Probing the mechanism: lending rate setting in a data-driven agent-based model

Papadopoulos, Georgios

University of Cyprus, Economics Research Centre

4 September 2020

Probing the mechanism: lending rate setting in a data-driven agent-based model

Georgios Papadopoulos*

Economics Research Centre, University of Cyprus, Nicosia, Cyprus

Abstract

The mechanism underlying banks' interest rate setting behaviour is an important element in the study of economic systems with important policy implications associated with the potential of monetary and -recently- macroprudential policies to affect the real economy. In the agent-based modelling literature, lending rate setting has so far been modelled in an ad-hoc manner, based almost exclusively on theoretical grounds with the specifics usually chosen in an arbitrary fashion. This study tries to empirically identify the mechanism that approximates the observed patterns of consumer credit interest rates within a data-driven, agent-based model (ABM). The analysis suggests that there is heterogeneity across countries, both in terms of the rule itself as well as its specific parameters and that often a simple, borrower-risk only mechanism adequately approximates the historical series. More broadly, the validation exercise shows that the model is able to replicate the dynamics of several variables of interest, thus providing a way to bring ABMs "close to the data".

JEL classification: C63; E21; E27; E43
Keywords: Agent-based modelling, Lending rate mechanism, Consumer credit, Model validation, Rule discovery

"... you don't know the rules of the game, but you're allowed to look at the board [...] and from these observations you try to figure out what the rules of the game are, what the rules of the pieces moving are."
Richard P. Feynman[1]

1 Introduction

The primary, empirical topic of this study is the investigation of the interest rate setting mechanism for consumer credit. Essentially, it tries to answer the question: "Which rule better approximates

*I would like to thank Marija Drenkovska, Malgorzata Mitka, Elisa Reinhold and internal Bank of Slovenia seminar participants for useful comments and suggestions. The views (and, of course, any errors) are mine and should not be attributed to the aforementioned individuals or institution.
Email: gpapad.gr@gmail.com

[1] "The Rules of the Game" in *The pleasure of finding things out.* Helix Books, 1999

banks' interest rate setting in a country-specific, data-driven agent-based model?" In a more fundamental level, it describes an attempt to discover the agent-specific rule which yields a close match between the simulated and historical timeseries.

The literature in agent-based modelling (ABM)[2] in Economics has evolved significantly over time, moving from the development of abstract models to elaborate descriptions of economic systems and phenomena providing very insightful contributions along the way.

In macroeconomic ABMs in particular, a large body of research can be grouped into a few main model families (Dawid and Delli Gatti, 2018). These models have been very successful in describing various aspects of the macroeconomy from the bottom-up and matching emergent patterns and stylized facts. Nevertheless, the ABM literature has so far focused, almost exclusively, on the replication of empirical regularities without being able to reproduce the evolution of observed timeseries. Indeed, matching real-world timeseries is a challenging task (Farmer and Foley, 2009) which only a handful of studies have managed to accomplish so far.[3]

One such study is from Delli Gatti et al. (2011) where the authors employ data from about 25,000 Italian firms in the period 1998-2002 to develop an ABM. Their validation exercise reveals a very good fit between the simulated and observed timeseries of firms' average interest rate of debt. Another study, focusing on the housing market, is from Geanakoplos et al. (2012). Based on data from 2.2 million households in Washington, DC during the period 1997-2010 manages to replicate the temporal evolution of several housing market indices. The third related study is a detailed ABM of the Austrian economy by Poledna et al. (2020). Using data from Eurostat for the period 1997Q1-2010Q1, compares favourably to its VAR and DSGE counterparts in out-of-sample forecasts of several macroeconomic aggregates. Finally, the study of Papadopoulos (2019) takes an alternative approach. Instead of a fully-fledged macroeconomic ABM, the author uses publicly available data from 2000-2018 and injects them into a small-scale ABM as a proxy of agents' interactions. The simulated output fits very well the historical series of household consumption and consumer credit for a number of countries.

This paper employs the latter approach to investigate the underlying mechanism which could possibly give rise to a certain, observed macro-pattern. In particular, a series which was used in the original model as an input, is replaced by an endogenous, agent-specific rule.[4] The idea behind the replacement of an input series is that if all but one element in the model follow their historical paths, then one can investigate which rule yields the most adequate approximation of the replaced series. This is done in the context of banks' consumer credit interest rate setting mechanism.[5]

In the ABM literature, interest rate setting is largely modelled in an ad-hoc manner (for an overview, see Table 7 in Appendix A and references therein). Nevertheless, there are some broad, common features among them. Typically, the charged interest rate is expressed as mark-up on a baseline rate, usually the policy rate of the central bank. Another common characteristic is that, with the exception of ABMs focusing on the housing market, the predominant loan type modelled is uncollateralised credit to firms.

With regard to the exact interest rate setting rule, in spite of the study-specific heterogeneity, four main mechanism groups can be distinguished. The first group treats the interest rate either

[2]In the paper, the "M" in the term *ABM* refers either to the modelling approach or a model itself.

[3]For a recent review of the advances in ABM validation and calibration, including a comparison of several calibration methods' performance, the work of Platt (2020) is an excellent reference.

[4]In this study both terms *rule* and *mechanism* refer to the mathematical representation which approximates the bank's interest rate setting behaviour. The particular equations implemented are elaborated in subsection 2.3.4.

[5]The terms *interest rate* and *lending rate* will be used interchangeably for the rest of the study, referring to the lending rate for uncollateralised consumer credit.

as constant or as a fixed mark-up on the baseline rate. This highly stylized mechanism leaves the monetary policy as the sole determinant of the lending rate thus, isolating it from the borrowers' and the lenders' financial conditions. The second group considers counterparty risk in the determination of the interest rate. The most frequently used proxy for borrower creditworthiness is their leverage, specified as a ratio of their total debt to some measure of their repayment ability such as their disposable income (for households) or net worth or cash flow (for firms). An increasing function of the chosen measure of borrower credit risk, on top of the baseline rate, defines the charged interest rate and quantifies the link between the premium asked for and the associated risk of credit supplied. Another mechanism group disregards borrower risk but incorporates lenders' financial health into the calculation of the interest rate. Since there are very few representatives of this type of mechanism, it is difficult to identify a general pattern regarding this rule's specifics. The final mechanism group combines all previous elements; the policy rate, a borrower- and a lender-related component. As before, these studies follow the consensus regarding the direction of the relationship between the interest rate and borrower risk; the higher the potential risk, the higher the premium required by the bank. The implementation of the bank-specific component varies across studies, ranging from a simple, stochastically varying term (proxying operational costs) to behavioural rules based on the bank's ability to lend all possible funds. The most frequent approach, though, is to link the interest rate with the lender's financial soundness via a decreasing function; i.e. the higher the net worth or the lower the leverage (in terms of total loans to equity share) of a lender, the lower the interest rate charged and vice versa. Interestingly, in many studies the choice of the specific functional form implemented is not thoroughly discussed and the associated parameters are usually chosen in an arbitrary fashion.

The empirical literature on the subject provides some mixed evidence. On the one hand it is reasonably established that the policy rate is the basis upon which banks' interest rates are formed (Freixas and Rochet, 2008). In addition, there is a general agreement about the positive relationship of the latter with borrowers' risk. On the other hand, the link between banks' capital adequacy (the ratio of capital plus reserves to total assets) is unclear. Some studies indicate a positive (Baugnet and Hradisky, 2004; Valverde and Fernández, 2007; Claeys and Vander Vennet, 2008; Gambacorta, 2008), others a negative (Brock and Franken, 2003; Horváth and Podpiera, 2012), some an insignificant relationship (Almarzoqi and Naceur, 2015) and others even a state-dependent one (Osborne et al., 2017). In addition to the previous studies, data from an online repository[6] of studies from the Bank for International Settlements (BIS) (Boissay et al., 2019) show that the majority of the literature suggests a non-negative impact of bank's capital adequacy on lending rates with only a few documenting a negative one.[7]

Combining the modelling choices from the ABM literature with the findings from the empirical studies, this paper attempts to discover the rule that potentially underlies banks' consumer credit interest rate setting behaviour. Section 2 describes the model, its initialisation and mechanics while section 3 the data used as an input. Section 4 presents and discusses the results and examines the ability of the model to replicate the historical dynamics of several variables of interest. Finally, section 5 concludes.

[6]https://stats.bis.org/frame/ [Accessed August 2020]

[7]From 16 studies, as of the latest update (March 2019), 38 from 41 estimates report a non-negative impact while only 3 a negative one.

2 The model

This section presents the structure of the model, providing a description of the agents and the rules that govern their interactions in the various markets. The model closely follows the data-driven, agent-based model developed in Papadopoulos (2019). It keeps most of the original model's inner workings as well as its data-driven nature unchanged, thus it retains its closeness to the observed time series.

However, from a modelling perspective, it has a fundamental, structural extension; an endogenous, behavioural rule replaces one of the input series. In particular, this study explores several interest rate setting mechanisms to identify which one most adequately approximates the historical series of consumer credit interest rates. The basic elements of the model and its extension are elaborated below.

The most sophisticated agents in the model are the households, followed by the banking sector which is modelled as a single agent representing the entire banking system. The central bank and firms are not explicitly modelled but their actions and interactions with the other agents are proxied by historical time series which are injected into the simulation. Finally, the government is modelled in a very elementary manner, providing an unemployment benefit to the households in need.

Households form their income expectations in a boundedly rational way and, based on those, devise their consumption plans. If their own financial resources are sufficient to meet their desired consumption level they do so and deposit any remaining funds. Alternatively, they ask for credit from the bank. The bank decides how much credit to extend and at what price on the basis of both internal properties as well as external factors. The former include the bank's risk tolerance and capital adequacy, while the latter are the potential borrower's risk and the central bank's policy rate. The market modelled in the most detail is the credit market, with endogenous interactions taking place between potential borrowers and the bank. The labour market is passively modelled, with the historical evolution of income and unemployment emulating the results of the interaction between households and firms.

The following subsections present in detail the mechanics of the model from its initialisation to a comprehensive description of the agents and the rules that underlie their behaviour.

2.1 Initialisation

The goal of the initialisation process is to reconstruct the prevailing economic environment at the chosen origin of the simulation and embed the model in it. In that way, it ensures that the simulation begins as closely as possible to the specific country it attempts to model.

The data-driven nature of the agent-based model starts with the initialisation of households' income. Based on historical data, income is distributed among households according to a $\Gamma(\alpha, 1/\lambda)$ distribution. Empirical studies show that this functional form adequately describes the distribution of income in several countries over time (Bandourian et al., 2002). Moreover, fine-tuning the distribution's shape (α) and scale (λ) parameters allows the simulated and historical distributions match in terms of minimum and average values as well as in their dispersion.

Another key variable, initialised according to historical data, is unemployment. The percentage of unemployed persons in the total working population in the simulation mirrors its real-world counterpart and thus determines the exact number of unemployed households. Therefore, after income is distributed, the appropriate number of random households is chosen and their income is fixed equal to the unemployment benefit.

After initialisation, the model runs for a burn-in period of N time steps without any changes in the amount of income that households receive nor their employment status. At the end of the burn-in period the model is considered to have reached its equilibrium state and the simulation starts. The evolution of every variable is thereafter either fed into the model as an exogenous input, or it is endogenously generated as a result of agents' interactions.

2.2 Sequence of events

The timeline of events in the simulation is presented below:

1. Historical data of household- and bank-related variables are updated.
2. Households collect interest from any deposits they might keep at the bank.
3. Employed households receive their monthly income according to the historical figures of income's growth.
4. The real-world data on unemployment determine the employment status of households.
5. Households form their income expectations.
6. Households try to meet any financial obligations they might have and form their desired consumption plans.
7. Households without the necessary financial resources to achieve the desired level of consumption ask for credit.
8. The bank estimates the maximum amount of credit it can offer along with the associated interest rate and makes its offer to the potential borrower.
9. Households decide how much to borrow and the loan's characteristics (size, interest rate) are established.
10. Households fulfill their consumption plans to the maximum extent possible and deposit any excess funds that might remain.

2.3 Agent and market description

The following subsections describe in detail the different types of agents and their behavioural rules that determine their interactions in the various markets.

2.3.1 Households

Expectations formation

Each time step in the simulation begins with households receiving -exogenously- their monthly income. Their first action is to form their income expectations. These constitute a fundamental element which subsequently determines households' desired consumption and drives their demand for credit.

Based on the respective literature on expectations formation, households form their income expectations in a boundedly rational way. More specifically, the heuristics switching model (HSM) described by Anufriev and Hommes (2012a,b) controls their expectations formation. Under the HSM, agents employ a small collection of simple forecasting heuristics and at each time step choose to use one based on its past performance and agent-specific, behavioural characteristics.

The HSM encompasses in a single framework many of the empirical regularities regarding expectations formation such as their heterogeneity within a population (Frankel and Froot, 1987; Allen and Taylor, 1990; Mankiw et al., 2003; Branch, 2004; Fehr and Tyran, 2008; Pfajfar and Santoro, 2010, among others) and the evolutionary selection of forecasting rules based on their performance (Marimon and Sunder, 1995; Arifovic and Sargent, 2003).

