

Exploration of the Parameter Space in Macroeconomic Agent-Based Models

Karl Naumann-Woleske (0000-0001-8182-9256)

Chair of EconophysiX and Complex Systems, Ecole Polytechnique Paris
LadHyX, UMR CNRS 7646, Ecole Polytechnique Paris
karl.naumann@ladhyx.polytechnique.fr

Max Sina Knicker (0000-0003-0732-6612)

Chair of EconophysiX and Complex Systems, Ecole Polytechnique Paris
Physics of Complex Biosystems, Technical University of Munich

Michael Benzaquen (0000-0002-9751-7625)

Chair of EconophysiX and Complex Systems, Ecole Polytechnique Paris
LadHyX, UMR CNRS 7646, Ecole Polytechnique Paris
Capital Fund Management, 75007 Paris, France

Jean-Philippe Bouchaud (0000-0002-0427-6456)

Chair of EconophysiX and Complex Systems, Ecole Polytechnique Paris
Capital Fund Management, 75007 Paris, France
Académie des Sciences, Quai de Conti, 75006 Paris, France

Abstract

Agent-Based Models (ABM) are computational scenario-generators, which can be used to predict the possible future outcomes of the complex system they represent. To better understand the robustness of these predictions, it is necessary to understand the full scope of the possible phenomena the model can generate. Most often, due to high-dimensional parameter spaces, this is a computationally expensive task. Inspired by ideas coming from systems biology, we show that for multiple macroeconomic models, including an agent-based model and several Dynamic Stochastic General Equilibrium (DSGE) models, there are only a few *stiff* parameter combinations that have strong effects, while the other *sloppy* directions are irrelevant. This suggests an algorithm that efficiently explores the space of parameters by primarily moving along the stiff directions. We apply our algorithm to a medium-sized agent-based model, and show that it recovers all possible dynamics of the unemployment rate. The application of this method to Agent-based Models may lead to a more thorough and robust understanding of their features, and provide enhanced parameter sensitivity analyses. Several promising paths for future research are discussed.

1 Introduction

The global economy can be seen as a complex adaptive system, with large numbers of heterogeneous agents interacting in parallel (Tesfatsion 2002, LeBaron & Tesfatsion 2008). The evolution of this system over time gives rise to emergent phenomena such as business cycles, inequality or even crises and crashes. In recent years, agent-based models (ABMs) have been developed as a tool to generate and understand such emergent phenomena (for deeper introductions to agent-based modelling see Kirman & Gallegati (2022), and Arthur (2022) in this volume, as well as Epstein (1999), LeBaron & Tesfatsion (2008), Dosi & Roventini (2019), Haldane & Turrell (2019)). An agent-based model is a large computational device for simulating artificial economies of many heterogeneous and interacting agents that has been successfully used to replicate and explain a wide array of different observed macroeconomic behaviours (for example, see the works of Gatti et al. (2011), Dosi et al. (2017)). The richness of phenomena that can occur in agent-based models makes them excellent *scenario-generators*, as advocated in Gualdi et al. (2015), Bookstaber (2017), Bouchaud (2020). That is, they can be used to gain an understanding of the spectrum of possible future paths of the economy, given its current state, and how these future scenarios may respond to policy interventions. Unfortunately, the versatility of these models also means that they suffer from the “wilderness of bounded rationality” (Sims 1980) both due to the large number of parameters and the different choices of behavioural rules. At worst, they have been described as “black boxes” able to generate any phenomena (see Fagiolo & Roventini (2017) for a discussion). To understand how robust the generated phenomena are, and consequently use agent-based models to generate scenarios, it is necessary to understand the full spectrum of possible dynamics that can be generated.

Our research provides a first answer to the question of *what are the different types of dynamics that an agent-based model can generate?* In the present paper, we first establish which parameter combinations, in the wilderness of parameters, are most crucial to the observed dynamics. Using this, we develop an algorithm to systematically uncover the set of possible dynamics that an agent-based model can generate, and apply this to a medium-sized ABM.

We begin by studying a set of macroeconomic models including a medium-sized ABM and five Dynamic Stochastic General Equilibrium (DSGE) models, which represent economists’ current “workhorse” models for monetary policy (Kaplan & Violante 2018). Each model is a mapping \mathcal{M} from a parameter space Φ to an observable output space Y , $\mathcal{M} : \Theta \mapsto Y$. Moving around in the parameter space changes the observed output. In systems biology, it has been shown that there are generally only a few *stiff* directions in the parameter space that significantly change the observable dynamics (Quinn et al. 2021).¹ Conversely, there are many *sloppy* directions in which there are no significant changes in the model dynamics, such that the resulting combination of parameters is only very loosely constrained by empirical data.

All of the models we consider appear to be sloppy, with their observable dynamics changing significantly only along a handful of directions. The linear combinations making up these directions also typically involve many parameters, and are thus generally not aligned with the bare parameter axes.

In the case of DSGE models, the key directions are actually dominated by the parameters governing the shock processes, which is a known dependence and limitation of DSGE models (Stiglitz 2018). Our finding implies the one-parameter-at-a-time sensitivity analysis approaches common in the ABM literature does not give an accurate representation of the sensitivity to parameter combinations. While sophisticated methods of sensitivity analysis with respect to multiple parameters exist (see Ligmann-Zielinska et al. (2020) for a recent review), these have not yet permeated the economic ABM literature or are focused on specific policy variables. Our approach significantly complements such sensitivity

analyses by considering the robustness of the phenomena at a given point in parameter space to its most sensitive direction.

To further study the implications of the sloppy nature of these models, we focus on the medium-scale Mark-0 ABM of Gualdi et al. (2015, 2017), inspired by the original work of Gatti et al. (2011). In the case of Mark-0, the authors uncovered several phase transitions (or tipping points) between different dynamics in the model’s output.² Focusing on the dynamics of unemployment, we find that in the two-dimensional case the most stiff parameter combinations point towards the nearest phase transition. In the presence of multiple close transitions, they point towards these directions sequentially according to the distance of the transition.

Our method therefore provides a useful tool to navigate the space of parameters and find such abrupt changes of behaviour. Exploiting this idea, we build an algorithm to step along these stiff directions and thus systematically explore the phase space within a given computational budget. In the case of the Mark-0 model, we show that this approach can efficiently recover all possible dynamics for the unemployment rate for both the low-dimensional and high-dimensional cases.

The issue of searching the parameter space of Agent-based Models is one that has received considerable attention in recent years. A prominent method and the closest to ours is meta-modelling (see Lamperti et al. (2018), van der Hoog (2019), Zhang et al. (2020), ten Broeke et al. (2021)).³ In this approach, the observable outputs are sampled at different points in the parameter space. A simpler model is then constructed that mimics the relationship between the parameters and a given set of outputs, which are typically a set of moments, or to classify the outcomes based on some criteria (e.g. see the neural network, as in Lamperti et al. (2018)). The benefit of this approach is that once a surrogate model has been constructed, further exploring the parameter space through predicting the class of unseen parameters becomes computationally cheap. The crux is that the approach is exploratory only with respect to the modeller’s criteria, such as a fat tailed distribution of GDP growth. This requires some ex-ante knowledge of the possible outcomes (or at least the desired outcomes) of the model.

Our approach instead aims to be exploratory and agnostic, with no prior assumptions about the dynamics we are seeking: we aim to provide an understanding of the full set of possible phenomena that can be generated by the model through intelligent sampling of the parameter space. This is crucial when some emergent behaviour is truly unexpected, as is indeed often the case in physics for example, but also in macroeconomic ABMs, see the detailed discussion of this point in Gualdi et al. (2015). Our method complements the meta-modelling approach, as the discovery of the different phenomena can facilitate a more precise criteria for the classification of parameters by the meta-model. In this way, one could identify new phases using our algorithm, and determine the phase transitions using a meta-model.

To develop the insights mentioned above, we first develop the numerical method for estimating the Hessian matrix, from which the parameter directions are extracted (Section 2). We then apply this to the two ABMs and five DSGE models introduced in Section 3 to assess their sloppiness (Section 4). Finally, we use these insights to develop our algorithm (Section 5) and test it on the Mark-0 ABM (Section 6), for which we unveil some behaviour that had escaped previous scrutiny.

