

A note on observational equivalence of micro assumptions on macro level

Alexey A. Ponomarenko

Abstract

The author sets up a simplistic agent-based model where agents learn with reinforcement observing an incomplete set of variables. The model is employed to generate an artificial dataset that is used to estimate standard macro econometric models. The author shows that the results are qualitatively indistinguishable (in terms of the signs and significances of the coefficients and impulse-responses) from the results obtained with a dataset that emerges in a genuinely rational system.

JEL B41 C63 D83

Keywords Microfoundations; bounded rationality; reinforcement learning; agent-based model

Authors

Alexey A. Ponomarenko, Bank of Russia, Moscow, PonomarenkoAA@cbr.ru

Citation Alexey A. Ponomarenko (2019). A note on observational equivalence of micro assumptions on macro level. Economics Discussion Papers, No 2019-51, Kiel Institute for the World Economy. <http://www.economics-journal.org/economics/discussionpapers/2019-51>

1. Introduction

The development of microfounded models has become a common topic in the economic literature. Notably, these models usually only serve as theoretical foundations and play little role in empirical analysis or forecasting. As pointed out by Wren-Lewis (2018), the researchers typically choose what developments they want to explain and develop the model accordingly. The task is claimed to be successfully accomplished if a researcher is able to produce a model that replicates some empirical stylized facts. These empirical descriptors are usually presented in terms of simple measures (i.e. regression coefficients or impulse–response functions from a vector autoregression model) estimated using aggregate data.

Following the famous papers by Lucas (1972) and Sargent (1973), the rational expectations approach became the central methodology for providing the microfoundations of economic models. The reasons for such dominance are not entirely clear. Although the rationality of expectations is a convenient assumption, whether it actually holds in practice is an open question. As discussed by Caverzasi and Russo (2018) and Haldane and Turrell (2018), faced with uncertainty, it is often rational for agents to rely on simpler decision rules. These are often called “rules of thumb” or heuristics. Use of such rules is sometimes thought to be arbitrary, sub-optimal or irrational. Yet in a world of uncertainty, that is far from clear. Heuristics may be the most robust means of making decisions in a world of uncertainty.

The ability of rational expectations models to reproduce some of the stylized facts observed in empirical aggregate data is far from unique. It is well known that seemingly simple systems may produce intricate and often efficient developments of the modelled variables. See for example seminal work by Arthur (1994), as well as Gode and Sunder (1993). As regards examples with a closer relation to macroeconomics, Andolfatto et al. (2008) and Ilek (2017) use Monte Carlo experiments to show that a model in which the expectations are not rational may generate artificial datasets where rationality would not be rejected by the textbook tests.

We contribute to this discussion by providing an example of how the aggregate indicators from a world populated by bounded-rational agents may be indistinguishable from rational developments based on the conventional econometric models. Unlike Andolfatto et al. (2008) and Ilek (2017), we use a radically different set-up to model bounded-rational agents. For this purpose, in the spirit of Schuster (2012), we set up an agent-based model populated by agents who learn with reinforcement. We generate several sets of artificial observations using

alternative specifications of the learning algorithm and investigate the properties of the standard econometric models estimated using these datasets.

The rest of this paper is structured as follows. Section 2 outlines the modelling set up. Section 3 presents the results of the experiments. Section 4 concludes.

2. The model

Our example is the market of a good where in each period the producers need to decide on whether to participate based on their costs and expected price.

There are n agents. Each of them may produce q_n goods.

The agents incur costs ($c_{n,t}$):

$c_{n,t} = \lambda_n C_t + \varepsilon_{n,t}$ where C_t is the trend component and $\varepsilon_{n,t}$ are random agent-specific innovations. The trend follows the exogenous autoregressive process

$$C_t = \alpha_0 + \alpha_1 C_{t-1} + v_t.$$

The price P_t of the good is determined by

$$P_t = D_t / Q_t$$

where Q_t is the sum of the q_n of the agents that decided to participate in the market. Here, D_t is the demand for the good. It follows an exogenous autoregressive process:

$$D_t = \beta_0 + \beta_1 D_{t-1} + \zeta_t$$

An agent's profit ($w_{n,t}$) is determined by

$$w_{n,t} = q_n (P_t - c_{n,t})$$

Note that $q_n = 0$ when an agent decides not to participate in the market.