Moreover, it has exhibited a very good fit to experimental as well as survey data in numerous studies and in a variety of contexts ranging from financial to macroeconomic in nature (Hommes et al., 2005, 2008; Assenza et al., 2019; Hommes et al., 2019).[8] A detailed description of the HSM's mechanics is provided in Appendix B.

Desired consumption

After forming their income expectations, households devise their consumption plans. Their desired consumption ($C^d_{h,t}$) depends on their expected income and their capacity to sustain themselves. The last condition hinges on their ability to have a consumption equal to the needed level of subsistence and influences their solvency. The rule in Equation 1 encapsulates this behaviour:

$$C^d_{h,t} = \begin{cases} C_{min,t}, & I_{h,t} + D_{h,t-1} < LP_{h,t} + C_{min,t} \\ max\{\alpha_y \cdot I^e_{h,t+1} + \alpha_w \cdot D_{h,t-1}, C_{min,t}\}, & I_{h,t} + D_{h,t-1} \geq LP_{h,t} + C_{min,t} \end{cases} \tag{1}$$

with $1 > \alpha_y > \alpha_w > 0$ denoting the marginal propensities to consume out of income and wealth respectively.

The amount of households' liquid wealth differentiates their behaviour; it determines a household's solvency and whether or not it will ask for credit. Primarily, households' current income ($I_{h,t}$) and accumulated deposits ($D_{h,t-1}$) must allow them to maintain a minimum level of consumption, equal to the subsistence level ($C_{min,t}$). However, if households cannot meet their monthly loan payment ($LP_{h,t}$) and survive, they decide to consume as little as possible and miss any potential loan payments due. Alternatively, they service their debt and attempt to consume an amount proportional to their expected income ($I^e_{h,t+1}$) and past deposits. In order to achieve their desired level of consumption, households first use their own financial resources. If spending their income and withdrawing their deposits proves insufficient, then they turn to the bank to ask for consumer credit.

The Modigliani consumption function (Modigliani and Brumberg, 1954) described in Equation 1 is frequently used by the ABM literature (Riccetti et al., 2013; Delli Gatti and Desiderio, 2015; Gualdi et al., 2015; Riccetti et al., 2015; Assenza et al., 2015; Caiani et al., 2016; Russo et al., 2016; Gurgone et al., 2018; Reissl, 2020a,b, among others).[9]. Furthermore, the results in Section 4.2 indicate that it yields a good match between the simulated and historical aggregate series of consumption.

[8]For reviews of the recent literature on the experimental study of expectations formation the works of Assenza et al. (2014) and Cornand and Heinemann (2014) provide excellent references.

[9]The interested reader is referred to the work of Dawid and Delli Gatti (2018) and references therein for an overview of consumption functions implemented in several macroeconomic ABMs.

2.3.2 Bank

A single bank agent represents the economy's banking system. The bank takes households' deposits paying an interest of r_t^D and provides them with consumer credit after evaluating their requests.

The size ($L_{h,t}$) and interest rate ($r_{h,t}^L$) of the provided credit are estimated by the bank for each potential borrower and depend on both bank-specific and external factors.

The external factors include the central bank's key policy rate and the potential borrower's risk. Both of them have a direct effect on the supplied credit's interest rate and thus affect the volume of credit in the economy and households' solvency. The key policy rate provides the floor above which the bank adds a mark-up to generate the offered interest rate (Freixas and Rochet, 2008). The mark-up associated with borrower's risk is a convex function of their leverage, i.e. their total debt as a share of their annual income (DTI). The rationale behind this functional form is that the bank requires higher compensation when it assumes higher risk. The possible new debt, in addition to any existing debt, is taken into account in the calculation of a household's leverage. The bank estimates households' annual income simply as 12 times their current monthly income, implicitly assuming that it will remain unchanged.

The internal factors are the bank's risk tolerance and its capital adequacy. The former is proxied by the maximum debt-service-to-income ratio (DSTI) that the bank is willing to accept for any potential borrower.[10] It predominantly influences the amount of credit that the bank is prepared to provide and changes over time. From a credit supply-side perspective, a related measure of the bank's risk tolerance is households' leverage. Similarly to the maximum DSTI, the bank sets a maximum DTI target above which refuses to provide credit. Finally, the bank's capital adequacy (CAR) is a key internal factor in determining the interest rate, $r_{h,t}^L$. Since the findings of the empirical literature are not conclusive regarding the direction of the relationship between a bank's capital adequacy and loan interest rates, two cases are considered; a positive relationship and a negative one.

It should be noted that due to their nature, bank-specific factors can be used as instruments for the conduct of macroprudential policy, hence they also have an exogenous aspect. Among the factors considered in the simulation, CAR is identical to the real-world series. Therefore, it is reasonable to assume that it already incorporates any imposed regulatory constraints.

The remaining two factors (DSTI and DTI) are calibrated so that the results replicate the observed patterns and not examined as extensive, counterfactual scenarios. The main reason behind this choice is that the purpose of the study is to uncover the interest rate setting mechanism. In addition to that, data from the European Systemic Risk Board indicate that until recently most borrower-based regulatory constraints applied to real estate lending, while consumer credit has been subject to less strict measures, if at all.[11] Thus, DSTI and DTI can be considered to mostly reflect the bank's risk attitude rather than the result of tight regulation. Nevertheless, this shows the model's flexibility to be applied on a topic with high policy interest such as the impact of specific macroprudential measures on credit growth and the evolution of interest rates.

2.3.3 Credit market

In the credit market, households without the financial resources to meet their desired consumption level ask for credit from the bank. The minimum amount of consumer credit that a household asks

[10] The DSTI is the ratio of interest payments plus amortisations over the potential borrower's monthly income.

[11] For the evolution of the application of various macroprudential measures across EU jurisdictions, the interested reader can consult the respective reports from the European Systemic Risk Board (2016, 2017, 2018, 2019, 2020).

for is determined by Equation 2:

$$L^{ask}_{h,t} = C^{d}_{h,t} - (I_{h,t} + D_{h,t-1}) \tag{2}$$

Subsequently, the bank determines the maximum size of the loan that is willing to extend conditional on two constraints. The first one is associated with the potential borrower's DTI:

$$L^{max,DTI}_{h,t} = DTI^{max} \cdot (12 \cdot I_{h,t}) - B_{h,t-1} \tag{3}$$

Equation 3 directly determines the maximum amount of credit that can be provided as a fixed fraction (DTI^{max}) of the borrower's annual income, taking into account any pre-existing debt they might currently have ($B_{h,t-1}$). The associated maximum monthly payment with $L^{max,DTI}_{h,t}$ is:

$$LP^{max,DTI}_{h,t} = \frac{r^{L,DTI}_{h,t} \cdot (L^{max,DTI}_{h,t} + B_{h,t-1})}{1 - (1 + r^{L,DTI}_{h,t})^{-m}} \tag{4}$$

The second constraint is linked with the household's DSTI. The respective maximum monthly payment is defined as:

$$LP^{max,DSTI}_{h,t} = DSTI^{max}_{t} \cdot I_{h,t} \tag{5}$$

Based on $LP^{max,DSTI}_{h,t}$, the largest amount of credit that the bank can provide is:

$$L^{max,DSTI}_{h,t} = \frac{LP^{max,DSTI}_{h,t} \cdot [1 - (1 + r^{L,DSTI}_{h,t})^{-m}]}{r^{L,DSTI}_{h,t}} - B_{h,t-1} \tag{6}$$

The maturity, m, of the loan is fixed and measured in months. The interest rates, $r^{L,DTI}_{h,t}$ and $r^{L,DSTI}_{h,t}$, can be a function of the central bank's policy rate, the potential borrower's leverage and the bank's capital adequacy. In general, the two constrains imply different values for the maximum interest rate. The detailed interest rate setting mechanism is elaborated in the next subsection.

The final loan offer is controlled by the most binding constraint between the DTI- and DSTI-related one:

$$LP^{offer}_{h,t} = min\{LP^{max,DSTI}_{h,t}, LP^{max,DTI}_{h,t}\} \tag{7}$$

Depending on which one is the minimum, all the remaining parameters such as the size of the loan ($L^{offer}_{h,t}$) and its interest rate are set accordingly.[12]

Finally, Equation 8 describes household's decision regarding the volume of debt it will assume:

$$L_{h,t} = \begin{cases} L^{offer}_{h,t}, & L^{ask}_{h,t} \geq L^{offer}_{h,t} \\ U(L^{ask}_{h,t}, L^{offer}_{h,t}), & L^{ask}_{h,t} < L^{offer}_{h,t} \end{cases} \tag{8}$$

Therefore, if the size of the loan offered is smaller than what the household demanded, it assumes as much as the bank supplies. Alternatively, it chooses a random amount in the range between $L^{ask}_{h,t}$ and $L^{offer}_{h,t}$. Once the household and the bank agree on the size of the loan, $L_{h,t}$, the interest rate $r^{L}_{h,t}$ and the monthly payment $LP_{h,t}$ are calculated and established.

[12]Equivalently, the Equation $L^{offer}_{h,t} = min\{L^{max,DSTI}_{h,t}, L^{max,DTI}_{h,t}\}$ can be used to determine the parameters of the loan offer.

The stochastic credit assumption mechanism in Equation 8 is the same as in Papadopoulos (2019). The intuition behind the stochastic part of the rule is twofold; the bank offers the largest permitted amount because it wants to maximize its interest revenues, and households consider getting more credit since this allows them to consume more. From a modelling perspective, this approximation addresses the problem of the implausibly small loans implied by Equation 2, while at the same time replicates the dynamics of consumer credit in a satisfactory way.

2.3.4 Interest rate setting mechanism

The key part of the model is the endogenous interest rate setting mechanism. The rule that the bank uses to determine it consists of several components. The basis on which the bank puts a mark-up is the central bank's policy rate, i_t^{CB}. Together with that, two additional mark-ups are examined in this study; a mark-up based only on borrower's risk, and one with an extra component associated with the bank's capital adequacy.

Lending rate and borrower's risk

The simplest mechanism ($\mathcal{M}_\phi^{DTI}$) connects borrower's risk with the lending rate. Household leverage ($DTI_{h,t}$) proxies their risk and a convex function links it with interest rate:

$$r_{h,t}^L = i_t^{CB} + r_{h,t}^\phi = i_t^{CB} + \phi^{DTI_{h,t}} \tag{9}$$

where $\phi > 1$ is a parameter determining how sharply does the bank raise the charged interest rate as the borrower's risk increases. The rationale behind this functional form is that the bank will require an increasingly higher premium for providing credit to more leveraged and thus more risky borrowers. Evidently, higher (lower) values of ϕ reflect lower (higher) risk appetite and tighter (looser) lending standards by the bank. Although in principle there is no upper limit for ϕ's values, a reasonable range would be $1 < \phi \leq 3$. In general, the floor[13] for the final interest rate is about 1 percentage point above the central bank's policy rate. However, $\phi \in (1, 3]$ implies that the maximum cost for a household taking a loan of the size of its annual income will be 3 p.p. above the key policy rate, while for a loan twice as large it will be 3 times as much.

Lending rate and bank's capital adequacy

The most common approach regarding the interest rate setting mechanism in macroeconomic ABMs is some variation of the previous rule, i.e. a monotonically increasing function of borrowers' leverage. However, some studies employ an additional part which connects the interest rate with the lender's financial soundness (Delli Gatti et al., 2010; Cardaci, 2018; Alexandre and Lima, 2020). In particular, the assumed relationship is a negative one, meaning that a more robust, from a financial point of view, banking system provides credit at more favourable terms.

Since the findings of the empirical literature are inconclusive about the direction of the relationship between lending rates and banks' financial health, two mechanisms are considered; a positive ($\mathcal{M}_\phi^\theta$) and a negative one ($\mathcal{M}_\phi^\tau$). The measure of the bank's financial condition is its capital adequacy ratio (CAR_t). Higher CAR_t figures indicate a better capitalised and therefore, sounder banking system.

Under the positively related rule a financially healthier bank will charge a higher interest rate on its borrowers as follows:

[13] Corresponding to very low-leverage households and any value of $\phi > 1$ or any leverage level and $\phi \approx 1$.

$$r^L_{h,t} = i^{CB}_t + r^{\phi}_{h,t} + r^{\theta}_t = i^{CB}_t + \phi^{DTI_{h,t}} + log_{\theta}(CAR_t) \quad (10)$$

Compared to Equation 9, the additional component in Equation 10 is r^{θ}_t. The logarithm of CAR_t links the interest rate with the bank's capital adequacy, with $\theta > 1$ being the logarithm's base. The closer the θ is to 1, the higher the associated mark-up. For instance, for $CAR_t \simeq 8\%$[14] the resulting additional mark-up for $\theta = 1.2$ is $r^{\theta=1.2}_t = 11.4$ p.p. while the respective figure for $\theta = 5$ is almost 10 times smaller at $r^{\theta=5}_t = 1.3$ p.p. A positive relationship between r^{θ}_t and CAR_t can be the outcome of various, non mutually exclusive, bank actions. One such action is associated with the cost of equity. The better a bank is capitalised in excess of the regulatory minimum, the more compensation it will seek for the higher cost of equity. Another behaviour producing the same result could be the motivation of a less capitalised bank to accept more risk in order to receive higher returns.[15] On the flip side, a well capitalised bank has the incentive to charge higher interest rates, and thus decrease the provided volume of credit, since there is more capital at risk.