2 Extracting Stiff and Sloppy Parameter Combinations

In a first step, we define stiff and sloppy directions in parameter space. These directions are given by the eigenvectors of the Hessian Matrix of the change in observable outputs, which represents the relative parameter importance. Specifically, the eigenvector corresponding to the largest eigenvalue represents the stiffest linear combination of parameters, while the eigenvector corresponding to the smallest eigenvalue is the sloppiest direction. This approach follows from the work of Gutenkunst et al.

(2007), and was most recently implemented in Hsu et al. (2020) for stochastic models.

More precisely, let $\Phi = \{\Phi_1, \dots, \Phi_p\}$ represent a point in P -dimensional parameter space. We quantify the change in observable outcomes due to a shift δ in parameters from Φ to $\Phi + \delta$ by the mean-squared loss function:

$$\mathcal{L}(\Phi, \delta) = \frac{1}{2SKT} \sum_s \sum_k \sum_t \left(\frac{y_{s,k,t}(\Phi + \delta) - y_{s,k,t}(\Phi)}{\|y_{s,k}(\Phi)\|} \right)^2, \quad (1)$$

where $y_{s,k,t}(\Phi)$ is the realisation of output variable $k \in \{1, \dots, K\}$ at time $t \in \{1, \dots, T\}$ for random realisation $s \in \{1, \dots, S\}$. The term $\|y_{s,k}(\Phi)\|$ is a normalisation to re-scale outputs to the same unit, and is taken to be the mean realisation. To avoid numerical issues, the output variables should all be of a similar order of magnitude.

The Hessian matrix for the mean-squared loss at point Φ is defined as

$$H_{i,j}^{\mathcal{L}}(\Phi) := \left. \frac{d^2 \mathcal{L}(\Phi, \delta)}{d \log \Phi_i d \log \Phi_j} \right|_{\delta=0}, \quad (2)$$

and we consider parameters $\Phi_i, \Phi_j \in \Phi = \{\Phi_1, \dots, \Phi_p\}$, where P is the number of parameters. We take the derivative with respect to the log parameters to focus on relative parameter changes, as most parameters have different units. Evaluating the Hessian at a predefined target point Φ yields

$$H_{i,j}^{\mathcal{L}}(\Phi) = \frac{1}{SKT} \sum_s \sum_k \sum_t \frac{1}{\|y_{s,k,t}(\Phi)\|^2} \frac{dy_{s,k,t}(\Phi)}{d \log \Phi_i} \frac{dy_{s,k,t}(\Phi)}{d \log \Phi_j}, \quad (3)$$

where the second derivative term vanishes as we evaluate $H_{i,j}^{\mathcal{L}}$ at $\delta = 0$. This reduces the computational burden to $\mathcal{O}(PS)$ function calls of the underlying model. Furthermore, since Eq. (3) is a sum of projectors, it is immediate to see that the matrix H is by construction positive definite.

Due to the simulation based nature of ABMs, we calculate the first-order derivatives numerically using a central difference approach.⁴ While numerical derivatives can be troublesome in sloppy models, we find that the eigenvalues of the Hessian matrix converge to a set of steady ratios as the simulation time T increases, which can be aided by introducing an equilibrating time T_{eq} before which all observations are dropped. We have generally observed eigenvalues to converge sequentially, with the largest eigenvalue converging first. Consequently, the values of the smallest eigenvalues are likely to be more noisy and closer together. In contrast to T and T_{eq} , the number of seeds S has little effect on convergence beyond a certain minimum threshold.

To assess the parameter directions of importance, we consider the eigenvalues $\{\nu_1, \dots, \nu_p\}$ corresponding to the sorted eigenvalues $\{\lambda_1, \dots, \lambda_p | \lambda_i \geq \lambda_j \forall i < j\}$, where λ_1 is the largest eigenvalue. The first eigenvector, ν_1 , represents the linear combination of parameters corresponding to the stiffest direction in parameter space.

3 Macroeconomic Models Considered

3.1 A simple Agent-Based Model: Mark-0

The first agent-based model under consideration is the Mark-0 model of Gualdi et al. (2015) that was expanded in Gualdi et al. (2017) and Bouchaud et al. (2018), with a recent application to the COVID crisis by Sharma et al. (2020).⁵ The Mark-0 model is a simplified medium-scale hybrid-ABM of a closed economy with an aggregate household and heterogeneous firms.⁶ The model can replicate several

macro-states depending on the parameters (including, for example, full employment and unemployment, endogenous crises, high-inflation high-output, and low-inflation low-output scenarios). While Mark-0's aggregate behaviour is most likely not quantitatively precise, it generates qualitatively plausible and generic dynamical behaviours. In the baseline model, Gualdi et al. (2015) have identified four distinct phases for the unemployment rate: full employment (FE), full unemployment (FU), residual unemployment (RU), and endogenous crises (EC). We chose the Mark-0 model because its' phase diagram has been well-studied in previous work. This allows us to compare our analysis of key parameter directions and the phase-space exploration algorithm (see Section 5) to previously studied and explored phases. We use the most recent version of Sharma et al. (2020), but exclude the central bank, which leaves a fourteen-dimensional parameter space ($P = 14$). Table 1 provides an overview of the $P = 14$ parameters that we consider in the Mark-0 model, in decreasing order of relevance to this paper.

For the purposes of this paper, we focus our attention on only the unemployment rate in the Mark-0 model (equivalent to output in Mark-0) because the phases identified by Gualdi et al. (2015) are most visible in the dynamics of the unemployment rate.⁷ For the approximation of the Hessian we use $S = 20$ seeds. We set a duration of $T = 30,000$ time steps with an initial cutoff of $T_{eq} = 10,000$, after which we find that the first two eigenvalues in the spectrum have converged to a steady state.

Table 1 – Parameters of the Mark-0 Model without Central Bank sorted in order of their relevance to this paper. The notation follows that of Sharma et al. (2020)

Symbol	Default	Definition
ρ^*	1%	Baseline interest rate
Θ	2.5	Default threshold
α_r	50	Loan rate effect on R
c_0	0.5	Baseline propensity to consume
R	2	Hiring/firing rate
$\frac{\gamma_w}{\gamma_p}$	1.0	Adjustment ratio
γ_p	0.1	Price-adjustment size
η_0	0.1	Baseline firing propensity
f	0.5	Bankruptcy effect on bank interest rates
δ	2%	Dividend share
β	2.0	Household intensity of choice
τ^R	0.5	Weight π^{ema} on $\hat{\pi}$
ϕ	0.1	Revival frequency
ω	0.2	EWMA memory

3.2 Dynamic Stochastic General Equilibrium (DSGE) Models

To contrast with the Mark-0 model, we also consider several standard DSGE models. Despite multiple shortcomings (see e.g. Fagiolo & Roventini 2017, Stiglitz 2018), DSGE models remain the working horse models of many macro-economists (Christiano et al. 2018, Blanchard 2018). The first three models are borrowed to chapters 2, 3 and 8 of Galí (2015), and represent the classic real business cycle and new Keynesian frameworks (chapter 8 being an open-economy model). These three models have $P = 9$, 12, and 13 parameters respectively, and have not been fitted to data explicitly. Rather they form the theoretical backbone of current DSGE models. We also consider three models that have been fitted empirically, namely the canonical Smets & Wouters (2007) model with $P = 35$ parameters, the financial

frictions model of Del Negro et al. (2015) ($P = 50$) and the term-premium model of Carlstrom et al. (2017) ($P = 51$). For the approximation of the Hessian in the DSGE models we use $S = 100$ seeds. We set a duration of $T = 5,000$ without an initial cutoff since these models are initialised to their steady state already. The outputs considered are the time-series of the log-deviations from the steady state.

4 Are Macroeconomic Models Sloppy?

Despite their different frameworks and levels of intricacy, the above models display the sloppy phenomenology found in systems biology models (Gutenkunst et al. 2007). Across all models, we observe that the eigenvalue spectrum (Figure 1) spans several decades. In particular, for the DSGE models (b-g) the distribution is uniform across decades, which is similar to models studied in systems biology. However, the presence of phase transitions in the Mark-0 agent-based model with all observable outputs lead to a different eigenvalue spectrum (a), which is non-uniformly distributed. Instead, there are two deviating eigenvalues and a grouping of the remainder around 10^{-5} . This occurs for two reasons: First, these stiff eigenvalues correspond to the directions of the closest phase transitions, which cause extreme changes in the model dynamics, while the latter ones cause only minor shifts in the model's steady state or within the phase. The Mark-0 was constructed to be a reduced version of the Mark-1 model (Gatti et al. 2011), such that only parameters affecting the phases remain. The second reason is that the amount of information, T , required to further distinguish the small eigenvalues in these models grows with their magnitude (as in Hsu et al. (2020)), such that small eigenvalues may not be accurately estimated and actually be more spread out than shown in Figure 1.

