReLU is non-smooth, the resulting network is also non-smooth, as are the agent's policies. Non-smooth policies may generate jumps in consumption. To ensure strict monotonicity in consumption, in line with the economic literature, the model assumes a polynomial fit to the output of the value function produced by the ReLU network, which the agent takes as their value function values when considering policies at states other than their current state in order to generate monotone policy functions. However, when just computing consumption at their exact current state/realized histories, the agent does not utilize this fit, including when computing realized MPCs out of transfers. This approach retains some of the flexibility of ReLU while smoothing the final output, in particular ensuring strict monotonicity of the consumption policy, which otherwise exhibits local oscillations. In Section A.2, I present the results of the neural network fit alone.

4 Parameterization

All dollar amounts are normalized to quarterly income units for an employed agent at the time of the COVID-19 stimulus payment. Income is chosen to match Ganong et al. (2024)'s average "primary" income. This represents 0.44 of total household income in their sample and captures the income earned while employed that is not earned while unemployed and

still waiting for UI benefits. This is estimated at a nominal basis of $9,042 on a quarterly basis, which becomes the simulation's baseline unit of 1. Additional unemployment benefits, on a quarterly basis, were $7,200 ($600 a week). In Ganong et al. (2024)'s sample, these come out to 34.5% of total average income, 78.4% of primary income. Unemployment is set at the replacement rate of 47.27% or around $4,275 on a quarterly basis, again following Ganong et al. (2024). I consider possible savings choices between 0 and 4.5 times the agent's average quarterly income (approximately between 0 and $90,000). I uniformly subdivide this grid into 8,750 feasible savings choices for policy decisions or approximately $10.25 increments of savings, on a quarterly basis.

4.1 Initialization

Neural Network During training for the initial fit, the ReLU network takes the average MSE over minibatches of size 32 collected from 500 savings choices and two income choices and does stochastic gradient descent on ϕ with ADAM, until these attain a tight fit on the rational solution value function globally. I initialized the agents' weights A^i for the neural network, before the fit from i.i.d. draws of a Gaussian with 0 mean and $\sqrt{2}$ standard deviation, as is required for the theoretical guarantees of Li and Yuan (2017) to apply. I initialized the agents' biases b^i at 0.[8]

The network was trained to minimize the maximum error over held-out asset values not included in the training sample. I stopped training when the maximum error of the neural network fit fell below 2.5×10^{-3}, below which achieving a given error at the current neural network size became noticeably more challenging.

Agents' initial asset holdings are drawn from an empirical distribution calibrated to the 2016 Survey of Consumer Finances (SCF). I use median estimates in the interquartile ranges to get financial asset values at the 12.5th, 37.5th, 62.5th, 87.5th, and 95th percentiles, which are then normalized to be multiples of quarterly income, the simulation's numeraire. For the lower 87.5 of the distribution, I apply piecewise linear interpolation between these data points to estimate the inverse CDF. For the upper tail, I fit a Pareto distribution with shape parameter α estimated from the ratio of the 95th to 87.5th percentile values using $\alpha = -\ln(1 - 0.95)/\ln(Q_{95}/Q_{87.5})$, where Q_p denotes the asset value at percentile p. This hybrid approach captures both the empirical distribution of typical households and the heavy-tailed nature of wealth concentration, with the Pareto tail accounting for extreme wealth

8. Xu and Gu (2020) also require a Gaussian initialization for their results to go through. They require a slightly different variance to get their theoretical analysis and particular convergence rates. I prefer Li and Yuan (2017), because Xu and Gu (2020) also require the assumptions of the network to be sufficiently overparameterized, which can be difficult to tell if it is satisfied.

holdings above the 87.5th percentile. Unlike the raw SCF, I truncate agents' distribution at 0, as no borrowing is allowed in the simulation, so I force the 0th percentile to be 0. I do not adjust other percentiles to compensate due to lacking the 0th percentile summary statistic. Due to the piecewise linear interpolation, my method may oversample the first octile of the distribution, slightly compared to the population distribution. Sampling occurs via the inverse CDF method, drawing uniformly from the interval [0,1] and applying the inverse linearly-interpolated CDF to obtain asset values. This ensures that the initial asset distribution reflects the empirical distribution of household assets while maintaining the no-borrowing constraint. All agents share the same initialization of weights, with the weights being chosen based on a combination of trying to achieve the minimal fit in terms of least-square error to the expected value function, crucially on held-out validation data (to avoid overfitting), and to ensure that the initial fit exhibits desirable properties such as smoothness and monotonicity. The latter I have not formally enforced and instead chose based on visual inspection of the fit, a weakness of the approach. There may be slight selection bias incurred as a result.

As a result of local fluctuations continuing to be present in policy, as mentioned previously, I also assume that subsequently the agent takes the neural network fit and fits a moderate-order polynomial *on the expected value function* after learning to achieve global smoothness. This helps to enforce the monotonicity of the consumption policy. At present, the model assumes the agent only uses these for considering counterfactuals, such as policy functions and not for learning. That is, explicitly, when computing marginal propensities to consume out of transfers and consumption and savings responses from cross-sectional data, smoothed values are not utilized. Instead, at present, these are solely used for displaying policies within an individual agent and the cross-sectional policy and MPC distributions.[9]

Agents are initialized at the same set of weight parameters so learning is the sole source of diversity of choices beyond assets. During the actual simulation, for the neural network, no minibatches or parallel evaluations are used. I continue to use ADAM for the learning step each period.

9. As an alternative, I considered using a local smoothness criterion or local quadratic variation in the derivative of the fitted value function to minimize perturbations. Unfortunately, this only minimized local perturbations in the fit on the expected value function and not on the policies themselves, which is what I care about when enforcing monotonicity. These instead had to do with local, small fluctuations in the slope of the expected value function in assets relative to the slope of the utility function that violated strict concavity of the value function. I also considered explicitly incorporating the polynomial fit into the learning process, but this led to limited variation in policies over time at low order polynomials and memory issues at higher order polynomials that prevented the fit computationally.

4.2 Parameters Summary

Key parameters are summarized in Table 4.1.

4.3 Experiment with Extreme Shocks

To highlight the mechanism, I conduct stylized experiments with two agents simulated in parallel. Two agents were simulated in parallel, with both agents starting with employment, the same asset level (the median). Then, in the subsequent period, one agent receives a series of unemployment realizations for four periods, while the other receives a series of employment realizations. The intent is to generate opposite extreme experiences in order to better observe how the impact of employment or unemployment realizations is affecting the agents' policies. Both are plotted relative to the rational benchmark. Figure 4.1 shows that for sufficiently high asset levels, both agents consume more than the rational benchmark. However, for low asset levels, the agent with more unemployment experiences exhibits significant scarring, represented by a shift down in the consumption policy curve. By contrast, the agent with more unemployment experiences consumes more than the rational benchmark everywhere, an "anti-scarring" effect due to too low risk. By comparison, see Fig. A3 for the initial fit.

Close to the no-borrowing constraint, an agent with more unemployment experiences appears to slightly decrease their consumption, while employed, more than the agent with employed experiences does while unemployed. This is a nice sanity check that the value-based mechanism works, since risk aversion and the borrowing constraint imply asymmetrically larger losses in the value function for the agent with more unemployment experiences under equivalent policies, partially reflected in the curvature of the consumption policy function. By contrast, the slope of the consumption policy for the agent with more unemployment experiences is generally steeper than the agent with more employment experiences, which is suggestive of a higher marginal propensity to consume (MPC) for the agent with more unemployment experiences; the agent's MPC is estimated as a local slope between any two points on the consumption policy curve separated by the transfer quantity. I verify this with a graph of the MPCs as a function of assets in Fig. 4.2 and see that in general, for values close to 0 assets, the agent receiving more employment experiences has higher MPCs, but for values further away, the agent receiving more unemployment realizations has higher MPCs.

Table 4.1: Key Configuration Parameters

Parameter	Value	Notes
Simulation Parameters		
number of agents	50	
number of time periods	50	
learning rate (simulation)	1.1×10^{-3}	
optimizer (simulation)	ADAM	less sensitive to learning-rate choice
polynomial degree	5	Calibrated, time 0 fit
# evaluation grid for fit	50 equally spaced	
max asset level (pre-shock)	4.5	
min asset level	0.0	
learning rate decay	$\mathcal{O}(t^{-1/2})$	Robbins and Monro (1951)
observation period	quarterly	
learning period	quarterly	
β	0.9703 (Q), 0.99 (M)	Ganong et al. (2024)
interest rate	0.985% (Q), 4% (Y)	Ganong et al. (2024)
initial asset distribution	piecewise linear inverse CDF, Pareto tail	SCF 2016, financial wealth
UE, UU	(0.392, 0.608) (Q) (0.28, 0.72) (M)	Ganong et al. (2024)
EE, EU	(0.939, 0.0607) (Q) (0.972, 0.028) (M)	Ganong et al. (2024)
replacement rate	0.472 ($4,268, 2020 dollars)	Ganong et al. (2024)
income value	1.0 ($9,042, 2020 dollars)	Ganong et al. (2024)
shock size (fraction of income)	0.784 ($7,200, 2020 dollars)	Ganong et al. (2024)
Neural Network Configuration		
size, hidden dimension	80	
# of hidden layers	2	
activation	Rectified Linear (ReLU)	
Grid Configuration		
asset grid size	500 states	