The final alternative associates the bank's capital adequacy and the offered interest rate with a negative relationship. This is expressed by the following rule:

$$r^L_{h,t} = i^{CB}_t + r^{\phi}_{h,t} + r^{\tau}_t = i^{CB}_t + \phi^{DTI_{h,t}} + (CAR_t)^{-\tau} \quad (11)$$

where $\tau > 0$ a parameter determining the strength of the relationship. A reasonable upper limit for τ would be $\tau < 1$ given that for common levels of CAR_t at around 8%, this would increase the final interest rate by $r^{\tau=0.9}_t = 9.7$ p.p. As in the previous case, there are several underlying motives which can potentially give rise to such behaviour. One theory suggests that bank capital acts as a disciplining device which forces banks to put more effort into loan monitoring. Another possibility is associated with an under-capitalised bank's motivation to lower its risk exposure in order to protect its franchise value in the event of failure. Finally, there is the hypothesis that a better capitalised bank will seek to increase its market share. All of these cases imply that bank capital is negatively associated with interest rates and positively with lending volumes.

2.3.5 Labour market

In a full-fledged ABM, firms and households would interact in the labour market determining, among others[16], the employment status and income of the latter. Since firms are not modelled explicitly, historical time series proxy their interactions with households in the simulation.

The monthly income that households are initially endowed with is calibrated to match the observed distribution of income of the simulated economies at the origin of the simulation. Subsequently, real-world data on income's growth govern its evolution in the model economy. It should be noted that the imposed path of income is applied in a uniform manner upon households. Thus, any change in the aggregate income series is translated into an equal change for every individual household's income.

Similarly to income, the observed series of unemployment rate control households' employment status. The appropriate number of households is affected by aggregate unemployment's changes in a random fashion. Unemployed households receive a dole equal to 80% of the contemporary minimum wage which allows them to maintain a minimum level of consumption.

[14] This was the average CAR_t across EU countries in early 2008.

[15] Assuming that higher returns will come from lower interest rates but increased lending volumes.

[16] The complex interactions would also affect firms' properties such as their productive capacity but this will not be discussed here.

This modelling choice is mainly dictated by the lack of more granular data on the monthly evolution of income and unemployment by households' income level. Nevertheless, the positive skewness of income's distribution implies that a randomly influencing unemployment will generally affect households at the bottom of income's distribution more than those at the top.

3 Data

A distinguishing feature of this study's ABM is its data-driven nature. Instead of modelling every part of the economy in the same detail, historical data proxy the actions of several agents and the result of their interactions in the respective markets.

Input data can be grouped in 2 major categories; scenario data and calibration data. The former are injected into the simulation and updated on every time step, whereas the latter are used to initialise the model according to the prevailing conditions of a certain economy at a specific point in time.

The country scenarios examined refer to the Cypriot, Slovenian and the UK economies, thus increasing the set of countries on which the model is applied to compared to the original (Papadopoulos, 2019). In addition to studying another economy, on top of Slovenia and the UK which were originally modelled, this is a test of the model's flexibility to accommodate different scenarios and an assessment of its performance.

An important observation regarding every data category is that their quality can significantly affect the results. Obviously, the closer the injected data are to the historical reality, the more similar are expected the simulated series to be to the realized ones. It follows that minimising data-related discrepancies between input and real-world data allows a cleaner study of the model's mechanics and in particular the interest rate setting mechanism.

3.1 Scenario data

The simulation begins in January 2000 and extends until mid-2019, covering almost twenty years of data. Each time step in the model represents one month in physical time, therefore any series collected in lower frequencies are converted to monthly using spline interpolation. Figure 1 and Figure 2 display the evolution of the country-specific, historical series which are injected into the simulation.

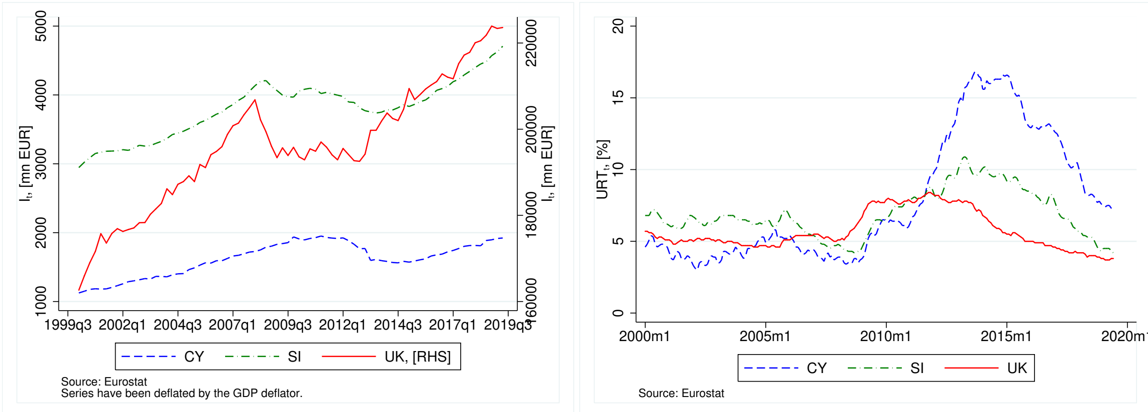

Figure 1: The original, household-related series of the macroeconomic scenario injected into the simulation. Wages & salaries (left) and unemployment rate (right).

The left panel in Figure 1 shows the series of wages & salaries[17] whose growth closely resembles the one of net disposable income which would be ideally used.[18] It is deflated by the GDP deflator, transformed into monthly frequency through spline interpolation and fed into the model in monthly growth rates. The right panel displays the evolution of unemployment (as a % of active population). It is already collected in monthly frequency and, like income, is injected into the simulation in monthly growth rates.

[17] Referred to simply as *income* for brevity for the remainder of the study.

[18] Comparison with available resources on net disposable income such as a related, yet discontinued, dataset from Eurostat (*ei_naia_q*) and another from the IDCM confirms the similarity in growth rates between the two series.

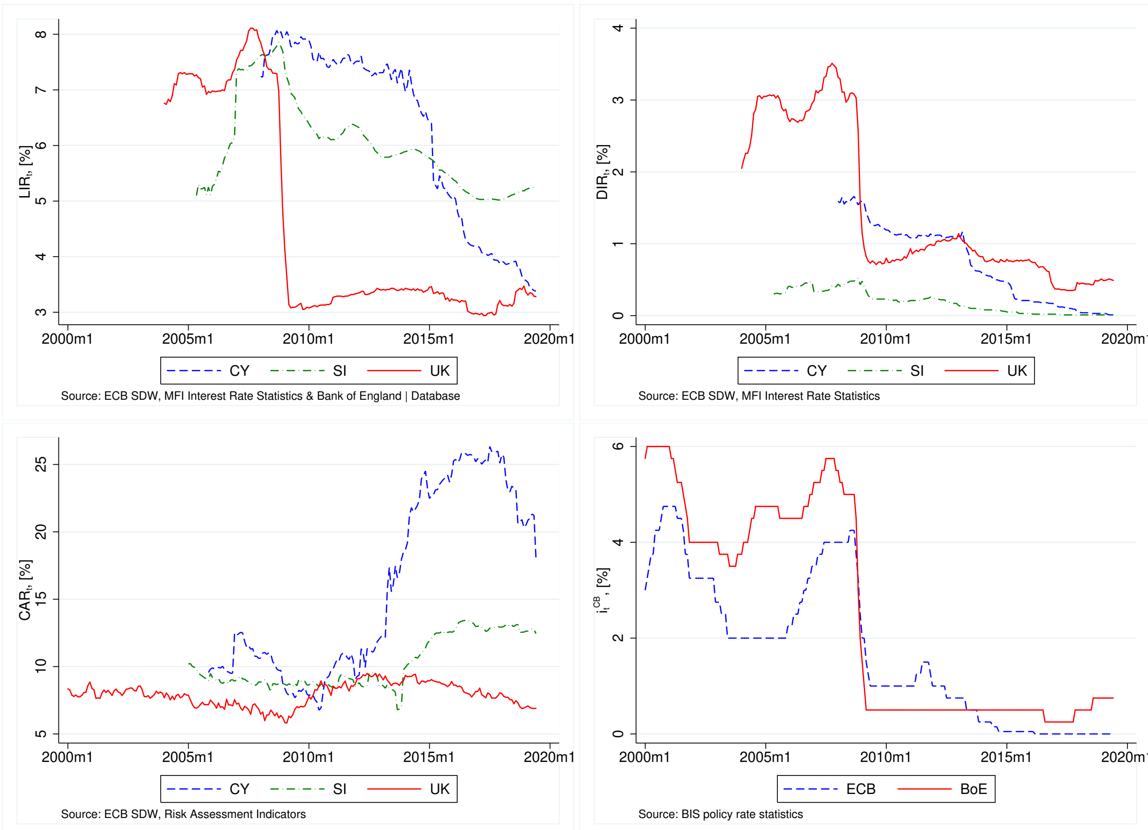

Figure 2: The original, bank-related series of the macroeconomic scenario injected into the simulation. Consumer credit interest rate (top left), overnight deposit interest rate (top right), capital adequacy ratio (bottom left) and policy rate (bottom right).

Every series shown in Figure 2 enters into the model in levels, without any interpolation. The only transformation applied is the conversion of the reported, annualized interest rates into monthly ones.[19]

The top row presents the evolution of bank lending (left) and overnight deposit (right) interest rate series.[20] The former have a dual role in the simulation; they are used as scenario input during the calibration of the DSTI series, while they are the benchmark against which the output from the various interest rate mechanisms is tested. In particular, for the case of Cyprus and Slovenia, the lending rate refers to credit for consumption and other lending with maturity above 1 and up to 5 years, denominated in euro. For the UK, it corresponds to personal loans with floating rate, in pound sterling. The top right panel depicts the evolution of overnight deposit interest rates from households in euro (CY and SI) and pound sterling (UK).

[19] CAR_t is injected as a percentage in the case of the positive interest rate rule and as a fraction $0 < CAR_t < 1$ in the negative one.

[20] In the simulation the missing values for any interest rate series have been replaced by the average of their first 12 observations.

The bottom left panel in Figure 2 displays the monthly series of the banking systems' CAR_t which is defined as the sum of capital and reserves divided by total assets.[21] Finally, the bottom right panel in Figure 2 shows the key policy rates set by the respective central banks; the European Central Bank (ECB) and the Bank of England (BoE). It should be noted that, for simplicity, it is assumed that the ECB policy rate applies for the whole period in the simulations using the scenarios for Cyprus and Slovenia. Thus, any part of the scenario that does not reflect the exact real-world settings can be considered to correspond to a fictional economy. It follows that model validation and the investigation of the interest rate setting mechanism are relevant only when the scenario mirrors the historical reality.

Overall, Figure 1 and Figure 2 demonstrate the difficult situation that households and banks had to go through during the Global Financial Crisis of 2008 (GFC) and its aftermath. The former faced declining incomes and rising unemployment while the latter slashed their interest rates and consequently their related profits.

3.2 Calibration data

Calibration data differ from scenario data in that they are static in nature and are used for initialising the simulation. The purpose of the latter is to embed the model into a specific economic environment and ensure that the simulation begins as closely as possible to the economy it attempts to model.

One subcategory of calibration data corresponds to households' income and employment status variables. The historical figures of minimum and average monthly income, as well as the Gini coefficient control the distribution of initial income among households. Gini coefficient's source is the UNU-WIDER, World Income Inequality Database (WIID)[22], whereas Eurostat's database provides the needed input on average and minimum monthly income.[23] In addition to income-related calibration data, Eurostat's database is used to calibrate initial monthly unemployment.

The second subcategory includes every bank-related variable such as loan and deposit interest rates, capital adequacy and the central bank's policy rate. In this case, calibration data and their sources are the same as the scenario data and only the observation one month prior to the simulation's start is used.

Table 1 reports the calibrated values of the initial setup for every country-specific scenario.

[21] In some studies this ratio is also referred to as leverage.

[22] Version date 6 May 2020.

[23] In particular, Eurostat's series on average annual gross earnings (*earn_gr_isco*) and monthly minimum wages (*earn_mw_cur*). For Cyprus, minimum wages are backdated from International Labour Organization's data on statutory nominal gross monthly minimum wages.

Table 1: Initial setup of the simulation for CY, SI and UK.

Description	Country scenario		
	CY	SI	UK
Average income	1480	1000	2430
Minimum income	600	350	950
Gini coefficient of income distribution (%)	28	23	35
Unemployment rate (%)	4.6	7.1	5.8
Consumer credit interest rate (%)	7.758	5.257	7.058
Overnight deposit interest rate (%)	1.594	0.348	2.621
Capital adequacy ratio (%)	-	-	8.641
Central bank policy rate (%)	3.0	3.0	5.5
DSTI (%)	42	50	56

It should be noted that the missing initial values of capital adequacy ratio in the CY and SI scenarios do not affect the results. The capital adequacy ratio is needed for the study of the interest rate setting mechanism described in Equation 11 and therefore its use is meaningful only after the two countries joined the euro area (in 2008 and 2007 respectively).