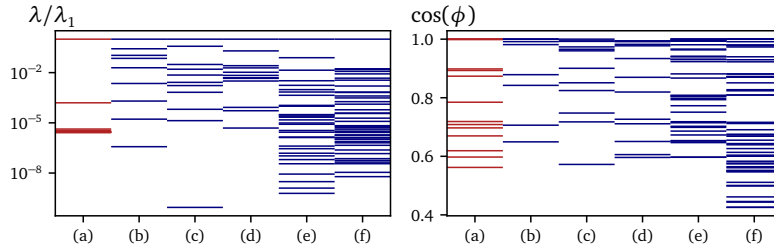


Figure 1 – (LHS) Eigenvalue Spectra relative to the largest eigenvalue (λ/λ_1), and (RHS) absolute cosine similarities between each eigenvector and its closest bare parameter axis. Note that the ranking of the cosine similarities does not necessarily correspond to that of the eigenvalues, see Table 2. The models include in red the Mark-0 ABM (Gualdi et al. 2015) with all variables evaluated in the full employment phase (a, $\Theta = 2.5$, $\rho^* = 1\%$), with the remaining parameters from Gualdi et al. (2015). In blue, the Dynamic Stochastic General Equilibrium models of (b) Galí (2015) chpt. 2, (c) chpt. 3, and (d) chpt. 8, as well as the estimated models of (e) Smets & Wouters (2007), (f) Del Negro et al. (2015), and (g) Carlstrom et al. (2017) evaluated at the authors' reported parameters.

To provide some intuition on the degree of sloppiness, the distance that can be travelled in a direction v_i without changing the observed output significantly (as measured by the loss function) is proportional to $1/\sqrt{\lambda_i}$ (Gutenkunst et al. 2007). Thus, for the Mark-0 model (Figure 1(a)), one must travel roughly $\sim 10^{5/2}$ times as far in the sloppiest direction as in the stiffest direction to achieve a change in observable output that is of similar magnitude (in the absence of any phase transitions). To put this into perspective, if one wants to empirically fit the parameters underlying these models by constraining each direction to within a 10% confidence interval, one would require 10^5 times the amount of data for the sloppiest direction in Mark-0 (Machta et al. 2013). In the context of macroeconomics, this implies a

Table 2 – Absolute cosine similarity between eigenvectors and the most aligned bare parameter axes for the first five eigenvectors of (a) the Mark-0 ABM (Gualdi et al., 2015) with all variables evaluated in the full employment phase (a, $\Theta = 2.5$, $\rho^* = 1\%$), and for only the unemployment rate in the full employment phase (FE, and FE II, $\Theta = 5.0$, $\rho^* = 1\%$), the residual unemployment phase (RU, $\Theta = 2.5$, $\rho^* = 5\%$, and RU II, $\Theta = 5.0$, $\rho^* = 5\%$), and the full unemployment phase (FU, $\Theta = 2.5$, $\rho^* = 15\%$, and FU II, $\Theta = 5.0$, $\rho^* = 12.5\%$). All remaining parameters are from Gualdi et al. (2015). The Dynamic Stochastic General Equilibrium models of (b) Galí (2015) chpt. 2, (c) chpt. 3, and (d) chpt. 8, as well as the estimated models of (e) Smets & Wouters (2007), (f) Del Negro et al. (2015), and (g) Carlstrom et al. (2017) evaluated at the authors’ reported parameters.

	<i>a</i>	<i>FE</i>	<i>FE II</i>	<i>RU</i>	<i>RU II</i>	<i>FU</i>	<i>FU II</i>	<i>b</i>	<i>c</i>	<i>d</i>	<i>e</i>	<i>f</i>	<i>g</i>
1	1.00	0.81	0.81	0.91	0.91	0.93	0.98	0.98	0.75	0.98	0.99	1.00	0.74
2	0.60	0.58	0.58	0.84	0.89	0.80	0.44	0.65	0.57	0.93	0.99	0.88	0.89
3	0.78	0.78	0.78	0.55	0.59	0.58	0.55	0.84	0.96	0.65	0.96	0.67	0.80
4	0.56	0.54	0.54	0.48	0.49	0.59	0.72	0.71	0.72	0.87	0.94	0.87	1.00
5	0.71	0.58	0.58	0.45	0.59	0.49	0.55	0.88	0.99	0.73	1.00	1.00	0.91

high data requirement that is infeasible. However, it may not be necessary for Agent-based Models to be fit to such a high degree of precision in regards to each of their parameters. As noted in Gutenkunst et al. (2007), “concrete predictions can be extracted from models long before their parameters are even roughly known” because observed outputs depend only on a few stiff directions.

The importance of the stiff directions suggests the question of what the major constituents of these linear combinations are. We first quantify the degree to which the eigenvectors correspond to the bare parameter axes by considering the cosine similarity between the eigenvectors and the bare parameter axes. Specifically, for each eigenvector we report the absolute cosine similarity of the bare axes with which it is most closely aligned.⁸ A value of one implies that the given eigenvector contains only a single non-zero parameter entry, rather than a combination of parameters. Figure 1 (right panel) shows the distribution of the cosine similarities for each eigenvector and its closest bare parameter axis. Note that the ranking here does not correspond to the ranking of eigenvalues; Table 2 shows instead the cosine similarities for the stiffest five directions.

4.1 Stiff Directions in DSGE models

Based on Figure 1, DSGE models can be described as sloppy models, with many eigenvalues orders of magnitude smaller than the largest one. This means that many combinations of parameters are in fact quite irrelevant for the dynamics of the system, and that one should be able to construct reduced models with similar performance.

The stiffest parameter directions frequently correspond to individual parameter axes, with cosine similarities close to 1. The stiffest eigenvectors are often simply the parameters related to the exogenous shock processes, such as their persistence or co-variances. For the estimated models (*f*, *g*, and *h* in Figure 1), almost all of the first 5 stiffest eigenvectors contain large elements relating to the shock processes, with a few also containing a single additional parameter (see Table 2). This reflects the adiabatic nature of these models, which are formulated as deviations from a given steady state, such that shocks in a given period dominate the dynamics of the model. In light of this, once the exogenous processes have been fixed, variations in the further parameters offer comparatively little discernible change in the models’ output.

4.2 Stiff Parameter Directions and Phase Transitions in Agent-based Models

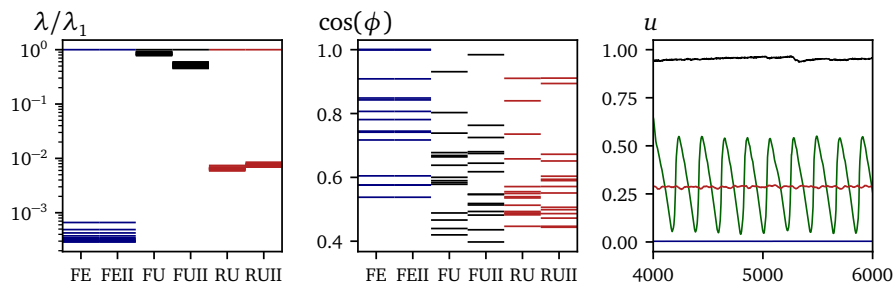


Figure 2 – (LHS) Eigenvalue Spectra relative to the largest eigenvalue (λ/λ_1), (Center) the absolute cosine similarities between each eigenvector and its most aligned bare parameter axis, and (Right) the unemployment rate for different phase, for the Mark-0 ABM (Gualdi et al. 2015) with only the unemployment rate in: the full employment phase (blue, with FE, $\Theta = 2.5$, $\rho^* = 1\%$, and FE II, $\Theta = 5.0$, $\rho^* = 1\%$), the residual unemployment phase (red, with RU, $\Theta = 2.5$, $\rho^* = 5\%$, and RU II, $\Theta = 5.0$, $\rho^* = 5\%$), and the full unemployment phase (black, FU, $\Theta = 2.5$, $\rho^* = 15\%$, and FU II, $\Theta = 5.0$, $\rho^* = 12.5\%$). The EC phase (green) is evaluated at $\Theta = 1.3$, $\rho^* = 1\%$. All remaining parameters are from Gualdi et al. (2015).