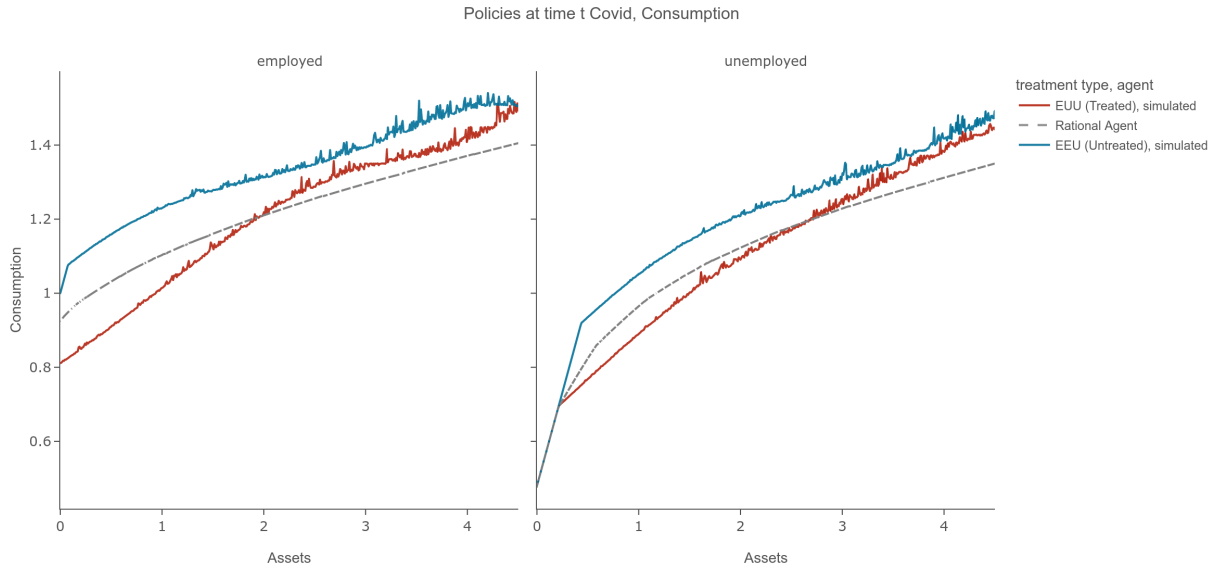


Figure 4.1: Consumption policies at $t = 5$ for the experiment with repeated employment with one agent receiving a series of unemployment realizations and the other receiving a series of employment realizations. The blue line is the agent with unemployment experiences, while the red line is the agent with employment experiences. The dotted gray line is the rational agent's policy. The figure on the left represents the agent's consumption policy as a function of their assets (x-axis) and income while employed y_e, while on the right it represents the agent's consumption policy as a function of their assets and income while unemployed y_u. Value functions use a polynomial fit.

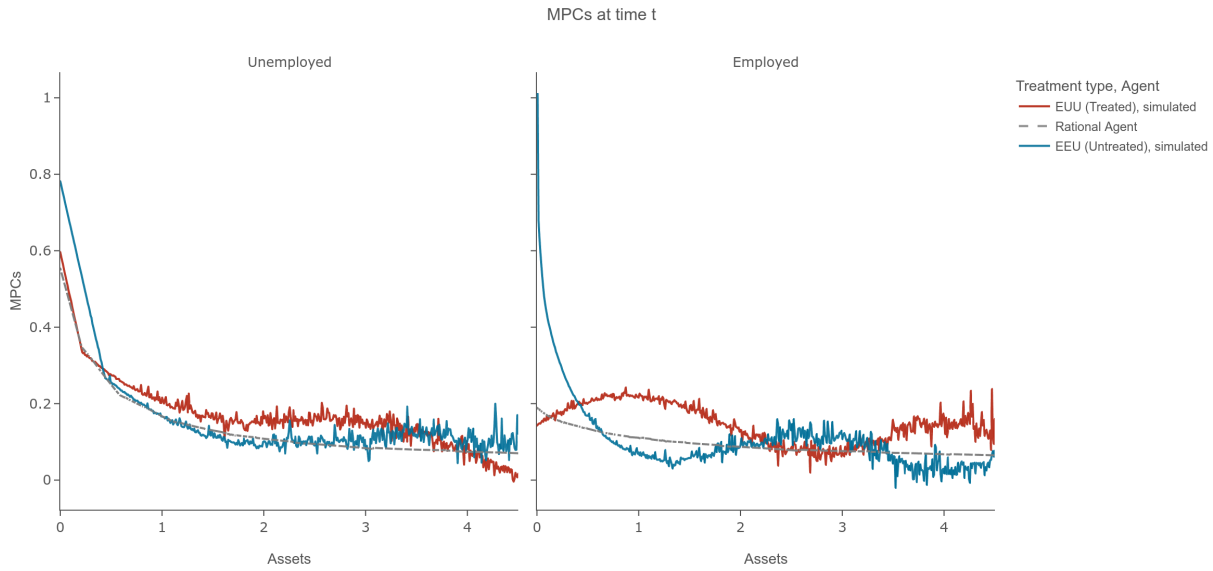


Figure 4.2: MPCs for the experiment with repeated employment and unemployment realizations after $t = 5$ periods, where one agent receives a series of unemployment realizations and the other receives a series of employment realizations. The blue line is the agent with unemployment experiences, while the red line is the agent with employment experiences. The dotted gray line is the rational benchmark's policy. The figure on the left represents the agent's marginal propensity to consume as a function of their assets (x-axis) and income while employed y_e, while on the left it represents the agent's consumption policy as a function of their assets and income while unemployed y_u. This uses the polynomial fit to the value function generating local oscillations from the polynomial.

5 Experiments

From our initial fit, I then let the agent learn for 50 periods, or about 12.5 years. I then repeat the analysis of Ganong et al. (2024) and Malmendier and Shen (2024).

5.1 Marginal Propensities to Consume

I repeat the primary setup of Ganong et al. (2024) with minor modifications.

Empirical Specification The primary identification strategy of Ganong et al. (2024) employs a difference-in-differences design exploiting administrative delays in unemployment benefit processing. The treated group consists of workers who became unemployed at the end of March 2020 and received benefits in April, while the control group includes workers who also became unemployed in March but experienced processing delays and did not receive benefits until June. The two-stage specification is:

First stage:

$$y_{i,t} = \alpha + \beta \text{Post}_t \times \text{Treat}_i + \text{Treat}_i + \text{Post}_t + \varepsilon_{i,t}$$

Second stage:

$$c_{i,t} = \phi + \text{MPC} \times \hat{y}_{i,t} + \text{Treat}_i + \text{Post}_t + \varepsilon_{i,t},$$

where $t \in \{\text{March}, \text{May}\}$, Treat_i indicates receipt of benefits in April, and Post_t indicates May observations. The coefficient MPC captures the marginal propensity to consume out of unemployment benefits.

To examine heterogeneity by liquidity status, households are split by their 2018 median liquidity buffer, defined as $(\text{checking balance}_{it} - 0.5 \times \text{spending}_{it})/\text{spending}_{it}$. This pre-pandemic measure avoids endogeneity concerns while capturing persistent differences in savings behavior. The paper finds MPCs of 0.53 for low-liquidity households versus 0.29 for high-liquidity households, with effects persisting even after benefit expiration.

Crucially, because I do not have unobserved characteristics and have direct access to asking agents about what they would have done if they had received a transfer at any given time, I do not need to repeat most of the regression of Ganong et al. (2024), which aims to control for unobserved heterogeneity and indirectly estimate group average MPCs. Instead, I "freeze" agents (so that they do not learn temporarily from new experiences) and directly provide agents with and without a transfer at time t and see how their consumption choices would vary given their value functions and policies.

Crucially, I still split agents at time 0 by whether they are above or below the median asset holding at time 0 in the simulation. Agents who fall above are called *high liquidity*

while those below are *low liquidity*. Because there is no exact analog of bank balance sheets, assets are used as the measure instead. I then compute group average MPCs out of transfers eight periods later, the same duration as the beginning of Ganong et al. (2024) until the payments. MPCs are computed by shifting agents' assets by the size of the COVID stimulus UI expansion, showing agents both asset levels at the same income level and evaluating how consumption changed relative to how assets changed. Because Ganong et al. (2024) is focused on MPCs for the unemployed, I only examine agents while they are unemployed. I consider all MPCs for any agent between quarters eight and nine, which covers the same duration of four months of FPUC transfers from Ganong et al. (2024).

5.2 Results for Marginal Propensities to Consume

In Table 5.1, I show the results of the MPC calculation for agents at eight periods after the initial period, when unemployed. In Table 5.2, I repeat this exercise but now with agents split up at different points in the simulation by liquidity level.

Table 5.1: Model versus Empirical Marginal Propensities to Consume (MPCs)

Liquidity Type	Model MPC	Empirical MPC[a]
Low	0.501	0.53
High	0.343	0.29

Notes: This table compares MPCs for unemployed agents, categorized by their initial asset level (liquidity) relative to the median. The MPC is measured for agents unemployed between periods $t = 8$ and $t = 9$.
[a] Empirical values and methodology replicate the "preferred specification" from Ganong et al. (2024).