3.3 DSTI data

The case of DSTI is a special one, sharing elements of both scenario and calibration data. Similarly to scenario data, DSTI[24] is updated in every time step in the simulation. However, its path as well as its initial values are calibrated in a process discussed in detail in Appendix C.

The calibration of DSTI is necessary in the absence of exact historical data on its evolution. The only publicly available source which provides an adequate approximation is the database for debt service ratios statistics from the BIS. This database contains quarterly series of aggregate DSTI for various countries and sectors in a consistent way. The methodology for its construction is described in (Drehmann et al., 2015) and suggests that although it captures DSTI's evolution in a satisfactory manner, nevertheless the reported levels do not reflect the accurate figures one would get from micro data. Therefore, both its trajectory and initial values need to be calibrated.

The calibration procedure is based on a grid-search attempting to find the DSTI path and its associated level which yield the maximum similarity between the model-generated and historical series of consumer credit.

A fundamental assumption underlying DSTI's calibration procedure is that the evolution of the actual series is similar to the approximated series of debt service ratios in BIS's database. This is particularly relevant for Cyprus and Slovenia which are not included there. Nonetheless, even if a country included in the simulation is covered by the BIS (as is the UK), it is not necessarily the case

[24]More precisely, the maximum DSTI accepted by the bank.

that the reported, aggregate DSTI series would be a good match to the specific, real-world DSTI series for consumer credit. As is exhibited in Papadopoulos (2019), housing and consumer credit DSTIs can have very different dynamics which in turn will affect the evolution of the respective credit series. Therefore, it is worth keeping such data-quality issues in mind when interpreting the simulation's results.

3.4 Fixed parameters

Finally, a set of parameters is fixed across every simulation and is not data-related, but their calibration is guided by the literature. These parameters are reported in Table 2.

Table 2: Fixed parameter values.

Description	Parameter	Value	Equation #
Number of households	H	1000	-
Burn-in period	N	170	-
Propensity to consume out of income	α_y	0.65	1
Propensity to consume out of deposits	α_w	0.05	1
WTR expectations extrapolation factor	ω_{wtr}	0.4	12
STR expectations extrapolation factor	ω_{str}	1.3	13
ADA expectations parameter	ω_{ada}	0.65	14
Households' memory strength	η	0.7	16
Households' expectations rule persistence	δ	0.9	17
Households' intensity of choice	β	0.4	17
Loan maturity (in months)	m	60	4, 6

Both propensities to consume assume values frequently used in the respective literature (Godley and Lavoie, 2016; Assenza et al., 2015; Meijers et al., 2018, among others). Regarding the HSM-related parameters, their values are set to those identified in the original studies by Anufriev and Hommes (2012a,b). Finally, loan maturity is fixed to 5 years which is a reasonable period of time for this type of credit.[25]

4 Simulation results

The focus of the simulation is the identification of the lending rate setting mechanism by the bank. Each mechanism ($\mathcal{M}_\phi^{DTI}, \mathcal{M}_\phi^{\theta}$ and $\mathcal{M}_\phi^{\tau}$) described in subsection 2.3.4 is implemented in every

[25]Nonetheless, loan maturity can have important implications for household DSTI and therefore credit growth. However, due to absence of specific data on loan maturity, the approximation of it being fixed at 5 years is considered to be reasonable.

country scenario and their output is recorded. Subsequently, the artificial interest rate series[26] are compared to the actual ones and the one which yields the closest match between the two is selected. Finally, everything is put together and the model's output is validated against the real-world data.

4.1 How do candidate mechanisms perform?

The examination of the mechanisms' performance is, in essence, the calibration of their associated parameters, ϕ, θ and τ, in order to obtain an output as close as possible to the observed lending rate series. A parameter with a broader scope, in addition to the rule-specific ones, which also needs calibration is DTI^{max}. This parameter controls the highest level of a household's leverage up to which the bank is willing to extend credit.

The calibration process consists of a grid-search over each interest rate rule's respective parameter space. The chosen ranges for each parameter lie within reasonable bounds, based on the discussion in subsection 2.3.4, while keeping a balance between computational cost and an adequately fine space. Table 3 reports the intervals within which the parameters vary across simulations.

Table 3: Lending rate mechanisms' parameter space.

Parameter	Min	Max	Step	Equation #
DTI^{max}	2	3	0.2	3
ϕ	1.2	2.4	0.2	9
θ	1.5	5	0.5	10
τ	0.1	0.6	0.1	11

For every rule all possible parameter configurations, according to the ranges shown in Table 3 are tested. This translates into 42 distinct pairs of $\{DTI^{max}, \phi\}$ for mechanism $\mathcal{M}_{\phi}^{DTI}$, 56 pairs of $\{\theta, \phi\}$ for $\mathcal{M}_{\phi}^{\theta}$ and another 42 pairs of $\{\tau, \phi\}$ for the third mechanism, $\mathcal{M}_{\phi}^{\tau}$. For the implementation of both lending rate mechanisms which are related to the bank's capital adequacy, DTI^{max} is fixed to the figures identified in the analysis related to $\mathcal{M}_{\phi}^{DTI}$. Practically, these values are sufficiently high and rarely reached in the simulations given the mechanics of the model, while they result in DTIs which fall within plausible ranges considering relevant historical data.[27]

Every rule configuration is applied to the 3 country scenarios and the simulation is executed 100 times. The average of each model-generated interest rate is estimated and compared to the historical series. It should be noted that the comparison period between the artificial and the historical interest rate series differs among scenarios. In particular, for CY and SI it begins at the dates when Cyprus and Slovenia joined the euro area, in January 2008 and 2007 respectively. This is reasonable since the policy rate prior to that period was set by each national central bank and

[26] The aggregate, model-generated interest rate is the average lending rate charged at each month.

[27] The closest data available are from Eurostat and the OECD. Both refer to total debt by households (including non-profit institutions serving households) as a share of their gross or net disposable income, respectively.

differed from the ECB's policy rate which is used in the simulation.[28]. On the contrary, in the UK scenario the validation period starts at the earliest possible date conditional on the availability of historical data.

From the wide range of available methods to assess a model's fit (for an interesting discussion and an approach focusing on the frequency domain see McAdam and Mestre (2008)), the distance and synchronisation between the artificial and real-world series is chosen. The former is quantified with the Mean Absolute Error (MAE) and the latter with Pearson's, zero-lag, cross-correlation coefficient (ρ). Finally, from each rule "family" (i.e. the borrower risk only and the 2 capital adequacy related rules) the best performing configurations are selected and compared to each other. This "horse race" reveals which rule "family" best approximates each banking system's interest rate setting mechanism as well as its particular set of parameters.

4.1.1 Borrower risk only rule ($\mathcal{M}_{\phi}^{DTI}$)

The simplest mechanism examined is the one described in Equation 9 which, in addition to the policy rate, links the bank's offered lending rate only with the borrower's risk. Figure 3 presents the results from the application of the borrower risk only rule in each scenario. The rows show the results for each country scenario. Column-wise, on the left is the distance measure and on the right the correlation between the simulated and the observed lending rates.

[28] A similar study could be performed for the period prior to the countries' adoption of the euro. However, lack of data on unsecured consumer credit and the associated lending rates prohibit the validation of the results from such analysis.

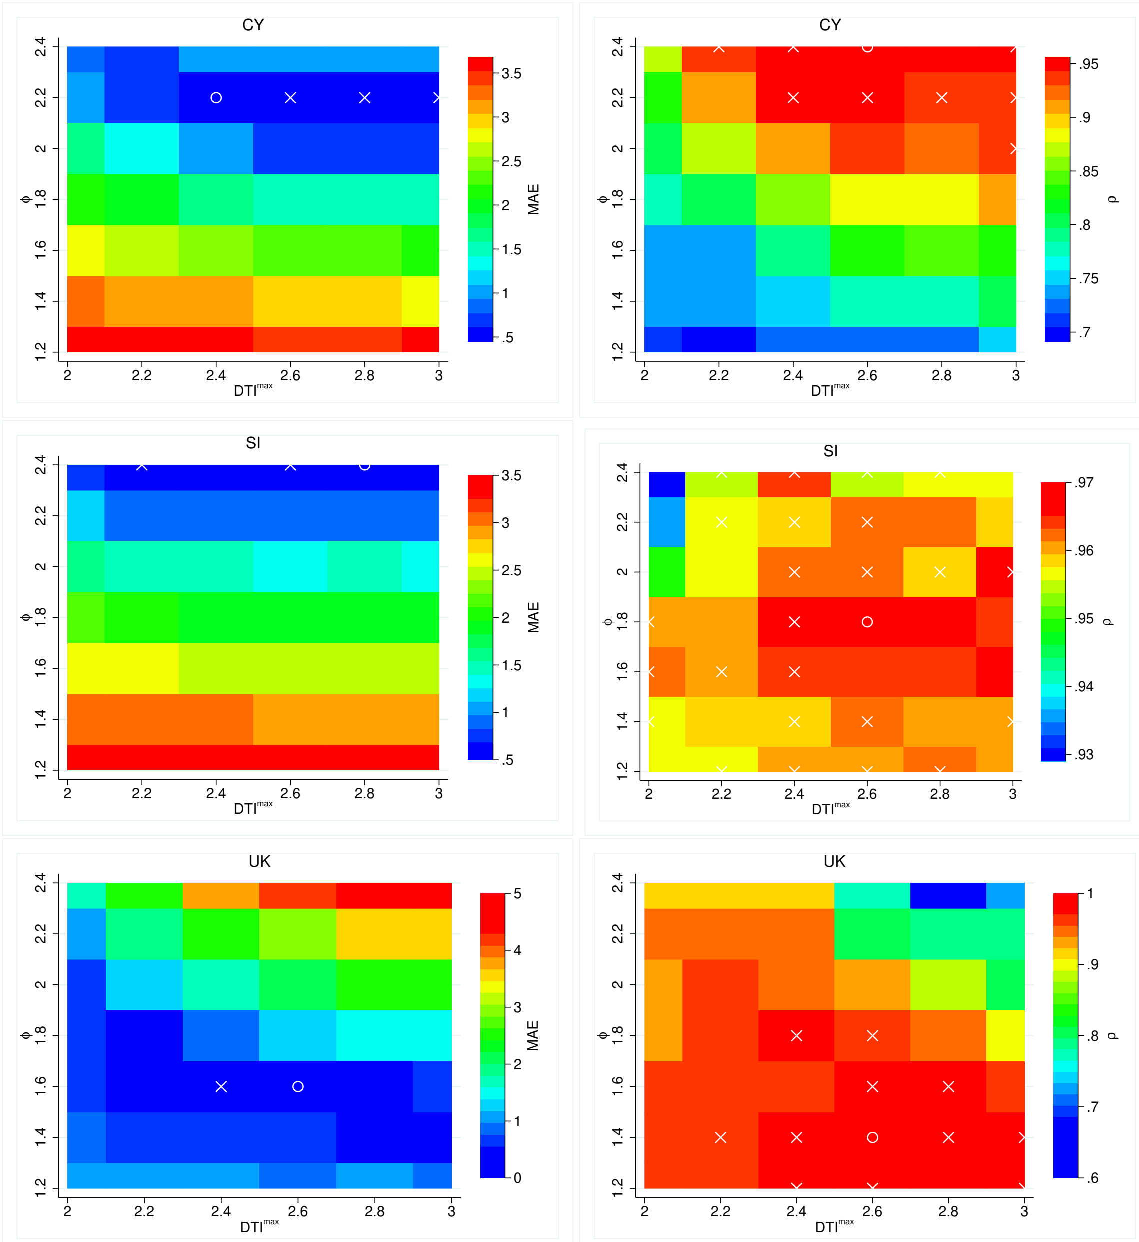

Figure 3: Distance and synchronisation measures between the simulated and historical consumer credit lending rates under the $\mathcal{M}_{\phi}^{DTI}$ mechanism. Circles denote the best performing configurations and Xs those with statistically equal performance to the former at 95% confidence level.

A notable pattern in the right panels of Figure 3 is the high level of synchronisation between the model-generated and the historical series. For most parameter pairs, ρ is above 0.8 in every country scenario. This is particularly evident in SI where the minimum ρ is around 0.93. Moreover, the analysis indicates that many configurations have, statistically, the same correlation with the best performing one. However, MAE figures in the left panels show a different picture.

In both CY and SI scenarios the best performing configurations for this rule are concentrated in the upper regions of the plot. In the UK scenario this pattern reverses and the lowest MAEs are exhibited by parameter pairs at the bottom of the plot. In general, the strongest improvement in performance is observed along the vertical axis, whereas MAE's decrease along the horizontal axis is much more weak. In particular, for CY and SI the rule's performance increases for higher ϕ and DTI^{max} values, while the opposite is true for the UK scenario.

Since correlation is adequately high for virtually every pair of $\{DTI^{max}, \phi\}$, the best performing candidates are selected based on their MAEs. Testing their absolute errors reveals that a few configurations have statistically equal MAEs with the best performing one.[29]

4.1.2 Positively related to capital adequacy rule ($\mathcal{M}_{\phi}^{\theta}$)

Based on the previous analysis, DTI^{max} is fixed for the implementation of both mechanisms related to the bank's capital adequacy. The values for the CY, SI and the UK scenarios are set to $DTI^{max} = 2.6$, $DTI^{max} = 2.8$ and $DTI^{max} = 2.6$ respectively. These figures are reasonably high and are among the best performing ones.