In the Mark-0 agent-based model, the stiffest directions are dominated by parameters relating to phase transitions in the model. To explore this, we focus on the phases of the unemployment rate: full employment (FE), residual unemployment (RU), endogeneous crises (EC), and full unemployment (FU) (see Figure 2). The eigenspectra for two points in three phases are shown in Figure 2. They can generally be described by an outlier first eigenvalue, followed by a dense body of the remaining eigenvalues. The degree of separation between the first and second eigenvalue depends on the phase considered, with a large deviation in the FE phase, and a small one in the FU phase.

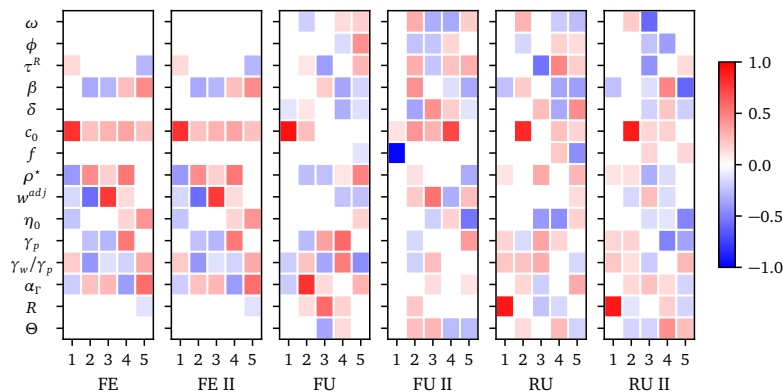


Figure 3 – Heatmaps of the components of the stiffest 5 eigenvectors of the Mark-0 ABM focused on the unemployment rate with all parameters. The sign of the eigenvector is such that the angle with the direction c_0 is positive. All components with a magnitude between -0.1 and 0.1 have been set to zero for exposition purposes. The directions were evaluated in the full employment phase (FE, $\Theta = 2.5$, $\rho^* = 1\%$, and FE II, $\Theta = 5.0$, $\rho^* = 1\%$), the residual unemployment phase (RU, $\Theta = 2.5$, $\rho^* = 5\%$, and RU II, $\Theta = 5.0$, $\rho^* = 5\%$), and the full unemployment phase (FU, $\Theta = 2.5$, $\rho^* = 15\%$, and FU II, $\Theta = 5.0$, $\rho^* = 12.5\%$). All remaining parameters are from Gualdi et al. (2015).

To understand the components of the stiffest directions, Figure 3 shows a heatmap of the stiffest five eigenvectors for the three phases considered. In each case, the direction for the first eigenvector is dominated by one of the variables which have been found to lead to phase transitions in the “handcrafted” analysis of Gualdi et al. (2015), including the parameters R and α_Γ (for a definition of these variables, see Table 1. In addition, we have here also identified c_0 to be a key variable in the phase transition. While the stiff directions are dominated by a single value, they are nonetheless unaligned with the bare parameter axes, as shown by the non-unity cosine similarities. In fact, the combination of parameters is generally one of multiple other phase-relevant parameters. For example, in the RU phase these mixture parameters include β , and γ_w and ρ^* , each of which are related to varying phase transitions. To illustrate this idea, Figure 4 shows the dynamics of the unemployment rate in the Mark-0 model starting in each phase (black lines), with log step in the stiffest (green) and sloppiest (red) directions (i.e. $\log \Phi' = \log \Phi + v_i$). Note that for the full employment case (left panel), we observe a transition to residual unemployment, while for the full unemployment phase (right panel) we observe a transition to the endogenous crises phase. For the residual unemployment phase there is no explicit transition, but a strong move towards full unemployment. Note that in all cases, there is almost no change in the dynamics when traversing the sloppy direction with an equal step size. The key observation is that the variables identified in the stiffest directions do indeed represent axes of nearby phase transitions. These results suggest that looking to the stiffest direction is a means to identify their different phases and the boundaries separating them.

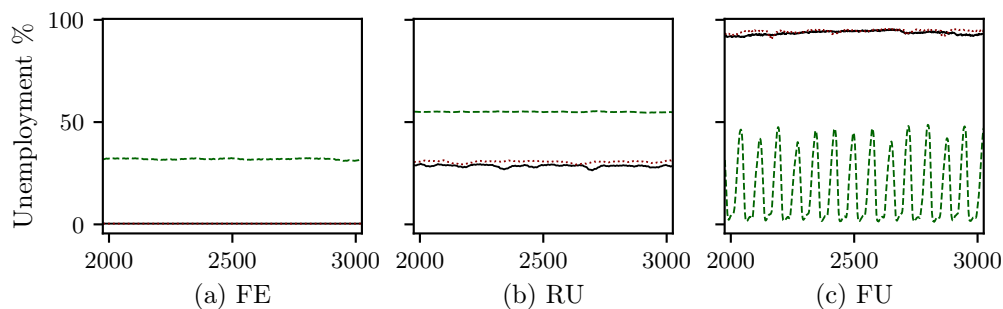


Figure 4 – Unemployment dynamics in different phases in their original form (black), with a unit log step in the stiffest (green) and sloppiest (red) directions identified in Figure 3, i.e. $\log \Phi' = \log \Phi + v_i$. The panels are: (left) the full employment phase (FE, $\Theta = 2.5$, $\rho^* = 1\%$), (center) the residual unemployment phase (RU, $\Theta = 2.5$, $\rho^* = 5\%$), and (right) the full unemployment phase (FU, $\Theta = 2.5$, $\rho^* = 15\%$). All remaining parameters are from Gualdi et al. (2015).

4.3 Transitions in Two Dimensions

For the unemployment rate dynamics in the two dimensional plane (bankruptcy threshold Θ , central bank’s baseline interest rate ρ^*), we find that the stiffest direction consistently points in the direction of a nearby phase transition. Figure 5 shows the approximate phase diagram in the $(\log \Theta, \log \rho^*)$ space (dashed black lines), together with the estimation of the first (red) and second (blue) eigenvectors for a grid of points.

The phase diagram (Figure 5 left panel) suggests that the line formed by the first eigenvector is also generally the shortest path to the phase transition. The second eigenvector then points in the direction of the second closest phase transition or parameter boundary. This is clearest when considering the

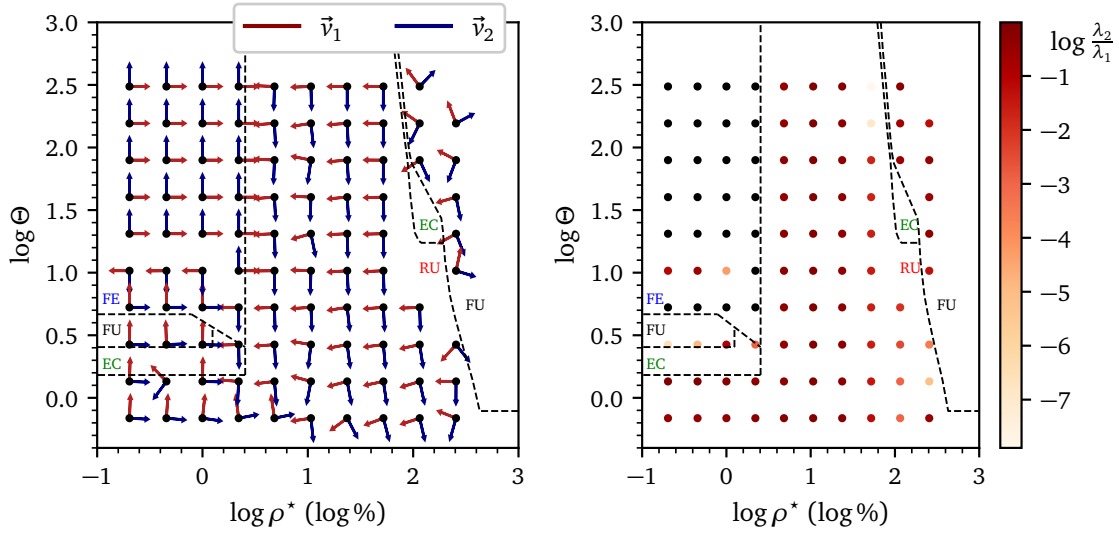


Figure 5 – (LHS) Phase diagram of the Mark-0 model in the $(\log \Theta, \log \rho^*)$ plane with projected eigenvectors. The eigenvectors, v_1 (red) and v_2 (blue), are based on the 2-dimensional Hessian with only Θ and ρ^* using unemployment as output variable. Dashed black lines indicate the locations of phase transitions. (RHS) logarithm of the ratio of the first to the second eigenvalue at different points in the phase space. The black dots indicate $\ln(\lambda_2/\lambda_1) < -10$. Simulation parameters: $T = 30,000$, $S = 20$, $T_{\text{eq}} = 10,000$. The parameters of the model, other than Θ and ρ^* , are from Gualdi et al. (2015)