Table 5.2: Model MPCs by Periods Experienced and Liquidity Type

	Empirical[a]	Model: Periods Experienced in Simulation Prior to Classification (quarters)								
Liquidity Type	Benchmark	0	5	10	15	20	25	30	35	40
Low	0.530	**0.501***	**0.589*****	**0.585*****	0.528	0.571	0.531	0.516	0.454	0.559
High	0.290	0.343	0.334	0.400	0.470	0.520	0.492	0.522	0.468	0.587
Difference	0.240	**0.158***	**0.255*****	**0.185*****	0.058	0.051	0.038	-0.007	-0.014	-0.028
		(2.15)	(4.44)	(5.35)	(0.77)	(0.65)	(0.50)	(-0.13)	(-0.19)	(-0.97)

Notes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Null hypothesis is no difference between low and high.
MPCs measured 8 quarters after liquidity classification, including for empirical. Differences show low minus high liquidity MPCs with t-statistics in parentheses (Welch's t-test with unequal variances).
[a] Empirical values and methodology replicate the "preferred specification" from Ganong et al. (2024).

5.3 Scarring

I repeat the analysis of Malmendier and Shen (2024) by computing a statistic for each agent at each time period that varies between 0 and 1 based on how recent past unemployment experiences were.

Empirical Specification An index of 0 indicates no unemployment experiences in an agent's past, while if an agent only experienced unemployment for all past periods, they would have an index of 1. Experiences further back in the past are down-weighted in how many time periods ago they occurred. The unemployment index is based on indicators for unemployment weighted by

$$w(t,k) = \frac{(t-k)}{\sum_{k=2}^{t-1}(t-k)}.$$

The numerator represents the relative importance while the denominator normalizes the weights to sum to 1. t is the current time period and k is the number of time periods ago the unemployment experience occurred. Ganong et al. (2024) begins with time period 1. The paper excludes the most recent two time periods (today and yesterday) to exclude possible shorter term, non-experience-based effects, which I also follow, though it is important to note that while I operate on quarterly frequency data, Malmendier and Shen (2024) work with biennial frequency data. Time periods in the authors' paper are calibrated to represent quarterly data in simulation and biennial samples of annual data in the data from PSID, with a maximum of 18 years of data from PSID (1999–2017; 9 periods), and 50 years in a lifecycle simulation (200 quarters). The current value of the index for an agent is given by p_{it} defined as

$$p_{it} = \sum_{k=2}^{t-1} w(t,k) W_{t-k},$$

where W_{t-k} is an indicator for whether an agent was unemployed at time $t-k$. The specification for the regression I utilize, with quarterly frequency, is based on the regression:

$$C_{it} = \alpha + \phi UE_{pers,i,t} + \gamma' x_{it} + \varepsilon_{it}$$

where C_{it} is the consumption of agent i at time t, ϕ is the coefficient on the unemployment experience index, $UE_{pers,i,t}$, and x_{it} are controls that consist of (a_{it}, y_{it}) asset income pairs for us.

Note that the original paper uses the full specification

$$C_{it} = \alpha + \psi UEP_{it} + \beta UE_{it} + \gamma' \mathbf{x}_{it} + \eta_t + \delta U_{st} + \varsigma_s + v_i + \epsilon_{it},$$

where x_{it} is a series of controls including demographic, income, and include: $\log(y_t)$, $\log(\log(y_t))$, $\log(y_{t-1})$, $\log(\log(y_{t-1}))$, $\log(a_t)$, $\log(\log(a_t))$ controls (shifted first from their minimal value, so that log is well-defined); UEP_{it} the personal measure of unemployment; η_t a time fixed effect; v_i an agent fixed effect; ςs a state fixed effect; U_{st} the state unemployment level at the present time; and UE_{it} the macroeconomic unemployment index, constructed identically to how I constructed the individual level index, except using the state-level unemployment percentage for W_t at time t. Given that I do not have confounding present in the simulation beyond asset and income experiences, as agents all begin with the same weighting, I do not include these additional controls in the regression. Similarly, because I do not have extreme skewness in wealth, I control for wealth only via the asset level a_{it} and do not include $\log(a_{it}), \log(\log(a_{it}))$ controls or income controls. Finally, Malmendier and Shen (2024) truncate agents at the tenth percentile of wealth and the ninetieth percentile of total income. I similarly do not do this.

5.4 Results for Scarring

In Table 5.3, I show the results of the regression of consumption on the unemployment experience index, as well as the average MPCs for agents with different unemployment experience indices. With approximately 4,700 observations (over multiple time periods), I find about a -0.0378% drop in consumption for a 1% increase in the unemployment index. This effect is small but significant at the 1% level; it is, however, about an order of magnitude smaller than Malmendier and Shen's estimate at the biennial frequency of -0.280. Given their difference in frequency, they may not be directly comparable in magnitude. When not controlling for assets, I find a larger drop in consumption of around 0.0884% at the 1% significance level. This suggests that both assets and unemployment experiences play an important role in determining consumption changes.

5.5 Experiments Summary

When I run the simulation, I get results qualitatively consistent with both sets of facts. In particular, after 50 periods of running the simulation, I appear to get overall higher consumption levels and higher MPCs than the rational case over 50 time periods. But once selecting for past unemployment experiences, I get evidence of scarring effects.

Table 5.3: Regression of personal unemployment experience on consumption with and without controls for asset levels.

	Dependent variable:	
	consumption	
	(1)	(2)
UEP_t	-0.0884^{***}	-0.0378***
	$(-0.108, -0.069)$	$(-0.052, -0.024)$
assets_t		0.1265^{***}
		(0.123, 0.130)
income_t	1.0229^{***}	0.9869^{***}
	(1.020, 1.026)	(0.985, 0.989)
Observations	4,692	4,692
R^2	0.658	0.835
Adjusted R^2	0.658	0.835
Log Likelihood	5,173.100	6,878.400
AIC	$-10{,}340.000$	$-13{,}750.000$
BIC	$-10{,}330.000$	$-13{,}730.000$
F Statistic	$9{,}015^{***}$ (df = 1; 4690)	$11{,}830^{***}$ (df = 2; 4689)
Note:	*** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$	
	95% confidence intervals in parentheses	

6 Discussion

The agents' MPCs quantitatively and qualitatively replicate the findings of Ganong et al. (2024). For Malmendier and Shen (2024), however, I find only broad qualitative agreement with simulation-based estimates of effect size generally about an order of magnitude smaller than the empirical effect size. However, one important thing to note is that the simulation generated data and Malmendier and Shen (2024) occur at different frequencies, with the simulation's data being quarterly and Malmendier and Shen's being biennial, representing longer durations and, hence, more time accumulated in the learning process. Unlike the higher perceived probability of unemployment like Malmendier and Shen (2024), I am able to generate both higher MPCs and lower average consumption levels. See Section A.5 for further details and a simulation result that shows that both MPCs and average consumption levels decline under the current parameterization when solely increasing probabilities of unemployment.

7 Conclusion

In this work, I show that agents utilizing reinforcement learning replicate the key findings of Ganong et al. (2024) and Malmendier and Shen (2024) in a model-free setting. This paper's learning-based mechanism, deep reinforcement learning, is similar to Malmendier and Shen (2024) in that it is experience-based, but differs in that it explicitly utilizes local "surprises" in utility to have an agent adjust their behavior, without having an explicit probability distribution that is updated separately from the agents' value function. This differing mechanism enables us to both generate higher MPCs and lower average consumption levels, something not possible when an agent's perceived probability of transitioning to unemployment occurs. Instead, the expected value function is the only estimator that the agent uses. The neural network estimator of the conditional expected value does not assume anything about the underlying income distribution. As a result, I do not need prior knowledge of a particular form of the distribution, such as a two-state Markov switching process. Regardless of the form, the agent will be able to update their expected value function, using this lower-dimensional statistic to guide their behavior. The practical trade-off is that without assuming a model, I also cannot extract clear beliefs on the part of the agent over future unemployment probabilities, limiting my ability to make precise statements about the causes of agents' behavior and how their beliefs are changing over time, as well as preventing us from using belief surveys to calibrate the agent from the data.

This is a limitation that can be addressed in future work, by either including a general

model-based reinforcement learning scheme or successor representation where an agent learns some flexible approximation of their probability of reaching different future states over time, using proxies for an agent's beliefs in the form of attention-based mechanisms that tell us what information the agent learns to view most saliently, or by using a model-free distributional or Bayesian reinforcement learning approach, where the agent includes a distribution over probabilities of future utility realizations/value functions they may face.

Another limitation of this work is that, at present, while I test against the two key papers, I do not fully qualitatively match both results, nor test against empirical data. This, too, presents a future direction of inquiry. In this paper, I do not touch on theoretical results on convergence or dynamics, which would be interesting to explore. I focus solely on the behavior of a single agent. A natural question would be what economic equilibria look like when many agents using reinforcement learning interact, such as in an Aiyagari model, what equilibria are learnable, and how they differ from the rational expectations equilibrium, if it is not learned.

Finally, the assumption of polynomial smoothed value functions can possibly be relaxed by instead using a policy function to interpolate choices across policies, such as via an actor-critic approach. This also presents a clearer separation between how agents judge the value of different choices, at a small set of values and how they decide, in a potentially suboptimal manner, which actions to take. At present, both mechanisms are combined in the same neural network, which makes it difficult to disentangle the separate motivations.