The values of the parameters are presented in the Annex.

In the next subsections we describe the alternative algorithms the agents use to decide on their market participation.

2.1. Rational agents

The first type of agents know all the data generating processes, distributions of parameters and past values of global variables. They also assume that so do all other agents. We label this type of agents as rational.

In the beginning of the period all agents get to know the realisation of their costs ($c_{n,t}$).

They are also able to calculate the expected values of demand (D_t^e) and trend costs (C_t^e) using the known values of D_{t-1} , C_{t-1} , α_0 , α_1 , β_0 and β_1 .

Next the agents calculate the expected output (Q_t^e) and price (P_t^e).

The stylised demand curve may be expressed as

$$Q_t^e = D_t^e / P_t^e . \quad (1)$$

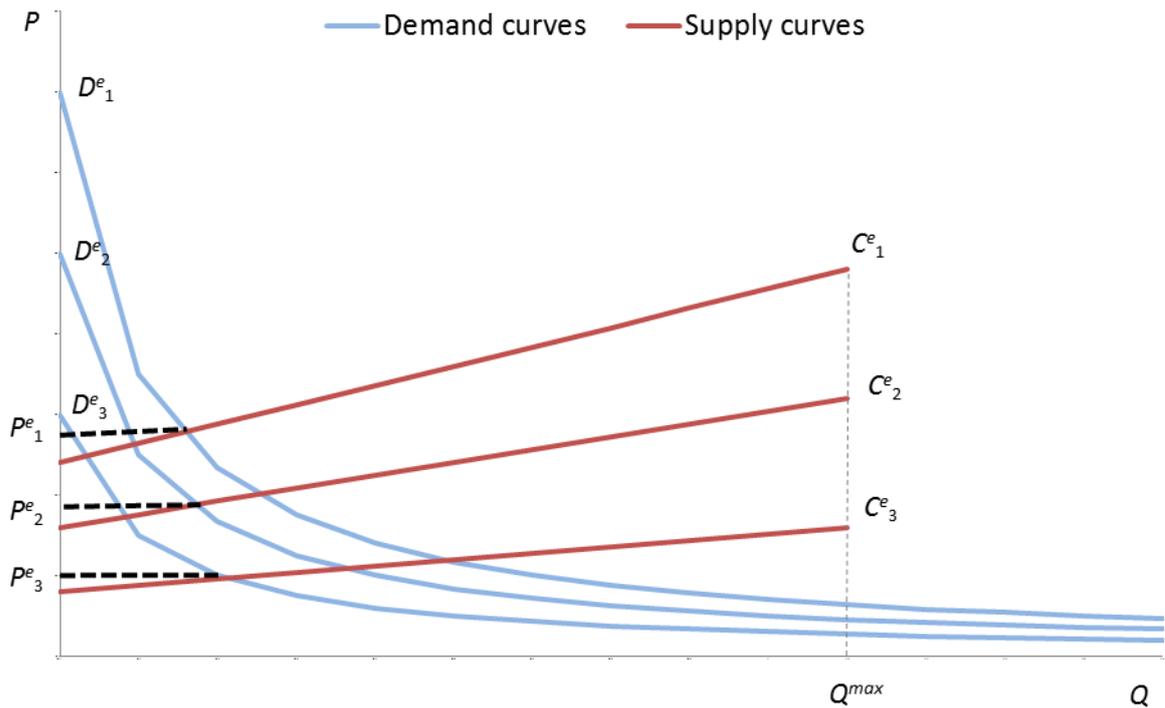
The stylised supply curve may be expressed as

$$P_t^e = C_t^e \lambda^{min} + C_t^e \frac{\lambda^{max} - \lambda^{min}}{\lambda^{min} Q^{max}} Q_t^e , \quad (2)$$

where λ^{min} and λ^{max} are the parameters of the distributions $\lambda_n \in (\lambda^{min}, \lambda^{max})$, and Q^{max} is the sum of q_n over all agents.