intersection of FE, RU and FU near $\log \Theta \approx 0.5$ and $\log \rho^* \approx 0$. Here, at $(0, 0.4)$ the stiffest direction is the closer $\text{FU} \rightarrow \text{FE}$ transition, while for $(0.2, 0.6)$ the stiffest direction is the closer $\text{FE} \rightarrow \text{RU}$ transition. While this is true near phase boundaries, it is also true further away from them. In the FE phase, far from the boundary, the stiffest direction still points towards the closer $\text{FE} \rightarrow \text{FU}$ transition.

The proximity of a phase transition appears, as one would expect, to lead to a strong splitting in the eigenvalues. The right panel of Figure 5 shows the ratio of the first (λ_1) and second (λ_2) eigenvalues corresponding to the grid in the left panel. Near the $\text{FE} \rightarrow \text{RU}$ transition as well as the $\text{RU} \rightarrow \text{FU}$ transitions the ratio λ_2/λ_1 tends to zero. This is likely due to the high value \mathcal{L} generated by a transition over the phase boundary, which eclipses other changes. On the other hand, it is also possible that this is simply a result of a drop in importance of all other parameters, as in the FE phase the second direction is irrelevant for almost all parameter values (except near phase transitions).

The implication of these findings is twofold: first, the stiff eigenvectors are indicators for the shortest path to the closest phase transition, second, the relationship between the eigenvalues is indicative of the distance to the phase transition, and potentially the number of different transitions.

5 A New Method for Systematic Phase Exploration

Using the formalism above, we now construct a probabilistic algorithm that exploits the insight that the stiffest parameter directions indicate the shortest path to a phase transition in order to generate a maximally diverse sampling of the observable output space at a relatively low computational cost. The general principle of the algorithm is to evaluate the stiffest direction at a given point, travel along them to a new point, and repeat the exercise for a given number of steps.

5.1 Formulation of the Algorithm

To begin with, the modeller must decide a number of steps N for which the algorithm should run, as well as a starting point $\Phi(0)$ in parameter space. Furthermore, the modeller should set the hyper-parameters for the Hessian estimation: the number of random realisations S , the simulation time T and the equilibrating time T_{eq} , such that the first two eigenvalues converge to a steady ratio. The algorithm then proceeds sequentially for each step $n \in \{1, 2, \dots, N\}$ as follows:

1. Evaluate the Hessian matrix $H^{\mathcal{L}}(\Phi(n))$ and decompose it into its eigenvalues and associated eigenvectors. For the purpose of the algorithm, we consider only the first two eigenvectors, v_1 and v_2 , together with their respective eigenvalues λ_1 and λ_2 . More can be considered, but an accurate estimation of their spectrum requires an exponentially increasing simulation time T (see e.g. Hsu et al. (2020)).
2. Select the direction, $v(n)$ of the next step via:

$$v(n) = \begin{cases} v_1, & \text{with probability } \lambda_1/(\lambda_1 + \lambda_2) \\ v_2, & \text{with probability } \lambda_2/(\lambda_1 + \lambda_2) \end{cases}, \quad (4)$$

such that $\lambda(n)$ is the corresponding eigenvalue. This generally selection steps in the stiffest direction, unless both directions are important, as could be the case in the presence of multiple close phase transitions for instance.

3. Since the sign of the eigenvector returned by the algorithm is ill-defined, ensure that the direction is consistent, that is

$$v(n) = \begin{cases} v(n), & \text{if } \arccos(v(n) \cdot v(n-1)) \leq 165^\circ \\ -v(n), & \text{otherwise} \end{cases}, \quad (5)$$

which prevents the algorithm from oscillating back and forth between two points across the same phase transition.⁹

In case this is the first step (i.e. $v(n-1)$ does not exist) we determine the direction of the first step by taking a step in both the $v(n)$ and $-v(n)$ directions, yielding parameter candidates $\Phi(n)^+$ and $\Phi(n)^-$ (see step 4 for the distance) and picking the candidate with the largest loss relative to the initial parameter point $\Phi(0)$ (see Eq. (1)).

4. Pick the next point in parameter space by travelling a distance d in the v^n direction which is neither too large nor too small:

$$d = \min\left(\frac{1}{\sqrt{\lambda(n)}} \max(\varepsilon, \varepsilon_{\min} \sqrt{\lambda_1}), \varepsilon_{\max}\right) \quad (6)$$

$$\log \Phi^{n+1} = \log \Phi(n) + d \cdot v(n) \quad (7)$$

where the distance depends on three given parameters: $\varepsilon_{\min} = 0.3$, the minimum step-size to take in log-parameter space, $\varepsilon = 0.1$, the distance to travel relative to the span of the interval over which the direction can vary ($\sqrt{\lambda(n)}$), and $\varepsilon_{\max} = 1.0$, the maximum distance that can be traversed (corresponds to a factor of at most e in bare parameter space). We introduce ε_{\min} and ε_{\max}

to counteract extreme eigenvalues, which may lead to extreme steps (unrealistically large parameters) or extremely small steps such that a large number of steps is required to generate sufficient changes. The goal is to maximise the local exploration while remaining in general proximity of the initial parameters.

The most likely path followed by this algorithm is to select the top eigenvector at each step, in this way, the algorithm continually explores the direction of the largest change. Only when the first eigenvalues are similar is there a chance of bifurcating. This implies that in close proximity to a phase transition, where the eigenvalues typically diverge from one another, the direction of the phase transition is selected. Meanwhile, in areas of the parameter space where there are no immediate transitions and the eigenvalues take on a more uniform distribution, it is more likely that the algorithm explores parameter combinations along the second eigenvector as well.

5.2 Computational Cost

The computational expense of agent-based models implies that there is typically an upper limit on the number of function calls that can be made. In regards to the algorithm presented, the most costly step is the calculation of the Hessian matrix, which requires $\mathcal{O}(SP)$ function calls. Consequently, the algorithm overall requires $\mathcal{O}(NSP)$ simulations of the agent-based model. For the two-dimensional Mark-0, this implies at least $8 \times 20 \times 2 \times 2 = 640$ evaluations with a duration of $T = 30,000$ time-steps are required (note one factor of two emerges from taking central difference derivatives). For reference, a Hessian matrix for the Mark-0 with $P = 14$ parameters, running $S = 20$ seeds in parallel, and $T = 30,000$, requires ~ 45 minutes of computation time on an AMD Ryzen 9 5950X processor with 32GB RAM, while a $P = 2$ Hessian requires ~ 6 minutes. A natural benchmark for our computation is the intelligent Latin Hypercube Sampling strategy (McKay et al. 1979, Cioppa & Lucas 2007).¹⁰ For the two-dimensional Mark-0, we find that ~ 29 samples of the parameter space are required to obtain all four phases. Considering $S = 20$ random samples to mitigate any noise effects, this corresponds to $29 \times 20 = 580$ function calls, which is on par with a seven-step estimation of the algorithm. As the dimension of the parameter-space increases, we expect the number of required parameter samples for the hypercube sampling to increase at a higher rate than the algorithm presented here. Simultaneously, the information gained from the Hessian matrices may prove more valuable than a quasi-random sampling.