Appendix: Supplementary Material

A.1 Environment

Formally, my setting can be defined as the Markov Decision Process (MDP) given by the tuple $\mathcal{M} = (\mathcal{S}, \mathcal{A}, P_{yy'}, u)$, where the state space is given by all possible asset, income pairs, $\mathcal{S} = \{a, y\}_{a \in \mathbb{R}^{++}, y \in \{y_u, y_e\}}$ and $\mathcal{A} = \{c\}_{c \in \mathbb{R}^{++}}$. Note that because of my special problem structure, the standard $P(s, s')$ for states $s \in \mathcal{S}$ reduces to $P_{yy'}$ as there is no exogenous uncertainty regarding how a' evolves. Similarly, a' is determined exactly by c and so can be omitted from the action space. The "reward" R depends solely on the agent's choice c and not on the surrounding state s, and is equivalent to $u(c) = \log(c)$.

A.2 Unsmoothed Setting

Consumption and MPCs without first using polynomial smoothing on the expected value function are given below in Fig. A1 and Fig. A2. They are broadly similar, except without guarantees of monotonicity.

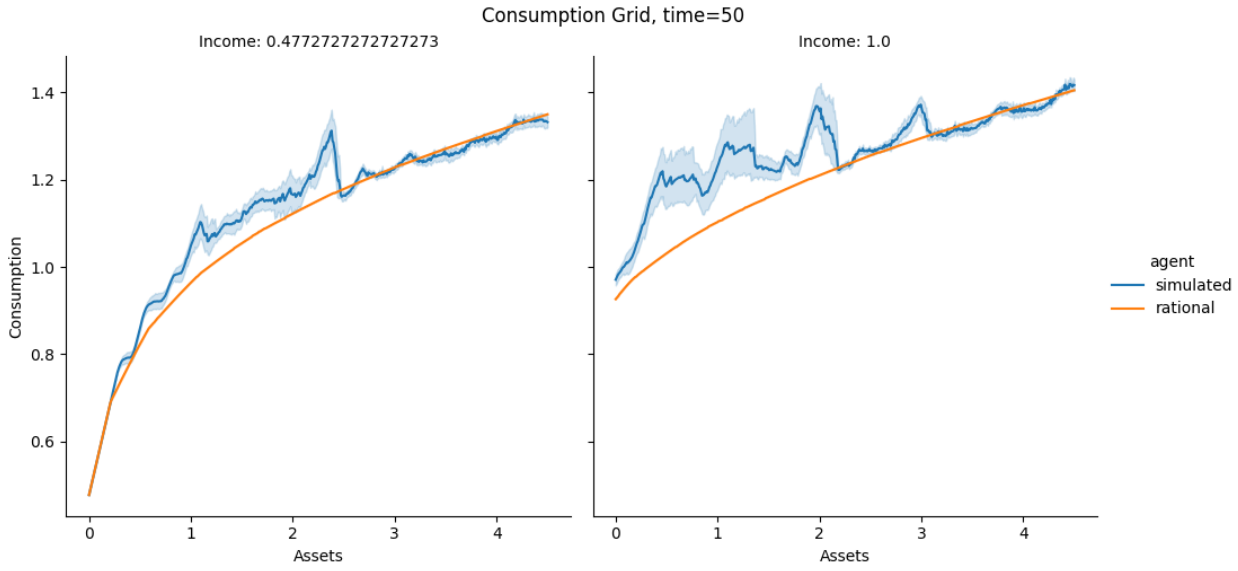


Figure A1: Consumption as a function of assets and income after 50 periods of learning. No polynomial smoothing is used. The blue line is the reinforcement learner, the orange line is the rational benchmark. 95th percentiles are plotted.

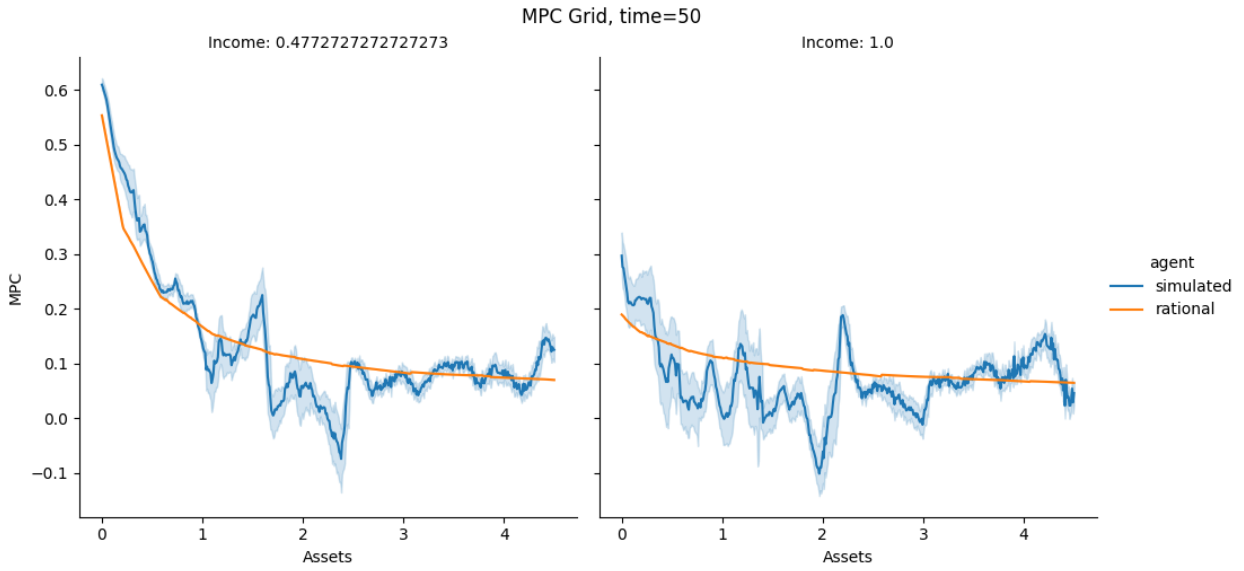


Figure A2: MPC as a function of assets and income after 50 periods of learning. No polynomial smoothing is used. The blue line shows the reinforcement learner, the orange line shows the rational benchmark. 95th percentile intervals are plotted.

A.3 Further Details on Experiment with Extreme Shocks

See Fig. A3 for the initial consumption policy fit for the experiment with extreme shocks, which shows that agents began with a tight-fit to the rational policy and with identical consumption policies.

A.4 Relaxing Markovian assumptions

Three approaches stand out:

1) directly modifying a Markovian algorithm, via a latent factor, conditional on which the problem again becomes Markovian
2) a separate learned representation of the environment;
3) a model-selection based approach

Latent Factor Representation Perhaps the most standard way is by taking all observations and capturing them in a low dimensional latent variable representing an agent's history h_t. This is commonly used in a non-Markovian setting called a Partially Observable Markov Decision Process (POMDP), where it is assumed that while observables may evolve in a non-Markovian fashion, there is some underlying latent hidden Markov state under which the model again becomes Markovian. To capture history dependence or limited observation, the agent is modified to use a recurrent neural network, and when updating their estimate for

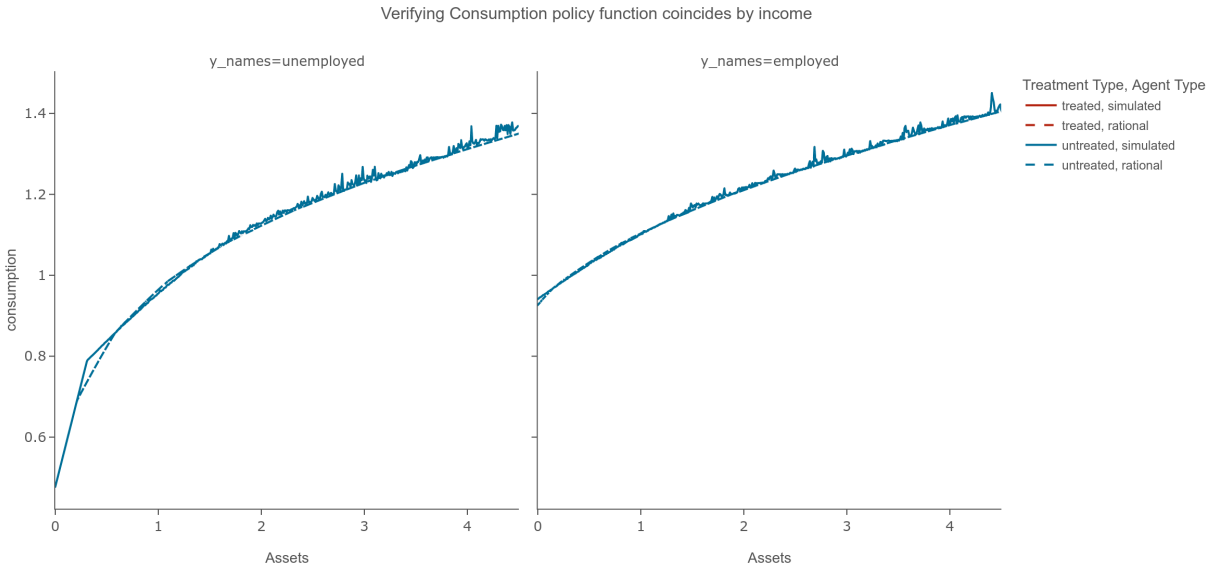


Figure A3: Consumption policies at $t = 0$ for the experiment with repeated employment with one agent receiving a series of unemployment realizations and the other receiving a series of employment realizations. The blue line is the agent with unemployment experiences, while the red line is the agent with employment experiences. The dotted gray line is the rational agent's policy. The figure on the left represents the agent's consumption policy as a function of their assets (x-axis) and income while employed y_e, while on the right it represents the agent's consumption policy as a function of their assets and income while unemployed y_u.

the continuation value, would use backpropagation through time to update their estimate of the sequence of past latent states simultaneously. More recent work has leveraged selective "attention"-based mechanisms, or diffusion-based mechanisms. See Tennenholtz et al. (2023) and Du et al. (2024) for examples. Leong et al. (2017) further show that humans neurologically use attention to selectively weight values while learning about their environment via reinforcement learning, which is a natural possible future direction of research. This is not attempted in this paper due to the additional layers of complexity and obscuring of the mechanism that would occur, and given this is meant to be an initial application of these approaches. This shares spiritual similarities with filtering, including the Kalman filter, and other estimation procedures for Hidden Markov Models.