The system of (1) and (2) may be solved for given D_t^e and C_t^e (see Figure 1 for a visualization) and P_t^e is calculated. Agents with $c_{n,t} < P_t^e$ will participate in the market.

Figure 1. Stylised demand and supply curves



2.2. Learning agents

An alternative paradigm is the learning procedure where the agents do not know the underlying data generating processes. Instead, the strategies that lead to losses tend to be abandoned, while strategies that lead to profit tend to be preferred.

In this paper we employ the reinforcement learning approach outlined in Sutton and Barto (1998) and implemented in an agent-based framework by Schuster (2012). For illustrative purposes, we have intentionally selected a concept that is simplistic and conceptually very different from the standard rationality assumption. This algorithm proposes a simple generic decision model for boundedly rational, adaptive artificial agents. It assumes that the agents start with very limited information about the world, and possess no causal model of how their actions affect themselves or their environment. The formal description of the algorithm follows.

The agents perceive the environment as being described by a collection of k attributes $\{\text{att}_1 \dots \text{att}_k\}$. Each attribute is represented by an observed variable and can take seven discrete values from extremely low to extremely high.¹ Accordingly, each situation may be classified by the agents as being in one of $s = 7^k$ possible states.

In this paper we employ two types of learning agents.

¹ The numeric ranges are presented in the Annex. The ranges were calibrated to ensure that observations were roughly equally distributed across possible states.

The first type of agents use three ($k = 3$) variables as state descriptors: current agent-specific costs in relation to the past price ($c_{n,t} - P_{t-1}$), past trend costs (C_{t-1}), and demand (D_{t-1}) indicators. We label this type of agent as learning agents with full information.

The second type of agents only use their current agent-specific costs in relation to the past price ($c_{n,t} - P_{t-1}$) as the state descriptor. We label this type of agent as learning agents with limited information.

After the state is classified the agents choose whether to participate in the market. Each agent will produce q_n goods with probability

$$\text{pr}_{n,t} = \frac{1}{1 + e^{-\theta f_{s,n,t}}},$$

where $f_{s,n,t}$ is the attractiveness (fitness) of participation in the currently observed state. This parameter is initially set to zero and in the subsequent periods, after the price, output and profits ($w_{n,t}$) are determined, the agents that participated in the market update this value:

$$f_{s,n,t} = \mu f_{s,n,t-1} + (1-\mu) w_{n,t}.$$

The parameters of the learning algorithm are reported in the Annex.

3. The experiments

We generate 5 artificial datasets using the models populated by the following types of agents:

- **Rational agents.**
- **Learning agents with full information.**
- **Learning agents with limited information.**
- **Mixed strategies.** The model is simultaneously populated (in equal proportions) by the three types of agents mentioned above.
- **Random strategies.** The agents choose to participate in the market with the invariant probability of 0.5. We report the results for this system for illustrative purposes to demonstrate the role of the decision making algorithms.

We conduct 500 independent model runs each producing 5000 observations. We discard the first 2000 observations and only use the remaining 3000 after the learning agents' systems have already arrived at the steady state.

The descriptive statistics for the obtained artificial series of aggregate output, prices and average profit per agent are reported in Table 1. The results show that the market participation

rate in the learning agents' system is higher than in the rational world. This is reflected in higher output, lower prices and lower profitability (although rational agents do not outperform the learning agents when they co-exist in the system with mixed strategies). The series are also somewhat more volatile in the learning agents' systems.