The mechanism $\mathcal{M}_{\phi}^{\theta}$ adds an extra element to the borrower-risk only rule which is a concave link with the bank's capital adequacy as described in Equation 10. The implementation of this rule yields the results shown in Figure 4.

[29] Given the large sample sizes (>130 observations), the test employed is a t-test with unequal variances using Welch's approximation and a 95% confidence level. Wilcoxon's non-parametric test yields similar results.

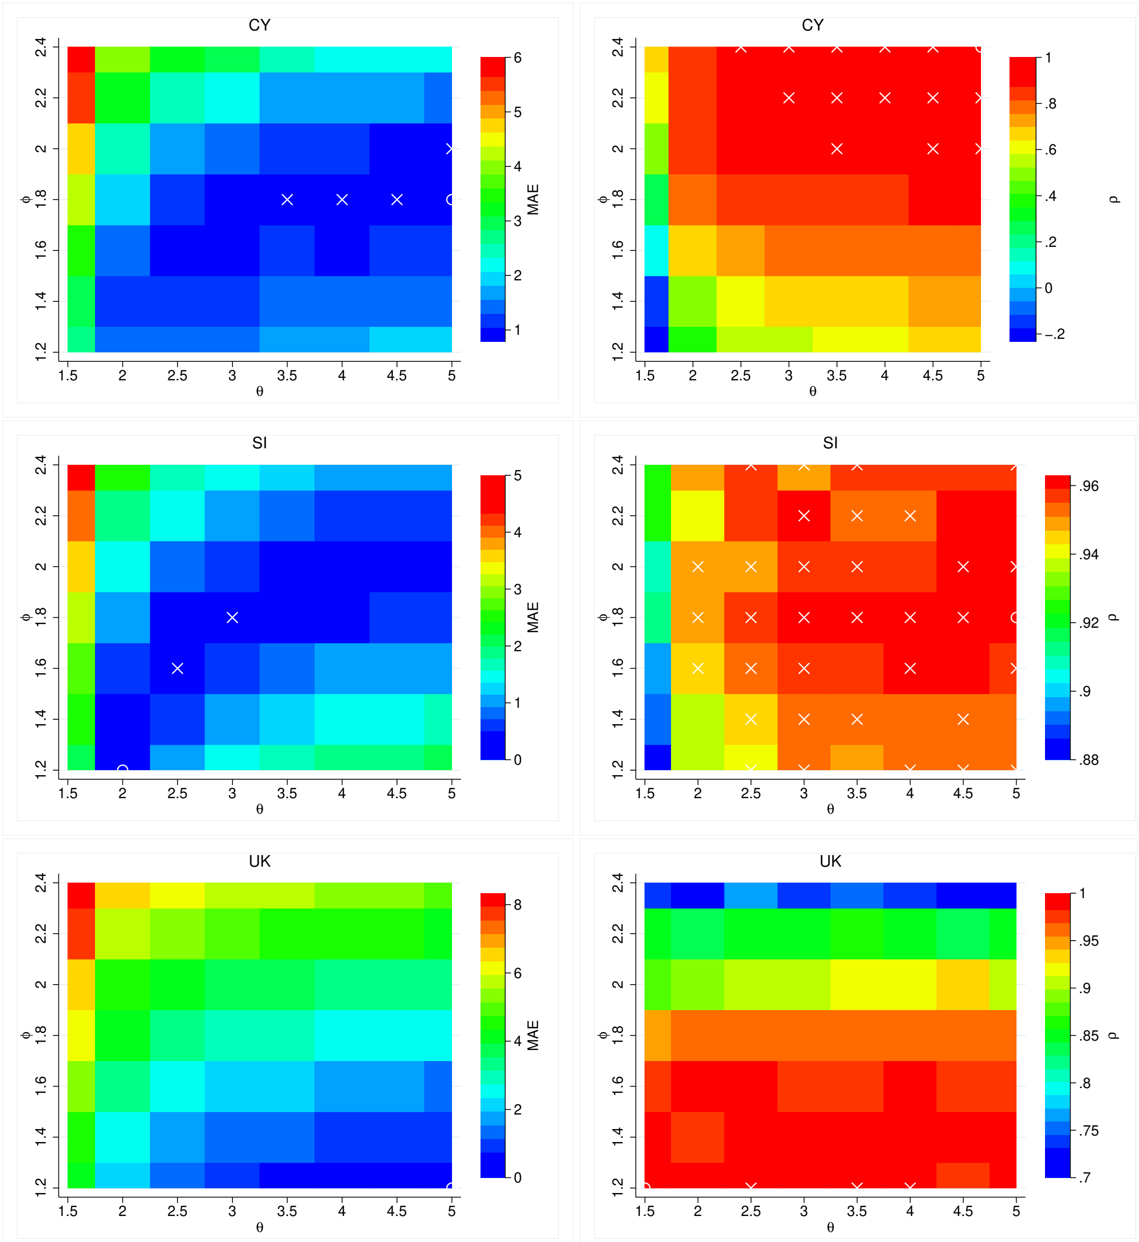

Figure 4: Distance and synchronisation measures between the simulated and historical consumer credit lending rates under the $\mathcal{M}_{\phi}^{\theta}$ mechanism. Circles denote the best performing configurations and Xs those with statistically equal performance to the former at 95% confidence level.

The results in the right panel of Figure 4 show that the output from the $\mathcal{M}_{\phi}^{\theta}$ mechanism broadly shares the same high degree of synchronisation with the observed lending rates as the

$\mathcal{M}_{\phi}^{DTI}$ mechanism.

In terms of distance though, the best-performance region in the left panels exhibits a wide "corridor" structure. The width of the corridor is possibly associated with the relatively large size of the parameter space's sampling step. The improvement path's direction is from the bottom-left to the top-right of the plot, i.e. towards both high ϕ and high θ values, a shape more apparent in the CY and SI scenarios. This is an expected pattern; higher values of θ imply a lower contribution from the capital adequacy element of the rule in the overall interest rate. Therefore, the remaining part is covered by the borrower risk related element, i.e. higher values for ϕ.

Interestingly, the configurations with the lowest MAE are located in different regions for every scenario. In CY they are found in the medium-high range of θ and medium levels of ϕ. On the contrary, in SI they are in the low end of both θ and ϕ. This indicates that, all else equal, the capital adequacy component is contributing more in the latter compared to borrower risk. Another alternative is that in the SI scenario borrowers have higher risk and/or the banking system lower capital adequacy than the CY one. In fact, as seen in Figure 2, CAR_t in CY is substantially higher -especially after 2012- compared to the rest scenarios, thus higher figures for θ (denoting lower contribution) should be probably expected. Finally, in the case of the UK scenario the best performing parameter combinations have low ϕs and high θs indicating a minimal contribution from both elements of the mechanism. This is most likely related to the exceptionally low levels of the historical lending rates. In the top-left panel of Figure 2, UK lending rates are the lowest hovering around 3% after 2010. Thus, implementing this mechanism requires very low contributions from both of its parts in order for the artificial lending rates to remain close to the observed figures.

4.1.3 Negatively related to capital adequacy rule ($\mathcal{M}_{\phi}^{\tau}$)

The final interest rate setting mechanism is described in Equation 11 and links the lending rate and the bank's capital adequacy with a monotonically decreasing function. Applying $\mathcal{M}_{\phi}^{\tau}$ on each scenario yields the results presented in Figure 5.

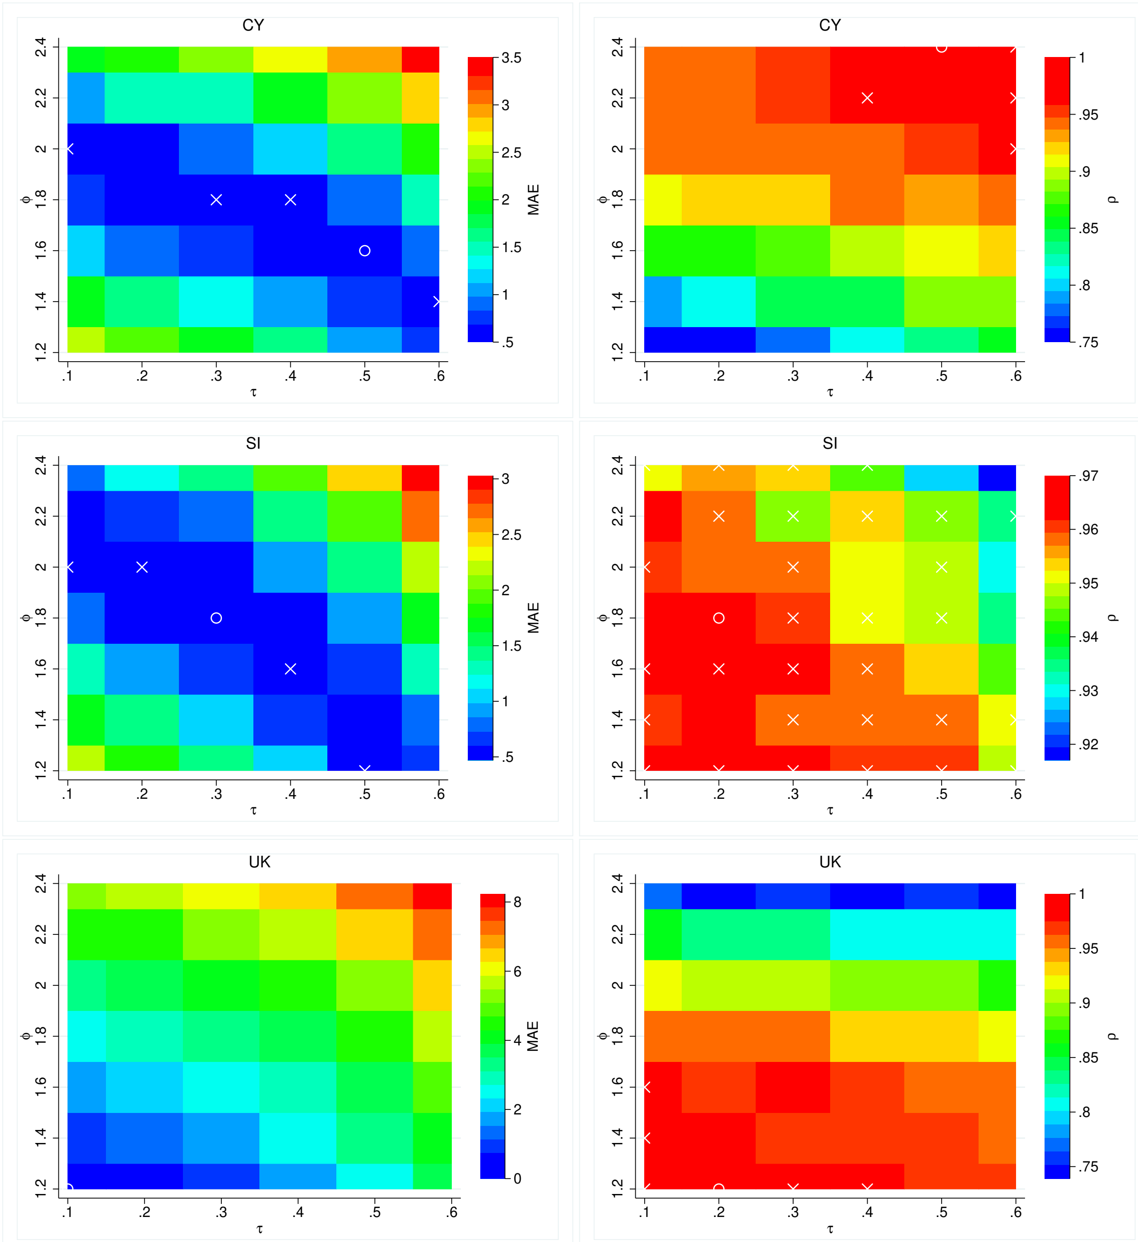

Figure 5: Distance and synchronisation measures between the simulated and historical consumer credit lending rates under the $\mathcal{M}_{\phi}^{\tau}$ mechanism. Circles denote the best performing configurations and Xs those with statistically equal performance to the former at 95% confidence level.

Like the previous cases, the simulated lending rates are very strongly correlated with the real-world series as the results in the right panels in Figure 5 show. The overall minimum ρ across all country scenarios is about 0.75 with the maximum values being above 0.97.

Distance-wise, the results from the implementation of $\mathcal{M}_{\phi}^{\tau}$ also exhibit a "corridor" structure. The corridor is narrower compared to $\mathcal{M}_{\phi}^{\theta}$'s possibly due to the finer sampling step. As expected, its direction is from high-ϕ/low-τ parameter combinations to low-ϕ/high-τ ones.

The statistical tests reveal that, in terms of lowest MAEs, the equally performing configurations are spread across each corridor's path in the cases of CY and SI. On the contrary, in the UK scenario the best performing parameter pairs are located in the bottom-left corner of the plot and none has statistically equal MAE to the best one. This particular characteristic is similar to the pattern observed for the same scenario in the $\mathcal{M}_{\phi}^{\theta}$ mechanism and has likely the same underlying cause; the low levels of the historical interest rates. In order for the simulated series to stay close to the real-world ones, it is necessary to keep the contributions from each mechanism's constituent parts to a minimum, translated to low values for ϕ and low(high) values for $\theta(\tau)$.