Model of Environment A second approach would be to model the agent's environment in a flexible manner using a low-dimensional representation for observations, updating the model over time. This is what Hafner et al. (2025) does recently, one of the first at-scale model-based RL procedures, using encoded latent factor representations to capture a "world model" generating observations. Alternatively, other approaches can be used to perform some kind of model selection procedure on a set of states. Two recent examples exploring this in a setting similar to my paper are Levine et al. (2024) and Lamb et al. (2023). These are more sophisticated in that they not only find a relevant subset of states, but also identify over time which states are controllable versus exogenous.

Non-Markovian These methods can also be applied to relax Markov settings to hidden Markov models, conditionally Markovian settings in continuous time, or fully non-Markovian settings. See Whitehead and Lin (1995), Kaelbling et al. (1998), Sutton et al. (1999), and Qin et al. (2023) for further details.

A.5 Changing Income Probabilities Experiment

I study what happens under the current parameterization in response to an increase in perceived transitions to unemployment, but without approximating the conditional expected value function directly. I use the existing parameterization Table 4.1, but solely run the rational model under two specifications, one with the original income transition probabilities, and one where the UU and EU probabilities are increased by 1.5 times their original value, and then the probabilities are re-normalized to sum to 1. Crucially, for each run, I use the perceived income distribution, but then solve the policy via iterating forward to a fixed point under this perceived income distribution, as in Malmendier and Shen (2024). The results are shown in Fig. A4 and Fig. A5. This shows a tilt down globally in the consumption policy for

the rational agent and a shift down in the marginal propensity to consume, as the agent is more pessimistic about their future income prospects. This is consistent with the intuition that pessimism leads to lower consumption, as the agent is more likely to be unemployed in the future and so saves more for that eventuality. Because the agent exhibits prudence in their utility function, a lower consumption policy everywhere in turn leads to a lower local slope of the consumption policy and lower MPCs per asset level.[10]

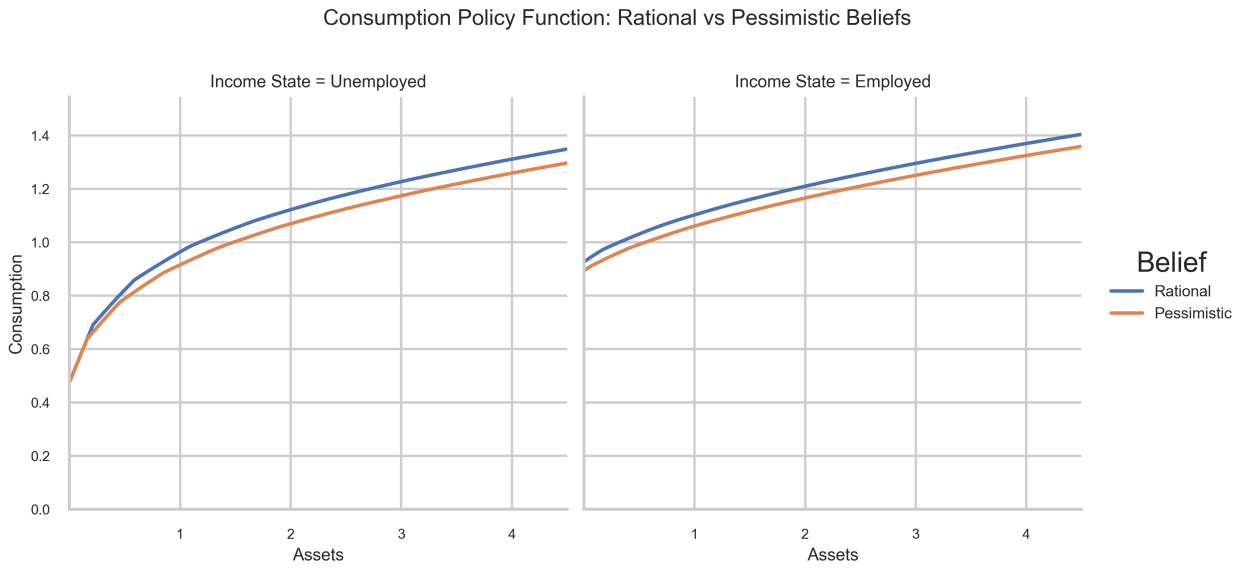


Figure A4: Orange line is rational agent solution under 1.5x pessimistic income transition probabilities, blue line is rational agent solution under original income transition probabilities. The figure on the left represents the agent's consumption policy as a function of their assets (x-axis) and income while unemployed y_u, while on the right it represents the agent's consumption policy as a function of their assets and income while unemployed y_e.

A.6 Long Run and Convergence

Policies appear to exhibit convergence in the long-run. In figures Fig. A7, the policy is examined at time 10, in Fig. A8 at period 240. Both are compared to time 0 initial fit in Fig. A6. Policies first drift away from rational solution before converging back.

10. Even without prudence, this could occur strictly due to the presence of the borrowing constraint. However, the author knows of no theoretical result that guarantees this at the time of writing.

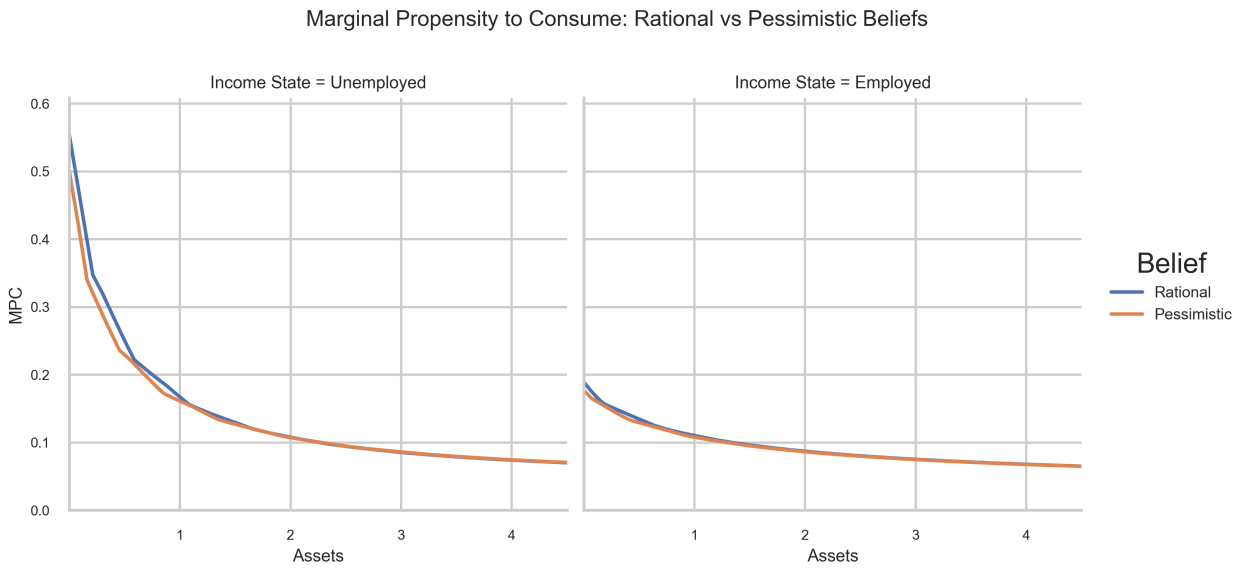


Figure A5: Orange line is rational agent marginal propensity to consume under 1.5x pessimistic income transition probabilities, blue line is rational agent solution under original income transition probabilities. The figure on the left represents the agent's consumption policy as a function of their assets (x-axis) and income while unemployed y_u, while on the right it represents the agent's consumption policy as a function of their assets and income while unemployed y_e.