Table 1. Descriptive statistics of artificial datasets: mean (standard deviation)

Variable	Agents' type				
	Rational	Learning (full inf.)	Learning (limited inf.)	Mixed strategies	Random strategies
<i>Output</i>	619 (43)	854 (66)	846 (51)	793 (54)	1000 (37)
<i>Price</i>	81 (6)	59 (4)	59 (3)	63 (4)	50 (2)
<i>Profit per agent</i>	Aggregate: 4.8 (0.6)				
	7.1 (0.1)	3.8 (2.2)	4.1 (1.6)	Rational: 4.8 (0.7)	-8.4 (3.2)
	Learning (both types): 4.9 (0.7)				

Obviously, these findings are model-dependent and as such serve little purpose other than to demonstrate there are noticeable differences between the dynamics emerging in the different systems. But are these differences sufficient to distinguish between rational and boundedly rational worlds by estimating the standard macro econometric models? We examine this issue in the next sub-sections by using the artificial datasets to estimate such models.

3.1. GMM regressions

We start by regressing the aggregate output variable (Q_t) on trend costs (C_t) and demand (D_t) variables.² The estimation is conducted via the Generalized Method of Moments (GMM) and the lags of the dependent and the explanatory variables are used as the instruments. The obtained coefficients are conventionally interpreted in the literature as the representation of agents' reactions to the fluctuations of the expected values of the variables in question (see, e.g. Gali and Gertler 1999).

The results for the datasets generated via alternative models are presented in Table 2. The coefficients in all systems have the expected sign.³ Interestingly, even the model estimated for the limited information agents indicate that output 'reacts' to fluctuations of demand, although we know that formally this is not the case. The models' fit (R^2) is not informative for distinguishing between rational and boundedly rational agents. Notably, the autocorrelation of

² For estimation of the models presented in Section 3.1 and 3.2 we pool the datasets across all 500 model runs and use the logs of the variables. There is a direct mechanical link between the two endogenous variables (output and price). There is no additional information contained in the models estimated for prices as the dependent variable. Therefore we only report the model estimates for aggregate output.

³ All reported coefficients are statistically significant at the 1% level.

the residuals is low in the system of the learning agents, although such a result is conventionally regarded as an indicator of rationality (Rich 1989).

Table 2. GMM regressions for output as the dependent variable

	Agents' type				
	Rational	Learning (full inf.)	Learning (limited inf.)	Mixed strategies	Random strategies
Costs	-0.71	-0.49	-0.18	-0.58	0.00
Demand	0.73	0.57	0.32	0.68	0.00
Intercept	1	2.2	3.9	1.4	6.8
R^2	0.18	0.14	0.17	0.79	0.00
Residuals' $AR(1)$	0.75	-0.02	0.13	0.18	0.00

3.2. Impulse–response analysis

We proceed by estimating the conventional vector autoregression (VAR) models⁴

$$Y_t = B(L)Y_{t-1} + u_t$$

$$Ae_t = u_t,$$

where Y_t is a vector of time series comprising output (Q_t), trend costs (C_t) and demand (D_t) variables; $B(L)$ is a matrix polynomial in the lag operator L ; u_t is a vector of residuals; and e_t is a vector of independent structural innovations. The identification scheme (matrix A) of independent innovations is structured as follows: C_t and D_t cannot be affected by any contemporaneous innovations, but the residuals of C_t and D_t may contemporaneously affect Q_t .

The impulse–response functions estimated for alternative datasets are reported in Figures 2 and 3. The results show that in all cases output is ‘affected’ by the innovations in the exogenous variables. Although the magnitude of the responses is somewhat different across the datasets, the general pattern is very similar.⁵

In summary, our experiments show that the appropriate (i.e. corresponding to rational behaviour) correlations between the endogenous and exogenous variables are very likely to emerge even when the developments of exogenous variables are not known to the agents. Even though the agents do not know the underlying data generating process, they may efficiently adapt through reinforcement learning. Also note that even when the agents do not directly observe the developments of the exogenous variables, the information about these developments is contained

⁴ Note that VAR models are convenient data descriptors and may be used to compare the output of various classes of models (see e.g. Minford et al. 2016).