4.1.4 Identifying the best performing mechanism

The previous analysis has identified several equally performing configurations within each scenario and mechanism family, exhibiting low distance and high correlation with the historical lending rate series.

To determine which mechanism better approximates the real-world data, the best performing configurations from the 3 mechanism families are compared to each other for every scenario. Table 4 presents the distance and synchronisation mesaures of the best performing configurations per mechanism family for CY, Table 5 for SI and Table 6 for the UK scenario.

Table 4: Distance and synchronisation measures of the best performing configurations per mechanism family in the CY scenario.

id #	Mechanism	MAE	ρ
1	$\mathcal{M}_{\phi=2.2}^{DTI=2.4}$	0.446	0.931
2	$\mathcal{M}_{\phi=2.2}^{DTI=2.6}$	0.454	0.938
3	$\mathcal{M}_{\phi=2.2}^{DTI=2.8}$	0.475	0.934
4	$\mathcal{M}_{\phi=2.2}^{DTI=3.0}$	0.462	0.931
5	$\mathcal{M}_{\phi=1.8}^{\theta=3.5}$	0.821	0.844
6	$\mathcal{M}_{\phi=1.8}^{\theta=4.0}$	0.810	0.843
7	$\mathcal{M}_{\phi=1.8}^{\theta=4.5}$	0.787	0.867
8	$\mathcal{M}_{\phi=1.8}^{\theta=5.0}$	0.783	0.865
9	$\mathcal{M}_{\phi=2.0}^{\theta=5.0}$	0.822	0.901
10	$\mathcal{M}_{\phi=2.0}^{\tau=0.1}$	0.564	0.928
11	$\mathcal{M}_{\phi=1.8}^{\tau=0.3}$	0.542	0.916
12	$\mathcal{M}_{\phi=1.8}^{\tau=0.4}$	0.558	0.930
13	$\mathcal{M}_{\phi=1.6}^{\tau=0.5}$	0.539	0.907
14	$\mathcal{M}_{\phi=1.4}^{\tau=0.6}$	0.597	0.877

The results in Table 4 reveal that the mechanism family with the lowest performance relative to the rest is $\mathcal{M}_{\phi}^{\theta}$ which exhibits the highest distance and lowest synchronisation with the historical interest rate series. The remaining two have a largely similar performances with ρs generally above 0.9 and low MAEs.

Table 5: Distance and synchronisation measures of the best performing configurations per mechanism family in the SI scenario.

id #	Mechanism	MAE	ρ
1	$\mathcal{M}_{\phi=2.4}^{DTI=2.2}$	0.582	0.958
2	$\mathcal{M}_{\phi=2.4}^{DTI=2.6}$	0.572	0.959
3	$\mathcal{M}_{\phi=2.4}^{DTI=2.8}$	0.538	0.960
4	$\mathcal{M}_{\phi=1.2}^{\theta=2}$	0.286	0.938
5	$\mathcal{M}_{\phi=1.6}^{\theta=2.5}$	0.297	0.952
6	$\mathcal{M}_{\phi=1.8}^{\theta=3}$	0.309	0.959
7	$\mathcal{M}_{\phi=2}^{\tau=.1}$	0.507	0.959
8	$\mathcal{M}_{\phi=2}^{\tau=.2}$	0.475	0.955
9	$\mathcal{M}_{\phi=1.8}^{\tau=.3}$	0.469	0.960
10	$\mathcal{M}_{\phi=1.6}^{\tau=.4}$	0.474	0.957
11	$\mathcal{M}_{\phi=1.2}^{\tau=.5}$	0.512	0.958

For the SI scenario, the reported figures in Table 5 indicate that $\mathcal{M}_{\phi}^{\theta}$'s configurations (ids 4 to 6) clearly outperform the rest in terms of MAE. On the contrary, candidate mechanisms are virtually indistinguishable from each other as regards their correlation with the observed lending rates.

Table 6: Distance and synchronisation measures of the best performing configurations per mechanism family in the UK scenario.

id #	Mechanism	MAE	ρ
1	$\mathcal{M}_{\phi=1.6}^{DTI=2.4}$	0.378	0.962
2	$\mathcal{M}_{\phi=1.6}^{DTI=2.6}$	0.342	0.972
3	$\mathcal{M}_{\phi=1.2}^{\theta=5}$	0.328	0.982
4	$\mathcal{M}_{\phi=1.2}^{\tau=.1}$	0.340	0.976

The results for the UK scenario in Table 6 are the least populated ones. The only mechanism represented by more than one candidate is the $\mathcal{M}_{\phi}^{DTI}$. In this case all four configurations are close to each other both in terms of distance as well as of synchonisation with the real-world series.

Figure 6 displays the output of the statistical comparison among the different mechanism configurations and the evolution of the simulated lending rate series.[30]

[30]Similarly to the previous analysis, the means equality test employed is a t-test with unequal variances using

Welch's approximation and a 95% confidence level. Wilcoxon's non-parametric test yields largely the same results.

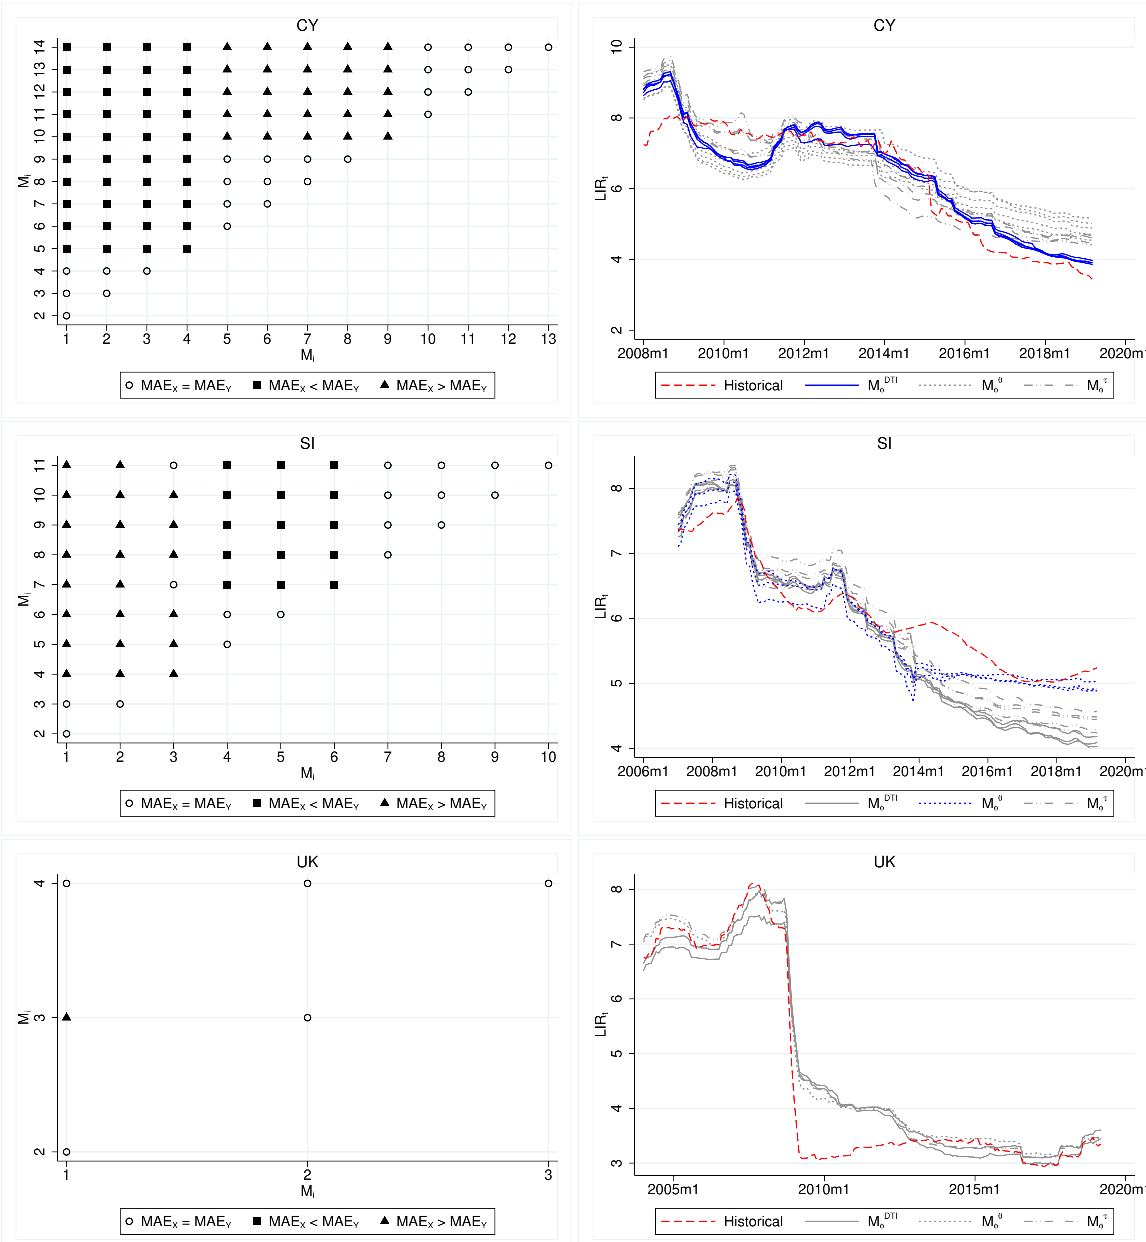

Figure 6: Comparison among different mechanism configurations and the evolution of the simulated lending rate series. Hollow circles (○) denote pairs with statistically equal MAEs, filled squares (■) pairs of which the mechanism in the horizontal axis performs better and filled triangles (▲) pairs of which the mechanism in the vertical axis outperforms its counterpart. Blue lines mark the dominant mechanism's configurations.

The left panels in Figure 6 present the results of the comparison among different mechanism configurations. The 3 mechanism families are clearly recognized in the respective regions of hollow circles near the diagonal within which the different configurations exhibit the same MAEs. However, different mechanisms exhibit the lowest, overall, distance from the historical interest rates for each country scenario. In the case of CY the configurations of the simple borrower-risk rule outperform those of both capital adequacy related ones, whereas in that of SI the best performing configurations belong to the $\mathcal{M}_{\phi}^{\theta}$ mechanism. In line with the results of Table 6, there is no clear winner in the UK scenario. All four configurations seem to have the same performance with the exception of $\mathcal{M}_{\phi=1.2}^{\theta=5}$ which shows a better fit when compared to $\mathcal{M}_{\phi=1.6}^{DTI=2.4}$.

The evolution of the simulated lending rates on the right panels shows the output from the best performing mechanisms along with the observed series (red, dashed line) for each scenario. For CY, $\mathcal{M}_{\phi}^{DTI}$'s configurations remain remarkably close to the historical data for most of the period under study. Interestingly, from early 2009 until 2011 all mechanisms but $\mathcal{M}_{\phi}^{\tau}$ under-predict the lending rate. During that period $\mathcal{M}_{\phi}^{\tau}$'s configurations (grey, dash-dotted lines) replicate the level and the dynamics of the historical lending rates more closely than their counterparts. In SI, virtually all configurations replicate adequately the levels and the dynamics of the observed interest rate series roughly until 2013. For the next 3 years no mechanism exhibits a good match up to 2016 when $\mathcal{M}_{\phi}^{\theta}$ seems to perform better -though not perfectly- than the rest. This result might be associated with the economic and financial turmoil in the country during that period and the related bank recapitalisations that took place in late 2013 and 2014 (Bank of Slovenia, 2015). These conditions might have caused a change in the banks' behaviour which is not properly captured by the rules or parameter ranges examined in this study. Finally, in the case of the UK, no single mechanism stands apart from the rest. All exhibit a similarly good fit in the period before October 2008 and after January 2013, but fail to follow the sharp drop in late 2008 and the dynamics in the subsequent four-year period.

4.2 Empirical validation

Guided by the previous analysis, an individual configuration from the best performing interest rate setting mechanism is implemented in each scenario and the model's output is compared to the historical series of three key variables; household consumption, consumer credit and the lending rate. The specific configurations chosen are $\mathcal{M}_{\phi=2.2}^{DTI=2.6}$ for CY, $\mathcal{M}_{\phi=1.6}^{\theta=2.5}$ for SI and $\mathcal{M}_{\phi=1.6}^{DTI=2.6}$ for the UK scenario. While in the first two country scenarios the statistical tests reveal a dominant mechanism family, this is not the case for the UK one. Therefore, the choice is based on the mechanisms' simplicity favouring $\mathcal{M}_{\phi}^{DTI}$ over the capital adequacy related ones.

Each scenario is executed 100 times and the 100-run averages of the simulated series are contrasted to the observed data. In order to facilitate comparison, two transformations are considered; rescaling based on the full-sample average (y_t/y^{mean}, where $y_t = \{C_t, L_t\}$) and annual growth rates.[31] As before, distance from the real-world series is quantified with MAE, while synchronisation with the the zero-lag, Pearson's correlation coefficient.