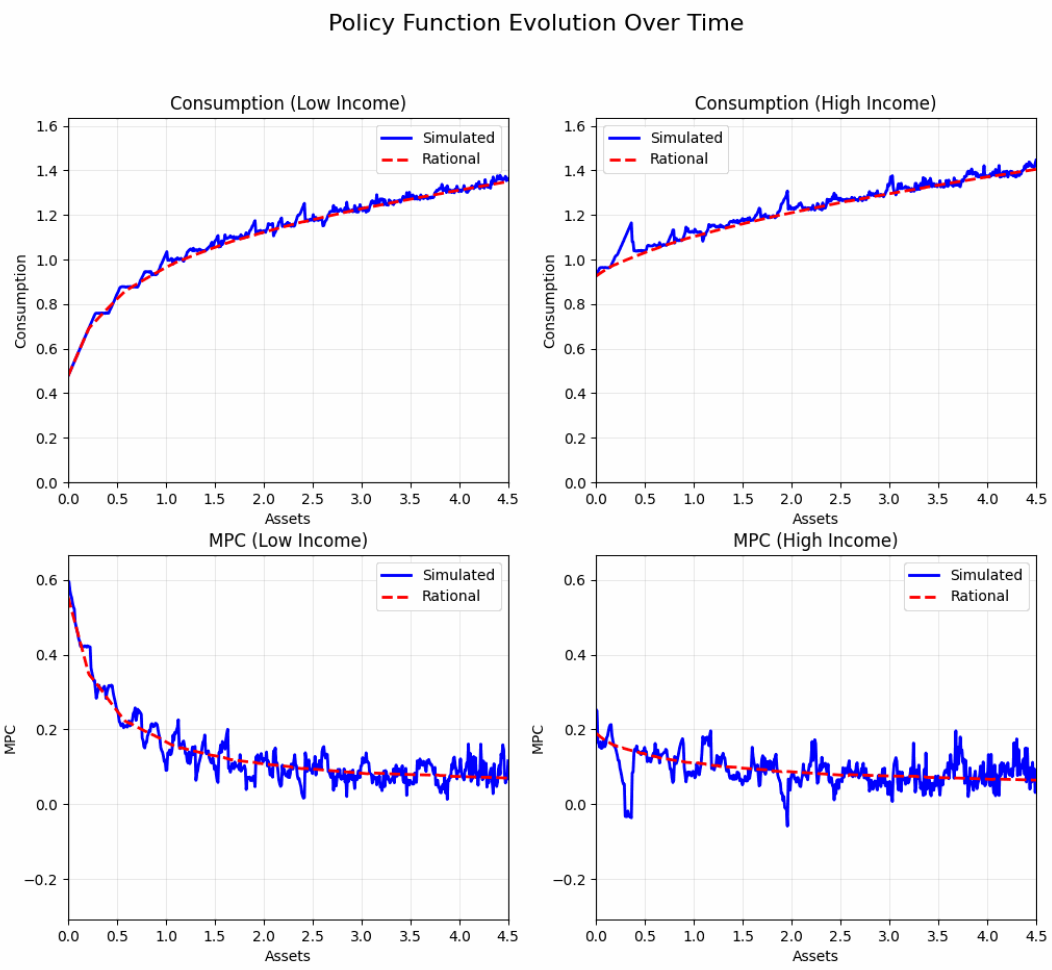


Figure A6: Initial consumption policy fit at $t = 0$ quarters for a single seed, no smoothing.

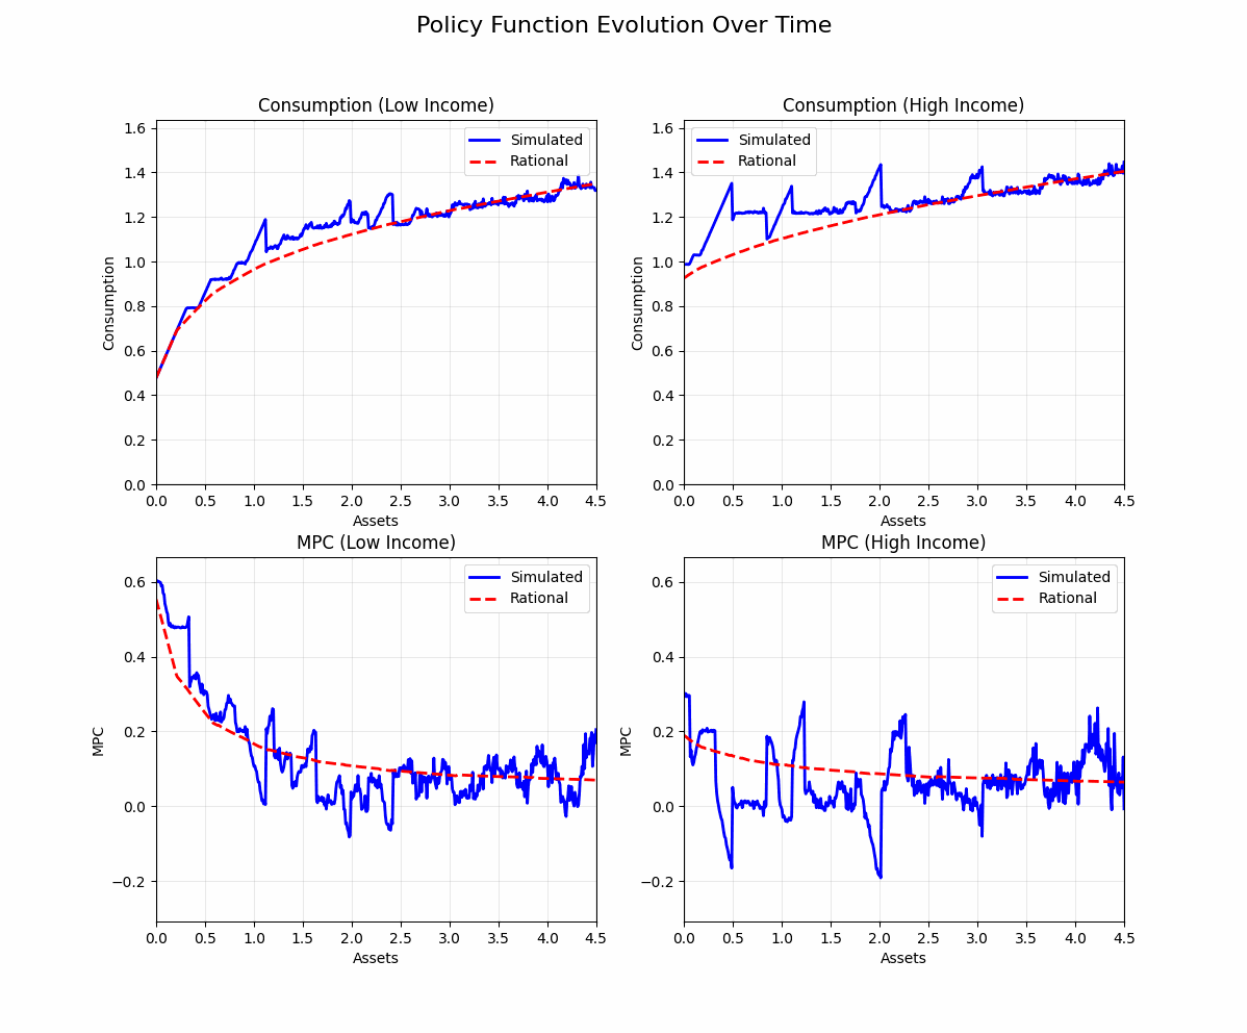


Figure A7: Consumption policy and MPC fit against rational at $t = 10$ quarters for a single seed, no smoothing, systematic tilting of policy function away from rational, as well as local fluctuations without smoothing.