6 An Application to the Mark-0 Agent-Based Model

The testing bed for the algorithm is the identification of the phases of unemployment dynamics in the Mark-0 model, for which the number of parameters is $P = 14$. We find that in two, three and sixteen dimensions, our algorithm can recover all the major phases of unemployment (full employment, full unemployment and residual unemployment) within eight steps only. Due to the small parameter region for which the EC phase exists, it is often the case that it is missed in the two-dimensional case. Thus, the recovery of the endogenous crisis phase is dependent on the starting point. Despite dependence on initial conditions, in almost all cases our algorithm recovers at least phase that is different to the phase of the initial parameter choice.

In the two dimensional (Θ, ρ^*) -space we find that the algorithm covers at least two phases irrespective of the starting point or the choice of path taken. Figure 6 shows two walks originating from each of three different starting points in the (Θ, ρ^*) phase diagram. The solid lines represent the most likely walk in which the algorithm exclusively follows the stiffest direction, v_1 , at each point. The dashed lines represent an alternative random walk, which contains at least one deviation in the direction of the

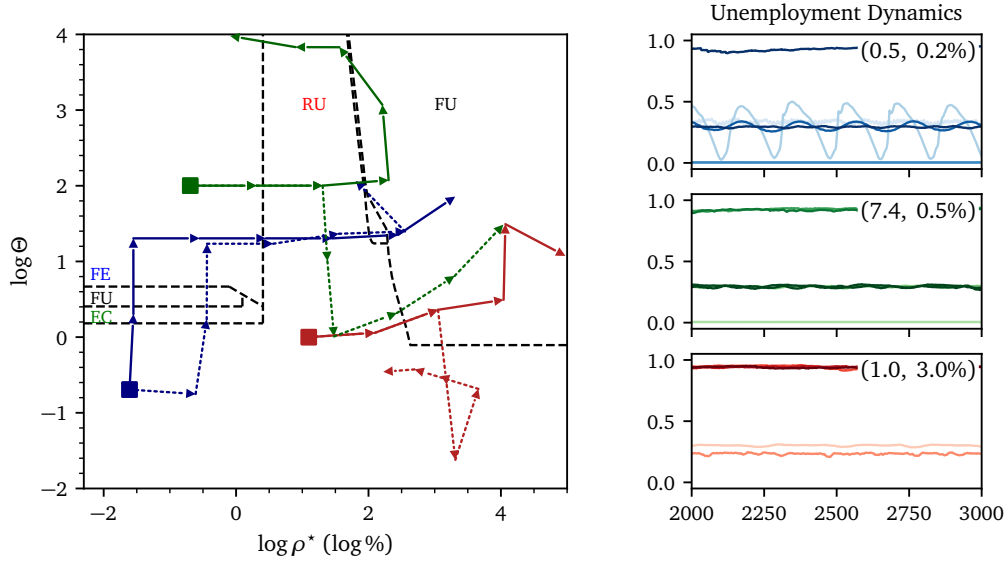


Figure 6 – (LHS) Phase diagram of the Mark-0 model without central bank in the $(\log \Theta, \log \rho^*)$ plane with two likely algorithm paths for three different starting values. Solid lines show the most likely path following only v_1 , while dotted lines indicate an alternate path with mixed steps. Dashed black lines indicate the locations of phase transitions. (b) Dynamics of the unemployment rate for different steps along the v_1 path for each starting point. Simulation parameters: $T = 30,000$, $S = 20$, $T_{\text{eq}} = 10,000$. The parameters of the model, other than Θ and ρ^* , are those of Gualdi et al. (2015). Algorithm parameters are $\varepsilon_{\min} = 0.3$, $\varepsilon = 0.1$, $\varepsilon_{\max} = 1.0$.

second stiffest direction. The sequence of choices, and the probabilities of choosing the stiffest direction at each point are shown in Table 3.

Table 3 – Probability of choosing eigenvector v_1 at each step in three paths taken from a starting point in the (Θ, ρ^*) -space. Simulation parameters: $T = 30,000$, $S = 20$, $T_{\text{eq}} = 10,000$. The parameters of the model, other than Θ and ρ^* , are those of Gualdi et al. (2015). Algorithm parameters are $\varepsilon_{\min} = 0.3$, $\varepsilon = 0.1$, $\varepsilon_{\max} = 1.0$.

Start (Θ, ρ^*)	Path	$\mathbb{P}(\vec{v}^n = \vec{v}_1)$						
		1	2	3	4	5	6	7
(0.5, 0.002)	1111111	0.53	0.70	1.00	1.00	0.98	0.62	0.64
(0.5, 0.002)	2111111	0.53	0.52	0.86	1.00	0.65	0.68	0.57
(1.0, 0.030)	1111111	0.51	0.91	0.68	0.94	0.60	0.91	-
(1.0, 0.030)	1121111	0.51	0.91	0.68	0.60	0.63	0.57	0.94
(7.4, 0.005)	1111111	1.00	1.00	0.56	0.69	0.66	1.00	0.52
(7.4, 0.005)	1122112	1.00	1.00	0.56	0.58	0.62	0.87	0.75

The effectiveness of the algorithm is dependent on its initial position within the parameter space.¹¹ The most effective walk begins near the parameter boundaries at $\Theta = 0.5$ and $\rho^* = 0.2\%$, and covers each of the unemployment phases (see the dynamics in the top right panel of Figure 6). The walk beginning in FE (green) covers the three major phases of FE, RU, and FU within three steps, while

missing out on the EC phase. This is likely due to the fact that within this two-dimensional space, the EC phase is only present in small pockets, such that it is easily missed. Finally, the second RU walk (red) is the least effective, as it recovers only the RU and FU phases before terminating at extreme ρ^* values.

The probability of following the stiffest direction (Table 3) reflects the observations made in Section 4, that the eigenvalues are drawn apart in the presence of a discontinuous phase transition. In particular, the probability of taking the stiffest direction tends to one when there is a singular close phase transition, while it tends to 50% when there is either two proximate phase transitions or no proximate phase transitions. We note that the FE phase may be an exception to this, as the second eigenvalue becomes vanishingly small throughout the phase. This may be due to the low-fluctuation steady-state nature of this phase.

While the full recovery of phases may depend on starting values, each walk identified at least two phases. This suggests that a potentially optimal exploration may require a combination of parameter space sampling combined with short iterations of the algorithm. Alternatively, since the direction of the eigenvectors v_i is indeterminate, an alternative option (when not close to parameter boundaries) would be to start the algorithm in both possible directions. In that case, starting points in the central RU phase would explore both the FE and FU phases.

6.1 A Tentative Extension to Higher Dimensions

The principal aim of our algorithm is to facilitate phase exploration of agent-based models in high-dimensional parameter spaces. We apply the algorithm to both the three dimensional and fourteen dimensional (all parameters) formulation of the Mark-0 model to identify the unemployment phases. Figure 7 shows the unemployment rate dynamics for two walks following exclusively the stiffest direction in the $P = 3$ (left) and $P = 14$ (right) cases. The starting points correspond to those of the two dimensional exploration (Figure 6).

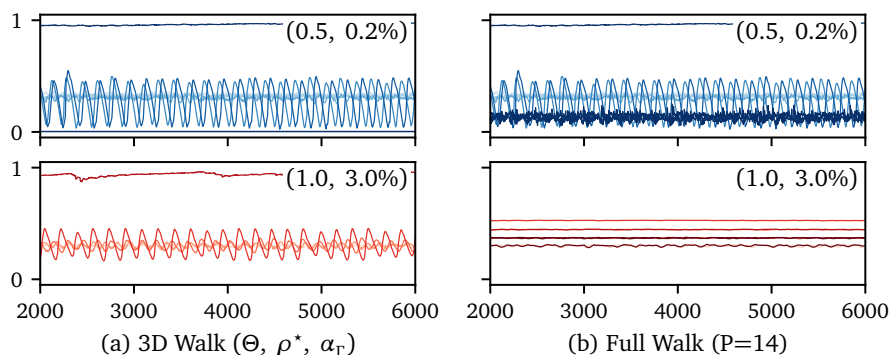


Figure 7 – Dynamics of the unemployment rate for different steps in the 3D case $(\Theta, \rho^*, \alpha_r)$ and the multi-dimensional ($P = 14$) case. We consider again the three most likely paths with 5 steps. The initial parameters of the algorithm in the (Θ, ρ^*) -space are shown in the top right-hand corner of each panel and correspond to those of Figure 6). Simulation parameters: $T = 30,000$, $S = 20$, $T_{eq} = 10,000$. All other parameters are those of Gualdi et al. (2015).