4.2.1 Lending rate validation

Since the focus of this study is the identification of the interest rate setting mechanism which better approximates the observed data, it is natural to begin with the validation of the lending rate series.

[31] $\%\Delta_{yoy}(y_t)$ where $y_t = \{C_t, L_t\}$ and $\Delta_{yoy}(y_t)$ for $y_t = \{LIR_t\}$.

Although it has already been studied in subsection 4.1, the examination of some additional elements such as growth rates and simulation ranges provides further information. On Figure 7 the panels on the left present the results for the level series, while the ones on the right for those in growth rates.

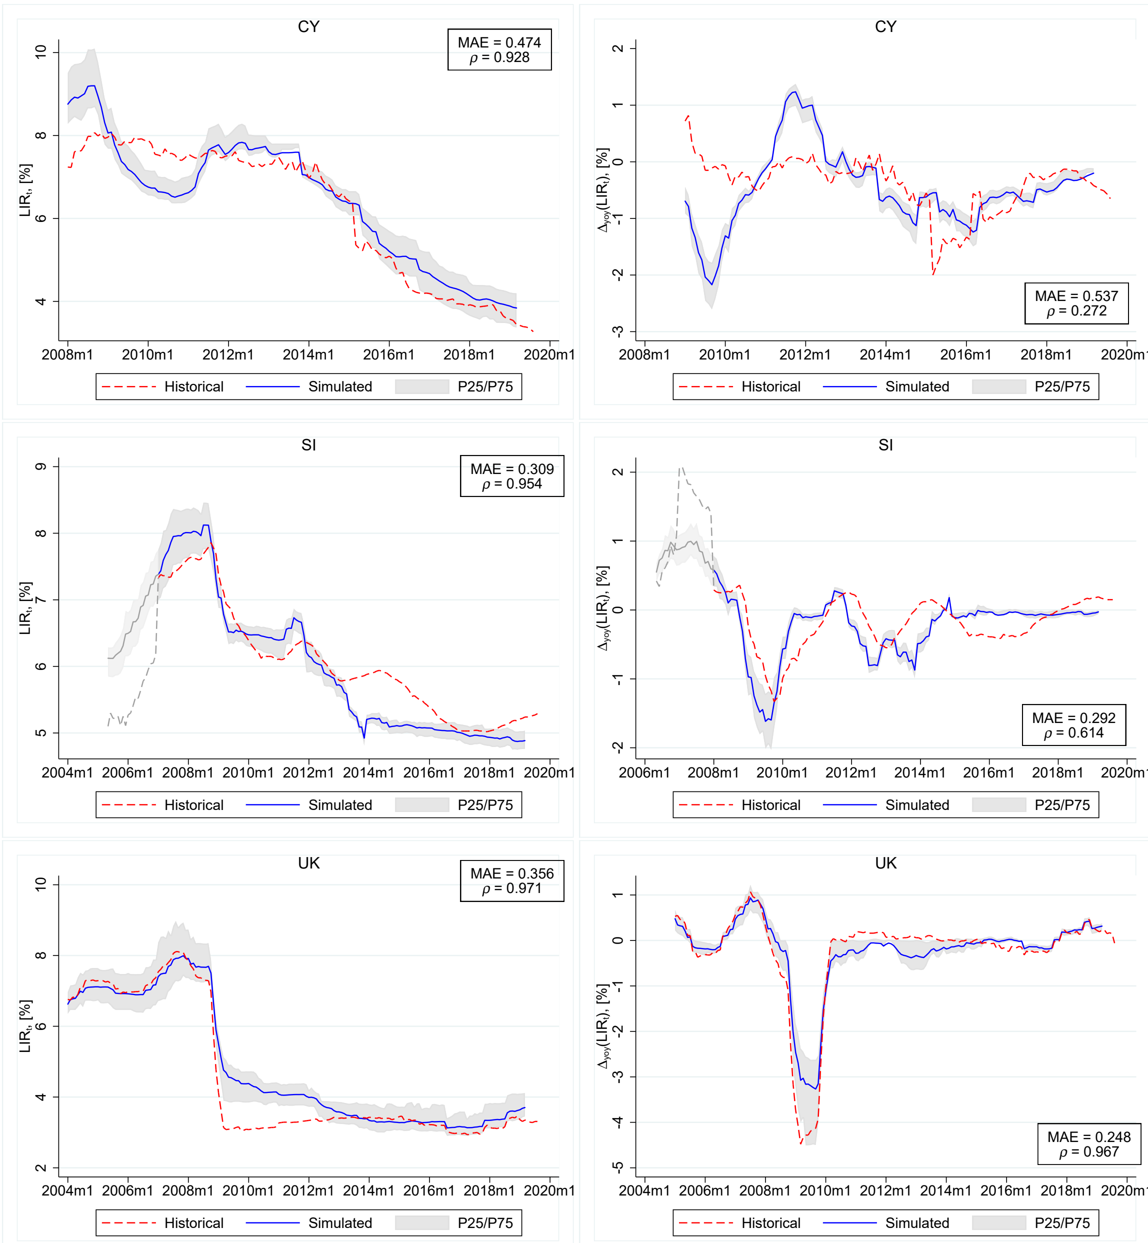

Figure 7: Model validation results for lending rates.

The plots in Figure 7 show a remarkably good fit between the model-generated and real-world lending rates. In levels, overall MAEs are lower than half a percentage point and correlations well above 0.9. In growth rates, distance ranges between about 2 p.p. and 5 p.p. while synchronisation varies but mostly stays at high levels.

An interesting result is the discrepancy between the artificial and historical series in the CY and SI scenarios during a specific part of the simulation period, roughly divided by the GFC and its aftermath. In each scenario, the implemented mechanism exhibits a good fit either post- (in CY) or pre-crisis (in SI), but does not perform as good in the other part. This could be due to changes in credit demand as embedded in Equation 8 or the interest rate setting mechanism itself. That being said, for the case of CY, the results in Figure 6 indicate that for some part of the pre-crisis period[32] some configurations from the $\mathcal{M}_{\phi}^{\tau}$ rule outperform their counterparts, including the selected one ($\mathcal{M}_{\phi=2.2}^{DTI=2.6}$). In the case of SI the analysis didn't identify an alternative, better performing configuration during the post-crisis period. Therefore, the results leave open the possibility of a different mechanism (function- or parameter-wise) being at work then.[33]

It should be noted that the reported distance and synchronisation figures in the boxes are estimated based on data availability (UK after 2004) and appropriateness (CY & SI after euro adoption). These are marked with light grey areas and lines and are not considered in the estimation of MAE and ρ.

4.2.2 Consumer credit validation

The next model-generated series compared to the observed ones are the volumes and growth rates of consumer credit. Figure 8 displays the evolution of the respective series along with the goodness-of-fit measures.

[32]Classifying the build-up period until the bailout programme for Cyprus between the Cypriot government and the European Commission, the ECB and the International Monetary Fund in March 2013 as "pre-crisis".

[33]Another possibility is that the 5-year, fixed-maturity assumption is no longer an adequate approximation of the real-world maturities to which the historical lending rates correspond.

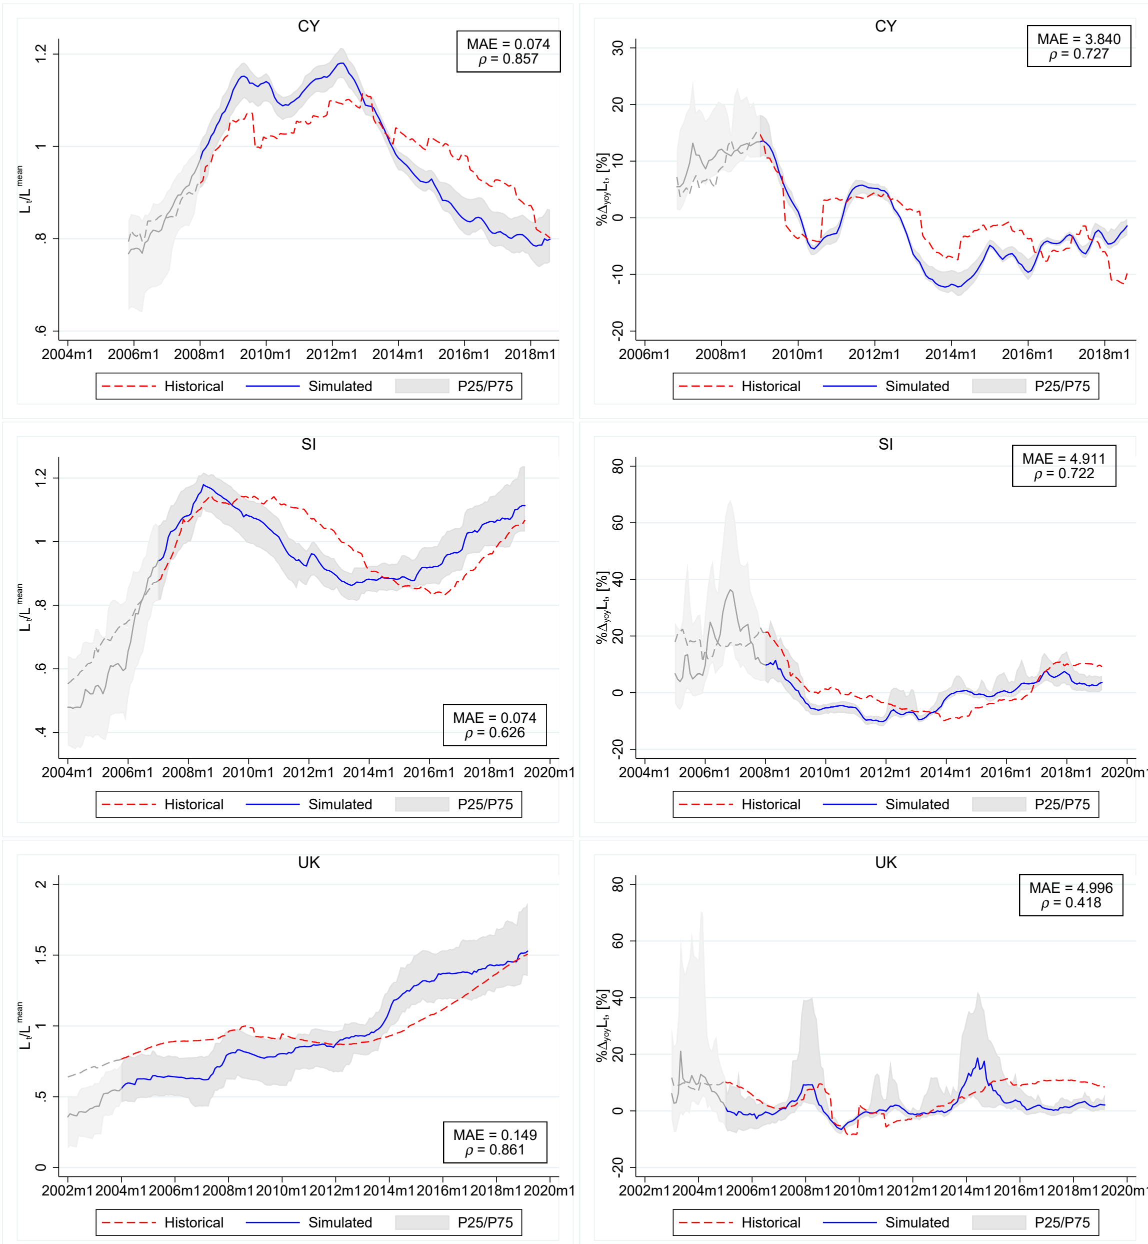

Figure 8: Model validation results for consumer credit.

Given that the evolution of credit is tied to the associated interest rates, it follows that validation is meaningful during the same period that the latter have been also validated. Thus, the light grey areas and lines are not considered for the estimation of the goodness-of-fit measures.[34] In the case

[34]Assuming that the implemented mechanisms apply for the whole simulation horizon and using all available

of CY an extra treatment has been applied on the data sample; observations after September 2018 have been discarded in order to eliminate a structural break in the real-world series which distorted the comparison.[35]

The results in Figure 8 are very promising. Both simulated levels and growth rates show a high level of synchronisation with the historical series with ρs spanning between 0.6 and above 0.85 in most cases. The only exception is the case of growth rates in the UK scenario where towards the end of the simulation horizon the observed consumer credit grows at a higher pace than the artificial one. Regarding distance, MAE figures indicate that model-generated and real-world data differ by 4 p.p. to 5 p.p. in absolute terms, on average. It should be noted that a key series affecting consumer credit is DSTI. Lacking accurate data on its level and dynamics introduces changes in the simulated series and makes model validation more challenging. Nevertheless, despite this limitation the model's output remains adequately close to reality.

4.2.3 Household consumption validation

Finally, Figure 9 shows the validation results for household consumption for each country scenario.

observations for consumer credit improves some results while leaves others roughly unchanged.

[35] In particular, a sharp shift of around 40% in month-on-month terms, possibly due to some structural or methodological reason.

Figure 9: Model validation results for household consumption.