⁵ We do not report the confidence bands of the responses as they are insubstantial in all cases.

in the observed endogenous variable (price). This information proves to be sufficient for the emergence of the corresponding correlations.

Figure 2. Responses of output to innovation in costs

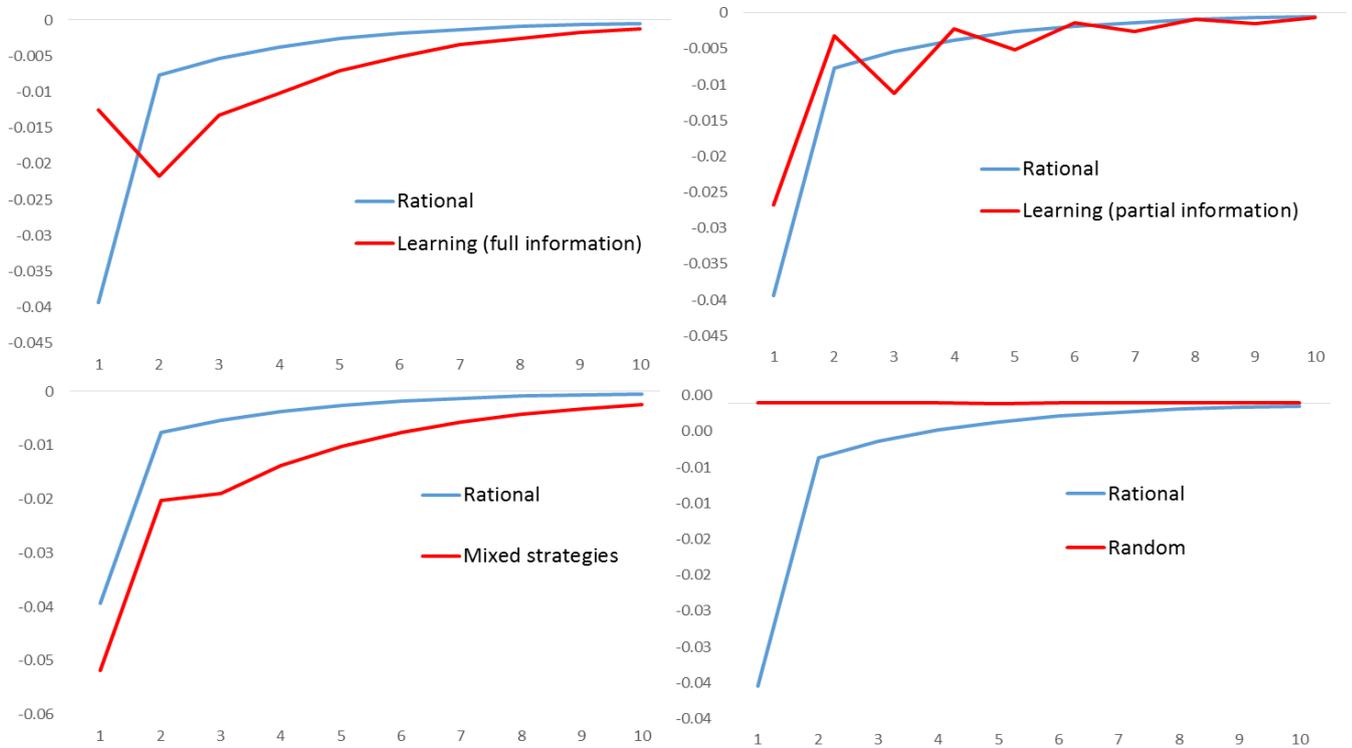
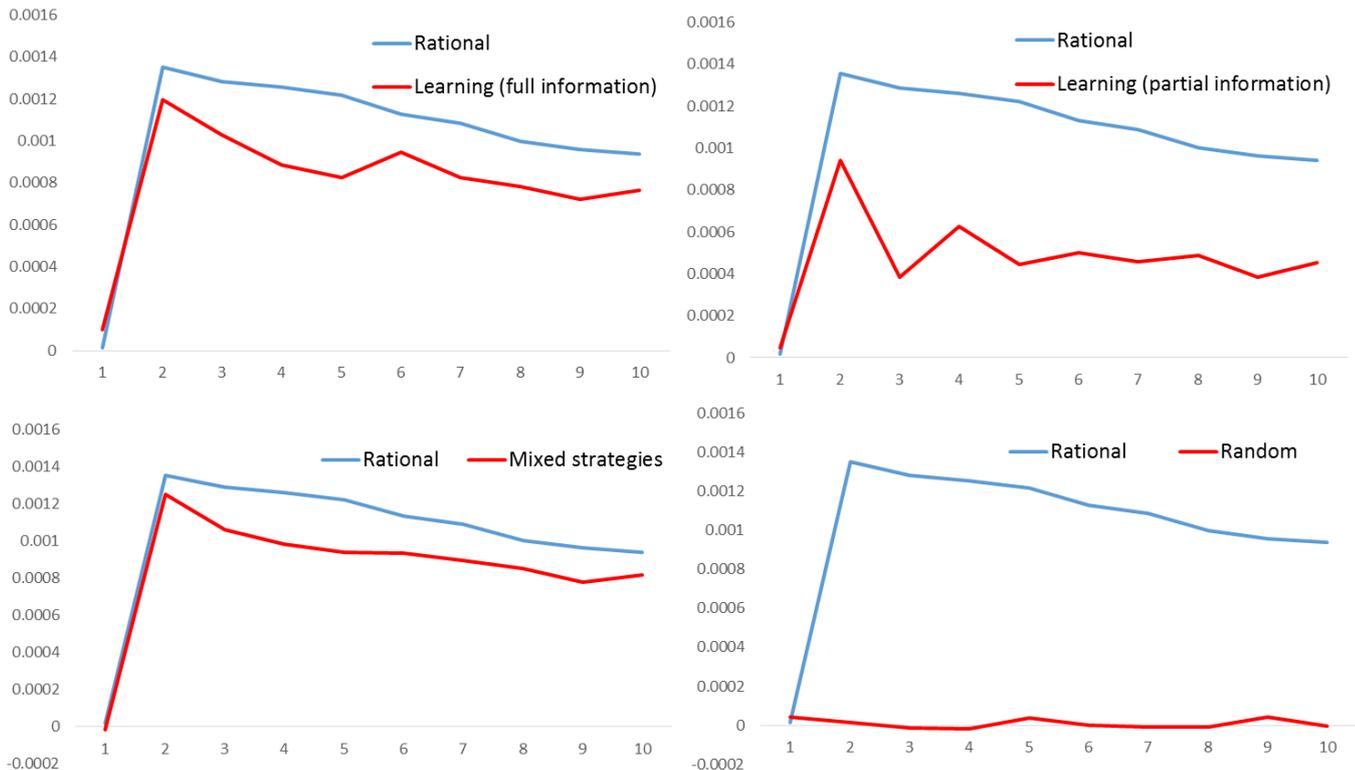