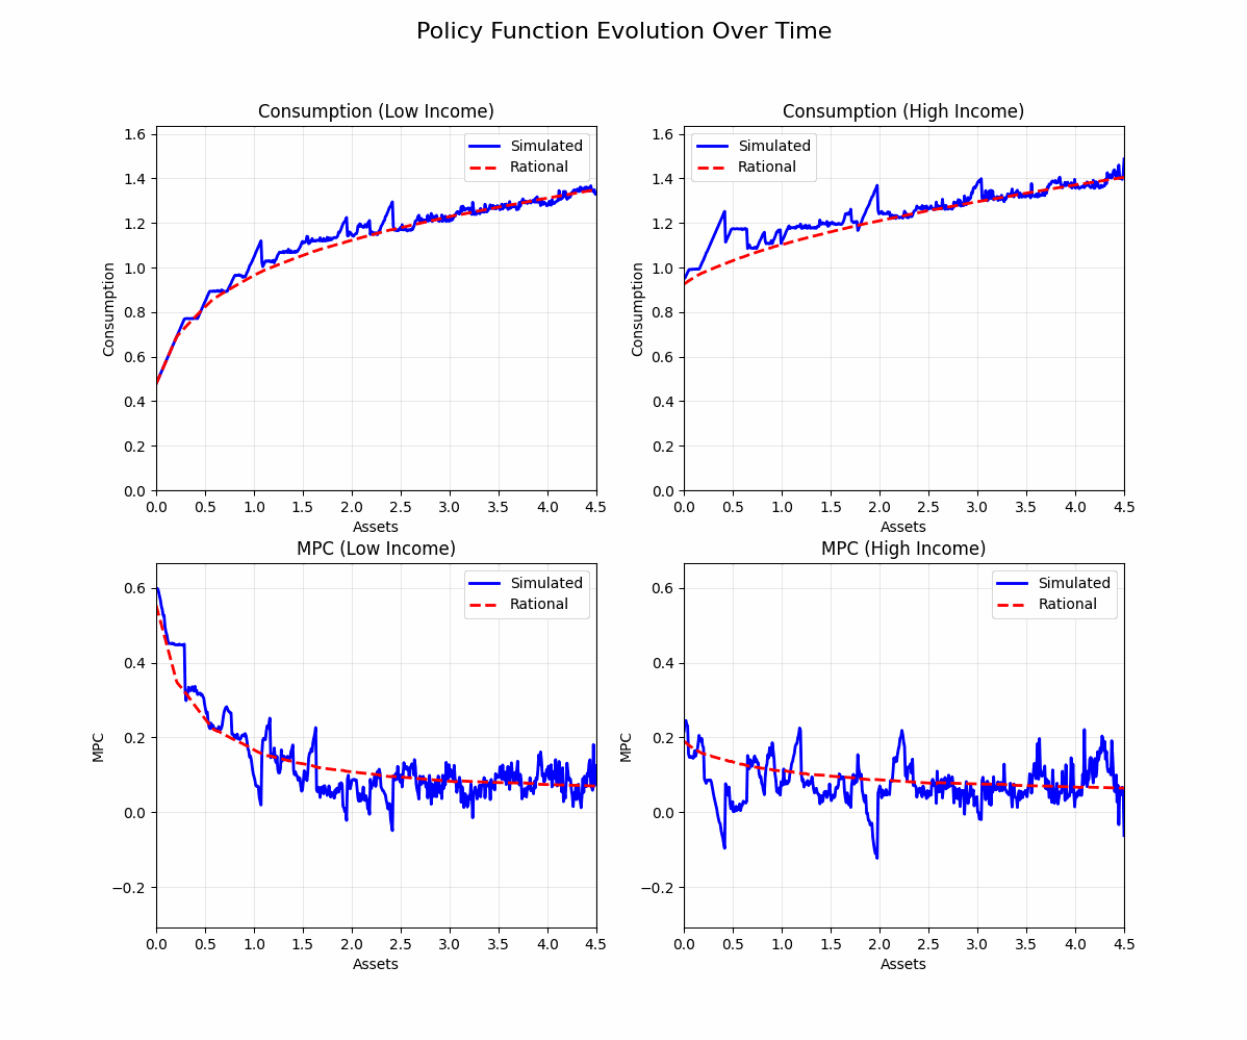


Figure A8: Consumption policy and MPC fit against rational at $t = 240$ quarters for a single seed, no smoothing. Local fluctuations remain elevated but systematic bias in policy has mostly vanished.

References

Barberis, Nicholas C., and Lawrence J. Jin. 2023. "Model-Free and Model-Based Learning as Joint Drivers of Investor Behavior." Pre-published, March. Working Paper. Accessed July 25, 2025. https://doi.org/10.3386/w31081. National Bureau of Economic Research: 31081. https://www.nber.org/papers/w31081.

Botvinick, Matthew, Jane X. Wang, Will Dabney, Kevin J. Miller, and Zeb Kurth-Nelson. 2020. "Deep Reinforcement Learning and Its Neuroscientific Implications." *Neuron* 107, no. 4 (August 19, 2020): 603–616. ISSN: 0896-6273, accessed August 5, 2025. https://doi.org/10.1016/j.neuron.2020.06.014. https://www.sciencedirect.com/science/article/pii/S0896627320304682.

Cogley, Timothy, and Thomas J. Sargent. 2008. "Anticipated Utility and Rational Expectations as Approximations of Bayesian Decision Making." *International Economic Review* 49 (1): 185–221. ISSN: 1468-2354, accessed July 23, 2025. https://doi.org/10.1111/j.1468-2354.2008.00477.x. https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1468-2354.2008.00477.x.

Dabney, Will, Zeb Kurth-Nelson, Naoshige Uchida, Clara Kwon Starkweather, Demis Hassabis, Rémi Munos, and Matthew Botvinick. 2020. "A Distributional Code for Value in Dopamine-Based Reinforcement Learning." *Nature* 577, no. 7792 (January): 671–675. ISSN: 1476-4687, accessed August 5, 2025. https://doi.org/10.1038/s41586-019-1924-6. https://www.nature.com/articles/s41586-019-1924-6.

Daw, Nathaniel D., Samuel J. Gershman, Ben Seymour, Peter Dayan, and Raymond J. Dolan. 2011. "Model-Based Influences on Humans' Choices and Striatal Prediction Errors." *Neuron* 69, no. 6 (March 24, 2011): 1204–1215. ISSN: 0896-6273, accessed August 5, 2025. https://doi.org/10.1016/j.neuron.2011.02.027. PMID: 21435563. https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3077926/.

Du, Hongyang, Ruichen Zhang, Yinqiu Liu, Jiacheng Wang, Yijing Lin, Zonghang Li, Dusit Niyato, et al. 2024. "Enhancing Deep Reinforcement Learning: A Tutorial on Generative Diffusion Models in Network Optimization." *IEEE Communications Surveys and Tutorials* 26 (4): 2611–2646. ISSN: 1553-877X, 2373-745X, accessed July 27, 2025. https://doi.org/10.1109/COMST.2024.3400011. https://ieeexplore.ieee.org/document/10529221/.

Evans, George W., Seppo Honkapohja, and Seppo Honkapohja. 1994. "Learning, Convergence, and Stability with Multiple Rational Expectations Equilibria." *European Economic Review* 38, no. 5 (May 1, 1994): 1071–1098. ISSN: 0014-2921, accessed July 23, 2025. https://doi.org/10.1016/0014-2921(94)90038-8. https://www.sciencedirect.com/science/article/pii/0014292194900388.

Friedman, Milton. 1957. "The Permanent Income Hypothesis." In *A Theory of the Consumption Function,* 20–37. Princeton University Press. Accessed July 23, 2025. https://www.nber.org/books-and-chapters/theory-consumption-function/permanent-income-hypothesis.

Ganong, Peter, Fiona Greig, Pascal Noel, Daniel M. Sullivan, and Joseph Vavra. 2024. "Spending and Job-Finding Impacts of Expanded Unemployment Benefits: Evidence from Administrative Micro Data." Pre-published, *American Economic Review* (Cambridge, MA) 114, no. w30315 (9 2024): 2898–2939. ISSN: 0002-8282, 1944-7981, accessed July 22, 2025. https://doi.org/10.1257/aer.20220973. National Bureau of Economic Research: w30315. https://pubs.aeaweb.org/doi/10.1257/aer.20220973.

Guvenen, Fatih. 2007. "Learning Your Earning: Are Labor Income Shocks Really Very Persistent?" *American Economic Review* 97, no. 3 (June): 687–712. ISSN: 0002-8282, accessed July 25, 2025. https://doi.org/10.1257/aer.97.3.687. https://www.aeaweb.org/articles?id=10.1257/aer.97.3.687.

Hafner, Danijar, Jurgis Pasukonis, Jimmy Ba, and Timothy Lillicrap. 2025. "Mastering Diverse Control Tasks through World Models." *Nature* 640, no. 8059 (April): 647–653. ISSN: 1476-4687, accessed July 27, 2025. https://doi.org/10.1038/s41586-025-08744-2. https://www.nature.com/articles/s41586-025-08744-2.

Havranek, Tomas, and Anna Sokolova. 2020. "Do Consumers Really Follow a Rule of Thumb? Three Thousand Estimates from 144 Studies Say "Probably Not"." *Review of Economic Dynamics* 35 (January 1, 2020): 97–122. ISSN: 1094-2025, accessed July 25, 2025. https://doi.org/10.1016/j.red.2019.05.004. https://www.sciencedirect.com/science/article/pii/S1094202518302849.

Ilut, Cosmin, and Rosen Valchev. 2024. "Learning Optimal Behavior through Reasoning and Experiences." Pre-published, March 27, 2024. Accessed July 25, 2025. https://doi.org/10.48550/arXiv.2403.18185. arXiv: 2403.18185 [econ]. http://arxiv.org/abs/2403.18185.

Jappelli, Tullio, and Luigi Pistaferri. 2010. "The Consumption Response to Income Changes." Pre-published, April 1, 2010. SSRN Scholarly Paper. Accessed July 25, 2025. https://doi.org/10.1146/annurev.economics.050708.142933. Social Science Research Network: 1693010. https://papers.ssrn.com/abstract=1693010.

Jovanovic, Boyan. 1979. "Job Matching and the Theory of Turnover." *Journal of Political Economy* 87 (5, Part 1): 972–990. ISSN: 0022-3808, accessed July 23, 2025. https://doi.org/10.1086/260808. https://www.journals.uchicago.edu/doi/abs/10.1086/260808.

Kaelbling, Leslie Pack, Michael L. Littman, and Anthony R. Cassandra. 1998. "Planning and Acting in Partially Observable Stochastic Domains." *Artif. Intell.* 101, no. 1 (May 1, 1998): 99–134. ISSN: 0004-3702, accessed July 22, 2025. https://doi.org/10.1016/S0004-3702(98)00023-X. https://www.sciencedirect.com/science/article/pii/S000437029800023X.

Kaplan, Greg, and Giovanni L. Violante. 2022. "The Marginal Propensity to Consume in Heterogeneous Agent Models." Pre-published, *Annual Review of Economics,* Working Paper Series, 14, no. 30013 (Volume 14, 2022 2022): 747–775. ISSN: 1941-1383, 1941-1391, accessed July 24, 2025. https://doi.org/10.1146/annurev-economics-080217-053444. https://www.annualreviews.org/content/journals/10.1146/annurev-economics-080217-053444.