For the previously successful walk beginning at $\Theta = 0.5$, $\rho^* = 0.2\%$ we again find that in the $P = 3$ case all phases have been recovered, while the full employment phase was missed in the $P = 14$ case. For the $\Theta = 1.0$, $\rho^* = 3.0\%$ case we find that in the $P = 3$ case we obtain the same results as in the $P = 2$ case, while in the $P = 14$ case there is only a series of level shifts within the residual unemployment

phase.¹² In particular, the lower right panel suggests that no different phases are identified (such as the large level shifts between RU and FE, or different oscillatory behaviours between EC and RU). The reason behind this is the agnostic mean-squared loss function (Eq. (1)) which directly compares the unemployment time-series. This means that qualitatively small changes may dominate. One example is the small shift in periodicity of the EC phase (Figure 7 top panels) that would lead to a large loss as T becomes large, though it is qualitatively still an EC phase. This suggests that our algorithm should be improved. Two directions come to mind to make the exploration of phase space more efficient: i) including more observables in the loss function, allowing one to better discriminate between different phases; ii) adding an element of randomness in the amplitude and direction of the random walk would remove the (deterministic) dependence on the initial condition. We leave these extensions for further investigations.

The introduction of further dimensions leads to a different picture on key parameters, as well as the effective dimensions of the model. The introduction of α_r ($P = 3$) leads to a strong representation of α_r in the stiffest direction, in place of ρ^* . This is because α_r controls the vertical boundary between the FE and RU phases shown in Figure 6, such that the distance to the phase transition in the $\log \alpha_r$ direction is likely shorter than in the $\log \rho^*$ direction. A similar situation occurs also in the $\Theta = 1.0$, $\rho^* = 3.0\%$ case, where the introduction of α_r leads to less frequent changes in ρ^* . This shift is also reflected in the $P = 14$ case. The parameters with the largest changes (R , α_r , β , Θ , γ_w/γ_p) are those that have previously been identified to induce changes in phases (Gualdi et al. 2015, 2017, Bouchaud et al. 2018). Meanwhile, for both starting points (Figure 7) at least five parameters were left unperturbed, suggesting that the model phases vary in a $Q = 9$ -dimensional space, suggesting one potential avenue for the reduction of the model.

7 Conclusion

In our paper, we have shown that both agent-based models and DSGE models have a sloppy phenomenology, where their dynamics depend only on a few key stiff parameter directions, much like many models in all fields of science (Quinn et al. 2021). In the case of the Mark-0 agent-based model, the stiff parameter directions point towards close phase transitions. Exploiting these key directions, we developed an algorithm to systematically explore the phase space of agent-based models, which should apply universally (perhaps at the cost of tweaking the exploration parameters defined in Eq. (6)). Applying this algorithm to the Mark-0 model, we have recovered the four phases that were identified in the prior literature, both in low and high-dimensional cases. This method of exploration also appears to be more computationally efficient as the number of parameter dimensions grows.

Since agent-based models are able to generate a rich phenomenology, this tool may aid in understanding the possible scenarios in an agent-based model, as well as their robustness to changes in the underlying parameterisation. This suggests several avenues of further exploration, including extending our analysis to a wider variety of observable outcomes and more complex agent-based models, as well as potentially combining it with a meta-modelling approach to not only infer the set of phases but also their boundaries. Another direction, suggested to us by J. Sethna, would be to follow sloppy directions instead of stiff directions, with the aim of removing irrelevant (combination of) parameters and constructing minimal models in a systematic way once the different behaviours have been classified, along the lines of Quinn et al. (2021).

8 Acknowledgements

We thank the members of the EconophysiX lab, as well as Francesco Zamponi for their comments and suggestions. We also thank James P. Sethna for his thoughtful comments and suggestions. Finally, we thank Davide Luzzati and the Editors for their time and effort. This research was conducted within the Econophysics & Complex Systems Research Chair, the latter under the aegis of the Fondation du Risque, the Fondation de l'École polytechnique, the École polytechnique and Capital Fund Management. Karl Naumann-Woleske also acknowledges the support from the New Approaches to Economic Challenges Unit at the Organization for Economic Cooperation and Development (OECD).

References

- Apgar, J. F., Witmer, D. K., White, F. M. & Tidor, B. (2010), 'Sloppy models, parameter uncertainty, and the role of experimental design', *Molecular BioSystems* **6**(10), 1890–1900.
- Arthur, B. (2022), Some Thoughts on Agent-based Modeling and the Role of Computation in Economics, in 'Handbook of Complexity Economics', Routledge, London, New York.
- Bargigli, L., Riccetti, L., Russo, A. & Gallegati, M. (2020), 'Network calibration and metamodeling of a financial accelerator agent based model', *Journal of Economic Interaction and Coordination* **15**(2), 413–440.
- Blanchard, O. (2018), 'On the future of macroeconomic models', *Oxford Review of Economic Policy* **34**(1-2), 43–54.
- Bookstaber, R. (2017), *The End of Theory*, Princeton University Press, Princeton, N.J.
- Bouchaud, J.-P. (2020), 'To make sense of complex systems, send in the agents', *Risk.net* .
- Bouchaud, J.-P., Gualdi, S., Tarzia, M. & Zamponi, F. (2018), 'Optimal inflation target: Insights from an agent-based model', *Economics: The Open-Access, Open-Assessment E-Journal* .
- Carlstrom, C. T., Fuerst, T. S. & Paustian, M. (2017), 'Targeting Long Rates in a Model with Segmented Markets', *American Economic Journal: Macroeconomics* **9**(1), 205–242.
- Chen, S. & Desiderio, S. (2021), 'A Regression-Based Calibration Method for Agent-Based Models', *Computational Economics* .
- Christiano, L. J., Eichenbaum, M. S. & Trabandt, M. (2018), 'On DSGE Models', *Journal of Economic Perspectives* **32**(3), 113–140.
- Cioppa, T. M. & Lucas, T. W. (2007), 'Efficient Nearly Orthogonal and Space-Filling Latin Hypercubes', *Technometrics* **49**(1), 45–55.
- Del Negro, M., Giannoni, M. P. & Schorfheide, F. (2015), 'Inflation in the Great Recession and New Keynesian Models', *American Economic Journal: Macroeconomics* **7**(1), 168–196.
- Dosi, G., Napoletano, M., Roventini, A. & Treibich, T. (2017), 'Micro and Macro Policies in the Keynes+Schumpeter Evolutionary Models', *Journal of Evolutionary Economics* **27**(1), 63–90.
- Dosi, G. & Roventini, A. (2019), 'More is different ... and complex! the case for agent-based macroeconomics', *Journal of Evolutionary Economics* **29**(1), 1–37.
- Epstein, J. M. (1999), 'Agent-based computational models and generative social science', *Complexity* **4**(5), 41–60.
- Fagiolo, G. & Roventini, A. (2017), 'Macroeconomic Policy in DSGE and Agent-Based Models Redux: New Developments and Challenges Ahead', *Journal of Artificial Societies and Social Simulation* **20**(1), 1.
- Galí, J. (2015), *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework*, second edn, Princeton University Press, Princeton, N.J.