Figure 3. Responses of output to innovation in demand



4. Conclusions

Developing microfounded models that are based on the rational expectations hypothesis and presenting these models as a theoretical foundation for empirically estimated macroeconomic models has become a common practice in the economic literature. Yet, the fact that the empirically established sets of correlations between macro variables are in line with those derived theoretically neither is validated nor validates the exact micro assumptions employed in the theoretical model.

We provided an example to illustrate this point. We demonstrated that a simplistic learning algorithm employing a minimal set of observed indicators is sufficient to produce a set of correlations between macro variables that is indistinguishable (via standard macro econometric models) from the set of correlations that emerges in a genuinely rational system.

Acknowledgements

The views expressed in this paper are those of the author and do not necessarily represent the position of the Bank of Russia.

References

- Andolfatto, D., S. Hendry, and K. Moran (2008) [Are inflation expectations rational?](#) *Journal of Monetary Economics* 55, 406–422.
- Arthur, B. W. (1994) [Complexity in economic theory. Inductive reasoning and bounded rationality.](#) *American Economic Review* 82(2), 406-411.
- Caverzasi, E. and A. Russo (2018) [Toward a new microfounded macroeconomics in the wake of the crisis.](#) *Industrial and Corporate Change* 27(6), 999–1014.
- Gali, J. and M. Gertler (1999) [Inflation dynamics: A structural econometric analysis.](#) *Journal of Monetary Economics* 44(2), 195-222.
- Gode, D. K. and S. Sunder (1993) [Allocative Efficiency of Markets with Zero-Intelligence Traders: Market as a Partial Substitute for Individual Rationality.](#) *Journal of Political Economy* 101(1), 119-137.
- Haldane, A. G. and A. E. Turrell (2018) [An interdisciplinary model for macroeconomics.](#) *Oxford Review of Economic Policy* 34(1-2), 219–251.
- Ilek, A. (2017) [A note on the rationality test.](#) *Macroeconomic Dynamics* 21, 2121–2137.
- Lucas, R. E. (1972) [Expectations and neutrality of money.](#) *Journal of Economic Theory* 4, 103–124.
- Minford, P., M. Wickens and Y. Xu (2016) [Comparing different data descriptors in Indirect Inference tests on DSGE models.](#) *Economics Letters* 145, 157–161.
- Rich, R. W. (1989) [Testing the rationality of inflation forecasts from survey data.](#) *Review of Economics and Statistics* 71, 682–686.
- Sargent, T. J. (1973) [Rational expectations, the real rate of interest and the natural rate of unemployment.](#) *Brookings Papers on Economic Activity* 2, 429–472.
- Schuster, S. (2012) [BRA: An Algorithm for Simulating Bounded Rational Agents.](#) *Computational Economics* 39(1), 51-69.
- Sutton, R., and A. Barto (1998) [Reinforcement Learning: An Introduction.](#) MIT Press, Cambridge, MA
- Wren-Lewis, S. (2018) [Ending the microfoundations hegemony.](#) *Oxford Review of Economic Policy* 34 (1–2), 55–69.

Annex

Table 3. Parameters of the model

<u>Description</u>	<u>Parameter</u>	<u>Value</u>
Costs and demand determination		
Number of agents	n	1000
Agent-specific cost factor	λ_n	$\in (1,2)$
Agent-specific cost innovation	$\varepsilon_{n,t}$	$\sim N(0,2)$
Trend costs (intercept)	α_0	10
Trend costs (autoregression)	α_1	0.7
Trend costs (innovation)	ν_t	$\sim N(0,4)$
Demand (intercept)	β_0	2500
Demand (autoregression)	β_1	0.95
Demand (innovation)	ξ_t	$\sim N(0,100)$
Learning algorithm		
Sensitivity of probability of participating in the market to changes in the fitness function	θ	20
Time discount in the fitness function	μ	0.2

Table 4. State descriptors

Variable	Value						
	Extremely low	Very low	Low	Neutral	High	Very high	Extremely high
$c_{n,t} - P_{t-1}$	< -12	(-12 : -8)	(-8 : -2)	(-2 : 2)	(2 : 8)	(8 : 12)	> 12
C_{t-1}	< 30	(30 : 31)	(31 : 32)	(32 : 34)	(34 : 35)	(35 : 36)	> 36
$D_{t-1}/1000$	< 49.7	(49.7:49.8)	(49.8:49.9)	(49.9:50.1)	(50.1:50.2)	(50.2:50.3)	> 50.3

Please note:

You are most sincerely encouraged to participate in the open assessment of this discussion paper. You can do so by either recommending the paper or by posting your comments.

Please go to:

<http://www.economics-ejournal.org/economics/discussionpapers/2019-51>

The Editor