Kasdin, Jonathan, Alison Duffy, Nathan Nadler, Arnav Raha, Adrienne L. Fairhall, Kimberly L. Stachenfeld, and Vikram Gadagkar. 2025. "Natural Behaviour Is Learned through Dopamine-Mediated Reinforcement." *Nature* 641, no. 8063 (May): 699–706. ISSN: 1476-4687, accessed August 5, 2025. https://doi.org/10.1038/s41586-025-08729-1. https://www.nature.com/articles/s41586-025-08729-1.

Kozlowski, Julian, Laura Veldkamp, and Venky Venkateswaran. 2020. "The Tail That Wags the Economy: Beliefs and Persistent Stagnation." *Journal of Political Economy* 128, no. 8 (August): 2839–2879. ISSN: 0022-3808, 1537-534X, accessed July 25, 2025. https://doi.org/10.1086/707735. https://www.journals.uchicago.edu/doi/10.1086/707735.

Krueger, D., K. Mitman, and F. Perri. 2016. "Chapter 11 - Macroeconomics and Household Heterogeneity." In *Handbook of Macroeconomics,* edited by John B. Taylor and Harald Uhlig, 2:843–921. Elsevier, January 1, 2016. Accessed July 28, 2025. https://doi.org/10.1016/bs.hesmac.2016.04.003. https://www.sciencedirect.com/science/article/pii/S1574004816300039.

Lamb, Alex, Riashat Islam, Yonathan Efroni, Aniket Rajiv Didolkar, Dipendra Misra, Dylan J. Foster, Lekan P. Molu, et al. 2023. "Guaranteed Discovery of Control-Endogenous Latent States with Multi-Step Inverse Models." *Trans. Mach. Learn. Res.* 2023. ISSN: 2835-8856, accessed July 27, 2025. https://openreview.net/forum?id=TNocbXm5MZ.

Leong, Yuan Chang, Angela Radulescu, Reka Daniel, Vivian DeWoskin, and Yael Niv. 2017. "Dynamic Interaction between Reinforcement Learning and Attention in Multidimensional Environments." *Neuron* 93, no. 2 (January 18, 2017): 451–463. ISSN: 0896-6273, accessed July 27, 2025. https://doi.org/10.1016/j.neuron.2016.12.040. PMID: 28103483. https://www.cell.com/neuron/abstract/S0896-6273(16)31039-X.

Levine, Alexander, Peter Stone, and Amy Zhang. 2024. "Learning a Fast Mixing Exogenous Block MDP Using a Single Trajectory." The Thirteenth International Conference on Learning Representations. October 4, 2024. Accessed July 27, 2025. https://openreview.net/forum?id=41WIgfdd5o.

Li, Yuanzhi, and Yang Yuan. 2017. "Convergence Analysis of Two-Layer Neural Networks with ReLU Activation." In *Advances in Neural Information Processing Systems,* vol. 30. Curran Associates, Inc. Accessed July 24, 2025. https://proceedings.neurips.cc/paper_files/paper/2017/hash/a96b65a721e561e1e3de768ac819ffbb-Abstract.html.

Malmendier, Ulrike, and Leslie Sheng Shen. 2024. "Scarred Consumption." *American Economic Journal: Macroeconomics* 16, no. 1 (January 1, 2024): 322–355. ISSN: 1945-7707, 1945-7715, accessed July 22, 2025. https://doi.org/10.1257/mac.20210387. https://pubs.aeaweb.org/doi/10.1257/mac.20210387.

Marcet, Albert, and Thomas J Sargent. 1989. "Convergence of Least Squares Learning Mechanisms in Self-Referential Linear Stochastic Models." *Journal of Economic Theory* 48, no. 2 (August 1, 1989): 337–368. ISSN: 0022-0531, accessed July 23, 2025. https://doi.org/10.1016/0022-0531(89)90032-X. https://www.sciencedirect.com/science/article/pii/002205318990032X.

Masset, Paul, Pablo Tano, HyungGoo R. Kim, Athar N. Malik, Alexandre Pouget, and Naoshige Uchida. 2025. "Multi-Timescale Reinforcement Learning in the Brain." *Nature* 642, no. 8068 (June): 682–690. ISSN: 1476-4687, accessed August 5, 2025. https://doi.org/10.1038/s41586-025-08929-9. https://www.nature.com/articles/s41586-025-08929-9.

Muller, Timothy H., James L. Butler, Sebastijan Veselic, Bruno Miranda, Joni D. Wallis, Peter Dayan, Timothy E. J. Behrens, Zeb Kurth-Nelson, and Steven W. Kennerley. 2024. "Distributional Reinforcement Learning in Prefrontal Cortex." *Nature Neuroscience* 27, no. 3 (March): 403–408. ISSN: 1546-1726, accessed August 5, 2025. https://doi.org/10.1038/s41593-023-01535-w. https://www.nature.com/articles/s41593-023-01535-w.

Muth, John F. 1960. "Optimal Properties of Exponentially Weighted Forecasts." *Journal of the American Statistical Association* 55 (290): 299–306. ISSN: 0162-1459, accessed July 23, 2025. https://doi.org/10.2307/2281742. JSTOR: 2281742. https://www.jstor.org/stable/2281742.

Qin, Aoyang, Feng Gao, Qing Li, Song-Chun Zhu, and Sirui Xie. 2023. "Learning Non-Markovian Decision-Making from State-Only Sequences." *Advances in Neural Information Processing Systems* 36 (December 15, 2023): 6596–6618. Accessed July 27, 2025. https://proceedings.neurips.cc/paper_files/paper/2023/hash/154926e0b66e2b2a8c1120852f31a12d-Abstract-Conference.html.

Rasamoelina, Andrinandrasana David, Fouzia Adjailia, and Peter Sinák. 2020. "A Review of Activation Function for Artificial Neural Network." In *2020 IEEE 18th World Symposium on Applied Machine Intelligence and Informatics (SAMI)*, 281–286. 2020 IEEE 18th World Symposium on Applied Machine Intelligence and Informatics (SAMI). January. Accessed July 24, 2025. https://doi.org/10.1109/SAMI48414.2020.9108717. https://ieeexplore.ieee.org/abstract/document/9108717.

Robbins, Herbert, and Sutton Monro. 1951. "A Stochastic Approximation Method." *Annals of Mathematical Statistics* 22, no. 3 (September): 400–407. ISSN: 0003-4851, 2168-8990, accessed July 23, 2025. https://doi.org/10.1214/aoms/1177729586. https://projecteuclid.org/journals/annals-of-mathematical-statistics/volume-22/issue-3/A-Stochastic-Approximation-Method/10.1214/aoms/1177729586.full.

Sousa, Margarida, Pawel Bujalski, Bruno F. Cruz, Kenway Louie, Daniel C. McNamee, and Joseph J. Paton. 2025. "A Multidimensional Distributional Map of Future Reward in Dopamine Neurons." *Nature* 642, no. 8068 (June): 691–699. ISSN: 1476-4687, accessed August 5, 2025. https://doi.org/10.1038/s41586-025-09089-6. https://www.nature.com/articles/s41586-025-09089-6.

Sutton, Richard S., and Andrew G. Barto. 2018. *Reinforcement Learning: An Introduction.* Second edition. Adaptive Computation and Machine Learning Series. Cambridge, Massachusetts: The MIT Press. ISBN: 978-0-262-03924-6. http://incompleteideas.net/book/the-book-2nd.html.

Sutton, Richard S., Doina Precup, and Satinder Singh. 1999. "Between MDPs and Semi-MDPs: A Framework for Temporal Abstraction in Reinforcement Learning." *Artificial Intelligence* 112, no. 1 (August 1, 1999): 181–211. ISSN: 0004-3702, accessed July 28, 2025. https://doi.org/10.1016/S0004-3702(99)00052-1. https://www.sciencedirect.com/science/article/pii/S0004370299000521.

Tennenholtz, Guy, Nadav Merlis, Lior Shani, Martin Mladenov, and Craig Boutilier. 2023. "Reinforcement Learning with History Dependent Dynamic Contexts." In *Proceedings of the 40th International Conference on Machine Learning,* 34011–34053. International Conference on Machine Learning. PMLR, July 3, 2023. Accessed July 27, 2025. https://proceedings.mlr.press/v202/tennenholtz23a.html.

Whitehead, Steven D., and Long-Ji Lin. 1995. "Reinforcement Learning of Non-Markov Decision Processes." *Artificial Intelligence,* Computational Research on Interaction and Agency, Part 2, 73, no. 1 (February 1, 1995): 271–306. ISSN: 0004-3702, accessed July 27, 2025. https://doi.org/10.1016/0004-3702(94)00012-P. https://www.sciencedirect.com/science/article/pii/000437029400012P.

Xu, Pan, and Quanquan Gu. 2020. "A Finite-Time Analysis of Q-learning with Neural Network Function Approximation." In *Proceedings of the 37th International Conference on Machine Learning,* 10555–10565. International Conference on Machine Learning. PMLR, November 21, 2020. Accessed July 24, 2025. https://proceedings.mlr.press/v119/xu20c.html.