- Gatti, D. D., Desiderio, S., Gaffeo, E., Cirillo, P. & Gallegati, M. (2011), *Macroeconomics from the Bottom-Up*, New Economic Windows, first edn, Springer Science & Business Media, Milano.
- Gualdi, S., Tarzia, M., Zamponi, F. & Bouchaud, J.-P. (2015), ‘Tipping points in macroeconomic agent-based models’, *Journal of Economic Dynamics and Control* **50**(C), 29–61.
- Gualdi, S., Tarzia, M., Zamponi, F. & Bouchaud, J.-P. (2017), ‘Monetary policy and dark corners in a stylized agent-based model’, *Journal of Economic Interaction and Coordination* **12**(3), 507–537.
- Gutenkunst, R. N., Waterfall, J. J., Casey, F. P., Brown, K. S., Myers, C. R. & Sethna, J. P. (2007), ‘Universally Sloppy Parameter Sensitivities in Systems Biology Models’, *PLOS Computational Biology* **3**(10), e189.
- Haldane, A. G. & Turrell, A. E. (2019), ‘Drawing on Different Disciplines: Macroeconomic Agent-Based Models’, *Journal of Evolutionary Economics* **29**(1), 39–66.
- Hsu, C.-T. J., Brouhard, G. J. & François, P. (2020), ‘Numerical Parameter Space Compression and Its Application to Biophysical Models’, *Biophysical Journal* **118**(6), 1455–1465.
- Kaplan, G. & Violante, G. L. (2018), ‘Microeconomic Heterogeneity and Macroeconomic Shocks’, *Journal of Economic Perspectives* **32**(3), 167–194.
- Kirman, A. & Gallegati, M. (2022), A Quarter of a Century of Complex Economics: How Far Have We Come?, in ‘Handbook of Complexity Economics’, Routledge, London, New York.
- Lamperti, F., Roventini, A. & Sani, A. (2018), ‘Agent-Based Model Calibration Using Machine Learning Surrogates’, *Journal of Economic Dynamics and Control* **90**, 366–389.
- LeBaron, B. & Tesfatsion, L. (2008), ‘Modeling Macroeconomies as Open-Ended Dynamic Systems of Interacting Agents’, *American Economic Review* **98**(2), 246–250.
- Ligmann-Zielinska, A., Siebers, P.-O., Magliocca, N., Parker, D. C., Grimm, V., Du, J., Genek, M., Radchuk, V., Arbab, N. N., Li, S., Berger, U., Paudel, R., Robinson, D. T., Jankowski, P., An, L. & Ye, X. (2020), ‘One Size Does Not Fit All’: A Roadmap of Purpose-Driven Mixed-Method Pathways for Sensitivity Analysis of Agent-Based Models’, *Journal of Artificial Societies and Social Simulation* **23**(1), 6.
- Machta, B. B., Chachra, R., Transtrum, M. K. & Sethna, J. P. (2013), ‘Parameter Space Compression Underlies Emergent Theories and Predictive Models’, *Science* **342**(6158), 604–607.
- Mannakee, B. K., Ragsdale, A. P., Transtrum, M. K. & Gutenkunst, R. N. (2016), Sloppiness and the Geometry of Parameter Space, in L. Geris & D. Gomez-Cabrero, eds, ‘Uncertainty in Biology’, Vol. 17, Springer International Publishing, Cham, pp. 271–299.
- McKay, M. D., Beckman, R. J. & Conover, W. J. (1979), ‘A Comparison of Three Methods for Selecting Values of Input Variables in the Analysis of Output from a Computer Code’, *Technometrics* **21**(2), 239–245.
- Quinn, K. N., Abbot, M., Tanstrum, M., Machta, B. B. & Sethna, J. P. (2021), Information geometry for multiparameter models: New perspectives on the origin of simplicity.
- Sharma, D., Bouchaud, J.-P., Gualdi, S., Tarzia, M. & Zamponi, F. (2020), ‘V-, U-, L-, or W-shaped recovery after COVID? insights from an Agent Based Model’, *PLoS One* **16**(3).

- Sims, C. A. (1980), ‘Macroeconomics and Reality’, *Econometrica* **48**(1), 1.
- Smets, F & Wouters, R. (2007), ‘Shocks and frictions in US business cycles: A Bayesian DSGE approach’, *American economic review* **97**(3), 586–606.
- Stiglitz, J. E. (2018), ‘Where modern macroeconomics went wrong’, *Oxford Review of Economic Policy* **34**(1-2), 70–106.
- ten Broeke, G., van Voorn, G., Ligtenberg, A. & Molenaar, J. (2021), ‘The Use of Surrogate Models to Analyse Agent-Based Models’, *Journal of Artificial Societies and Social Simulation* **24**(2), 3.
- Tesfatsion, L. (2002), ‘Agent-Based Computational Economics: Growing Economies from the Bottom Up’, *SSRN Electronic Journal*.
- Transtrum, M. K., Machta, B. B., Brown, K. S., Daniels, B. C., Myers, C. R. & Sethna, J. P. (2015), ‘Perspective: Sloppiness and emergent theories in physics, biology, and beyond’, *The Journal of Chemical Physics* **143**(1), 010901.
- Transtrum, M. K., Machta, B. B. & Sethna, J. P. (2010), ‘Why are Nonlinear Fits to Data so Challenging?’, *Physical Review Letters* **104**(6), 060201.
- Transtrum, M. K., Machta, B. B. & Sethna, J. P. (2011), ‘Geometry of nonlinear least squares with applications to sloppy models and optimization’, *Physical Review E* **83**(3), 036701.
- van der Hoog, S. (2019), ‘Surrogate Modelling in (and of) Agent-Based Models: A Prospectus’, *Computational Economics* **53**(3), 1245–1263.
- Zhang, Y., Li, Z. & Zhang, Y. (2020), ‘Validation and Calibration of an Agent-Based Model: A Surrogate Approach’, <https://www.hindawi.com/journals/ddns/2020/6946370/>.

Notes

¹ In particular we refer to the seminal work of Gutenkunst et al. (2007), with follow-on work by Apgar et al. (2010), Transtrum et al. (2010, 2011, 2015), Machta et al. (2013), Mannakee et al. (2016), and Hsu et al. (2020)

² A phase transition describes the change in state of a medium, such as the conditions at which water freezes (liquid → solid). In this case we refer to tipping points, where small changes in parameters drastically alter system dynamics. For example, changing from full employment to full unemployment.

³ There are also applications to parameter calibration in Bargigli et al. (2020) and Chen & Desiderio (2021)

⁴ Central Differences imply $\frac{dy_{s,k,t}(\Phi^*)}{d \log \Phi_i} \approx [y_{s,k,t}(\log \Phi^* + \delta e_i) - y_{s,k,t}(\log \Phi^* - \delta e_i)] / 2\delta$, where e_i is a vector of zeros such that only element i equals to one, and δ is the step-size taken in log space. The choice of δ is arbitrary and potentially model dependent. In our case, we have chosen $\delta = 0.1$ for the agent-based models to ensure a sufficiently large step to avoid noise, while remaining small enough to avoid numerical instability. For the DSGE models we found $\delta = 0.1$ to be unstable, and so reduced to $\delta = 0.001$.

⁵ We refer to these papers for detailed pseudo-code as well as a mathematical description of the model and its properties. In particular, Sharma et al. (2020) provide a full description.

⁶ Hybrid refers here to the characteristic that only one set of agents (firms in the Mark-0 model) are represented by a multiplicity of agents, while other groups (e.g. households) are simply formulated as an aggregate dynamic or representative agent

⁷ In Figure 1(a) we do consider all possible output series when calculating the Hessian. To ensure consistent scales when evaluating all output series, we apply the transformation $\hat{x}_t = \ln(x_t + c)$ with $c = 10^7$, to re-scale each of the observable output series x . Without this transformation, some variables (e.g. inflation) dominate as the $\|y_{s,k}(\Phi^*)\|$ term in Eq. (1) becomes extremely small.

⁸ The cosine similarity between two vectors is defined by $\cos \phi(v_i, v_j) = \frac{v_i \cdot v_j}{\|v_i\| \|v_j\|}$. In our case, we compare each eigenvector v_i to an identity vector e_p representing parameter direction p , where $e_{ip} = 0 \forall i \neq p$ and $e_{pp} = 1$. From this, we extract the cosine similarity of the eigenvector to the most aligned parameter axis: $sim_i = \max_{p \in P} |\cos \phi(v_i, e_p)|$.

⁹It is possible to choose the sign of the eigenvector randomly. This non-ballistic motion might lead to a denser exploration of the phase space that is less sensitive to the initial conditions and choice of direction.

¹⁰For Latin Hypercubes, the parameter space is divided into equally probable intervals and then sample points are placed such that there is only one sample in each axis-aligned hyperplane

¹¹This could potentially be mitigated by randomising the first direction, or simply walking in both directions

¹²Note that increasing the step size ϵ to 0.3 leads to a wider array of levels explored between the RU and FU phases, at the cost of missing the EC and FE phases in the $\Theta = 0.5, \rho^* = 0.2\%$ case