A noticeable pattern across all scenarios, is the extremely narrow interquartile range of the simulated series. Marked by the grey area around the average, it is barely visible. The stochastic element controlling the width of the interquartile range is debt assumption by households and its subsequent spending for consumer needs. Therefore, this pattern indicates that consumer credit in the model is a small fraction of household consumption and therefore has a small effect on it.[36] By extension, this implies that the effect of interest rate changes on consumption will be modest at best, a result in line with other classes of macroeconomic models (McAdam and Morgan, 2003).

As seen in Figure 9 the simulated series reproduce in a satisfactory manner the dynamics of the real-world ones both in levels and in annual growth rates. This is particularly evident in SI and UK where they exhibit low figures of MAEs and high correlations. The results for CY display some discrepancies, especially after 2009, but they show an acceptable fit. The latter could be due to some part of the model not being a sufficient approximation of real-world processes after 2009; either the underlying consumption function or the implemented income scenario could differ from what has actually happened, leading to divergent results. However, overall, the model provides a decent replication of the evolution of household consumption.

5 Conclusions

This study has examined the capacity of three different mechanisms, commonly used by the ABM literature, to approximate the historical series of consumer credit interest rates. The mechanisms were implemented within a data-driven agent-based model and their performance was studied in three country-specific scenarios.

Naturally, the primary result is the identification of the best performing mechanism and its distinct parameter values. Nonetheless, the analysis has uncovered three more generic patterns which extend beyond the determination of the specific mechanism per country. First, as expected, there is heterogeneity across country scenarios regarding both the mechanism family and its particular parameter figures that better approximate the observed data. However, it seems that, in two cases, the simple candidates outperformed their more complicated counterparts. More precisely, in CY, the borrower-risk only rule exhibited a better fit than any of the capital adequacy related rules. In the UK scenario, candidate configurations from all three mechanism families performed equally well. Therefore, given its lower complexity, the borrower-risk only rule is considered to be the dominant one. Finally, an interesting result is the possible existence of a dynamic behaviour on the bank's side, likely influenced by the prevailing economic conditions. In two scenarios (CY and SI), the best-performing mechanisms were largely identified according to their performance on a certain part of the sample. In both cases the GFC and its repercussions seemed to be the catalytic event, differentiating behaviour in the sub-periods determined by its occurrence.

From a policy-simulation perspective, this model can provide valuable, country-specific insights regarding the interest rate transmission mechanism and the effect of various policies on lending rates and credit growth. It allows the implementation of alternative monetary policy scenarios or the introduction of macroprudential regulations and provides a quantitative assessment of their impact on the variables under focus.

Further investigation of the lending rate mechanism's stability over time is one of the many ways this model can be used and extended. Another extension could be the examination of the bank's liability side; the evolution of deposits and their associated interest rates. Finally, a more challenging research avenue would be to gradually replace every historical input series with endogenous rules by modelling each part of the economy in greater detail (e.g. housing market, labour market, private/public sector etc).

[36] Thus, using observations beyond the validation period of the lending rate and consumer credit series does not have a material impact on the results.

More broadly, this study has shown that using a data-driven ABM in which some parts of the economy follow a specific historical scenario, instead of explicitly modelling every agent, enables the focused investigation of specific parts of agents' behaviour and the calibration of their parameters. It thus represents a small step towards the development of a fully-fledged macroeconomic ABM which is "closer to the data" than the current standard in the literature.

References

Alexandre, M. and Lima, G. T. (2020). Combining monetary policy and prudential regulation: An agent-based modeling approach. *Journal of Economic Interaction and Coordination*, 15:385–411.

Allen, H. and Taylor, M. P. (1990). Charts, noise and fundamentals in the London foreign exchange market. *The Economic Journal*, 100(400):49–59.

Almarzoqi, R. and Naceur, S. B. (2015). Determinants of bank interest margins in the Caucasus and Central Asia. *IMF Working Paper No.15/87*, 15(87).

Ampudia, M. and Ehrmann, M. (2017). Macroeconomic experiences and risk taking of euro area households. *European Economic Review*, 91:146–156.

Anufriev, M. and Hommes, C. (2012a). Evolution of market heuristics. *The Knowledge Engineering Review*, 27(2):255–271.

Anufriev, M. and Hommes, C. (2012b). Evolutionary selection of individual expectations and aggregate outcomes in asset pricing experiments. *American Economic Journal: Microeconomics*, 4(4):35–64.

Arifovic, J. and Sargent, T. J. (2003). Laboratory experiments with an expectational Phillips curve. In *Evolution and Procedures in Central Banking*, pages 23–55. Cambridge University Press.

Ashraf, Q., Gershman, B., and Howitt, P. (2016). How inflation affects macroeconomic performance: An agent-based computational investigation. *Macroeconomic Dynamics*, 20:558–581.

Ashraf, Q., Gershman, B., and Howitt, P. (2017). Banks, market organization, and macroeconomic performance: An agent-based computational analysis. *Journal of Economic Behavior & Organization*, 135:143–180.

Assenza, T., Bao, T., Hommes, C., and Massaro, D. (2014). Experiments on expectations in macroeconomics and finance. In *Experiments in Macroeconomics, Research in Experimental Economics*, volume 17, pages 11–70. Emerald Group Publishing Limited.

Assenza, T., Delli Gatti, D., and Grazzini, J. (2015). Emergent dynamics of a macroeconomic agent based model with capital and credit. *Journal of Economic Dynamics and Control*, 50:5–28.

Assenza, T., Heemeijer, P., Hommes, C. H., and Massaro, D. (2019). Managing self-organization of expectations through monetary policy: A macro experiment. *Journal of Monetary Economics*.

Bandourian, R., McDonald, J., and Turley, R. (2002). A comparison of parametric models of income distribution across countries and over time. *LIS Working Paper Series, No. 305*.

Bank of Slovenia (2015). Report of the Bank of Slovenia for the National Assembly. Technical report, Bank of Slovenia.

Baugnet, V. and Hradisky, M. (2004). Determinants of Belgian bank lending interest rates. Technical report, National Bank of Belgium.

Boissay, F., Cantú, C., Claessens, S., and Villegas, A. (2019). Impact of financial regulations: Insights from an online repository of studies. *BIS Quarterly Review*.

Botta, A., Caverzasi, E., Russo, A., Gallegati, M., and Stiglitz, J. E. (2019). Inequality and finance in a rent economy. *Journal of Economic Behavior & Organization*.

Branch, W. A. (2004). The theory of rationally heterogeneous expectations: Evidence from survey data on inflation expectations. *The Economic Journal*, 114(497):592–621.

Brock, P. and Franken, H. (2003). Measuring the determinants of average and marginal bank interest rate spreads in Chile, 1994-2001. *Journal Economía Chilena (The Chilean Economy)*, 6(3):45–65.

Caiani, A., Godin, A., Caverzasi, E., Gallegati, M., Kinsella, S., and Stiglitz, J. E. (2016). Agent based-stock flow consistent macroeconomics: Towards a benchmark model. *Journal of Economic Dynamics and Control*, 69:375–408.

Caiani, A., Russo, A., and Gallegati, M. (2019). Does inequality hamper innovation and growth? An AB-SFC analysis. *Journal of Evolutionary Economics*, 29(1):177–228.

Cardaci, A. (2018). Inequality, household debt and financial instability: An agent-based perspective. *Journal of Economic Behavior & Organization*, 149:434–458.

Cardaci, A. and Saraceno, F. (2019). Between Scylla and Charybdis: Income distribution, consumer credit, and business cycles. *Economic Inquiry*, 57(2):953–971.

Central Bank of Cyprus (2013). Directive on loan origination processes and processes of reviewing existing loans of 2013. Technical report, Central Bank of Cyprus.

Cincotti, S., Raberto, M., and Teglio, A. (2012). The EURACE macroeconomic model and simulator. In *Complexity and Institutions: Markets, Norms and Corporations.*, pages 81–106. Palgrave Macmillan, New York.

Claeys, S. and Vander Vennet, R. (2008). Determinants of bank interest margins in Central and Eastern Europe: A comparison with the West. *Economic Systems*, 32(2):197–216.

Cornand, C. and Heinemann, F. (2014). Experiments on monetary policy and central banking. In *Experiments in Macroeconomics, Research in Experimental Economics*, volume 17, pages 167–227. Emerald Group Publishing Ltd.

Dawid, H. and Delli Gatti, D. (2018). Agent-based macroeconomics. *Handbook of Computational Economics*, 4:63–156.

Dawid, H., Harting, P., van der Hoog, S., and Neugart, M. (2019). Macroeconomics with heterogeneous agent models: Fostering transparency, reproducibility and replication. *Journal of Evolutionary Economics*, 29(1):467–538.

Delli Gatti, D. and Desiderio, S. (2015). Monetary policy experiments in an agent-based model with financial frictions. *Journal of Economic Interaction and Coordination*, 10(2):265–286.

Delli Gatti, D., Desiderio, S., Gaffeo, E., Cirillo, P., and Gallegati, M. (2011). *Macroeconomics from the bottom-up*, volume 1. Springer Science & Business Media.

Delli Gatti, D., Gallegati, M., Greenwald, B., Russo, A., and Stiglitz, J. E. (2010). The financial accelerator in an evolving credit network. *Journal of Economic Dynamics and Control*, 34(9):1627–1650.

D'Orazio, P. (2019). Income inequality, consumer debt, and prudential regulation: An agent-based approach to study the emergence of crises and financial instability. *Economic Modelling*, 82:308–331.

D'Orazio, P. and Giulioni, G. (2017). From micro behaviors to macro dynamics: An agent-based economic model with consumer credit. *Journal of Artificial Societies and Social Simulation*, 20(1).

Dosi, G., Fagiolo, G., Napoletano, M., and Roventini, A. (2013). Income distribution, credit and fiscal policies in an agent-based Keynesian model. *Journal of Economic Dynamics and Control*, 37(8):1598–1625.

Dosi, G., Fagiolo, G., Napoletano, M., Roventini, A., and Treibich, T. (2015). Fiscal and monetary policies in complex evolving economies. *Journal of Economic Dynamics and Control*, 52:166–189.

Drehmann, M., Illes, A., Juselius, M., and Santos, M. (2015). How much income is used for debt payments? A new database for debt service ratios. *BIS Quarterly Review*, page 89.

Ehrmann, M. and Tzamourani, P. (2012). Memories of high inflation. *European Journal of Political Economy*, 28(2):174–191.

European Systemic Risk Board (2016). A review of macroprudential policy in the EU in 2015. Technical report, European Systemic Risk Board.

European Systemic Risk Board (2017). A review of macroprudential policy in the EU in 2016. Technical report, European Systemic Risk Board.

European Systemic Risk Board (2018). A review of macroprudential policy in the EU in 2017. Technical report, European Systemic Risk Board.

European Systemic Risk Board (2019). A review of macroprudential policy in the EU in 2018. Technical report, European Systemic Risk Board.

European Systemic Risk Board (2020). A review of macroprudential policy in the EU in 2019. Technical report, European Systemic Risk Board.

Farmer, J. D. and Foley, D. (2009). The economy needs agent-based modelling. *Nature*, 460(7256):685–686.

Fehr, E. and Tyran, J.-R. (2008). Limited rationality and strategic interaction: The impact of the strategic environment on nominal inertia. *Econometrica*, 76(2):353–394.

Frankel, J. A. and Froot, K. A. (1987). Using survey data to test standard propositions regarding exchange rate expectations. *The American Economic Review*, 77(1):133–153.

Freixas, X. and Rochet, J.-C. (2008). *Microeconomics of banking*. MIT press.

Gaffeo, E., Delli Gatti, D., Desiderio, S., and Gallegati, M. (2008). Adaptive microfoundations for emergent macroeconomics. *Eastern Economic Journal*, 34(4):441–463.

Gambacorta, L. (2008). How do banks set interest rates? *European Economic Review*, 52(5):792 – 819.

Geanakoplos, J., Axtell, R., Farmer, D. J., Howitt, P., Conlee, B., Goldstein, J., Hendrey, M., Palmer, N. M., and Yang, C.-Y. (2012). Getting at systemic risk via an agent-based model of the housing market. *American Economic Review*, 102(3):53–58.

Giri, F., Riccetti, L., Russo, A., and Gallegati, M. (2019). Monetary policy and large crises in a financial accelerator agent-based model. *Journal of Economic Behavior & Organization*, 157:42–58.

Godley, W. and Lavoie, M. (2016). *Monetary economics: An integrated approach to credit, money, income, production and wealth*. Springer.

Gualdi, S., Tarzia, M., Zamponi, F., and Bouchaud, J.-P. (2015). Tipping points in macroeconomic agent-based models. *Journal of Economic Dynamics and Control*, 50:29–61.

Gurgone, A., Iori, G., and Jafarey, S. (2018). The effects of interbank networks on efficiency and stability in a macroeconomic agent-based model. *Journal of Economic Dynamics and Control*, 91:257–288.

Hommes, C., Massaro, D., and Weber, M. (2019). Monetary policy under behavioral expectations: Theory and experiment. *European Economic Review*, 118:193–212.

Hommes, C., Sonnemans, J., Tuinstra, J., and van de Velden, H. (2005). Coordination of expectations in asset pricing experiments. *Review of Financial Studies*, 18(3):955–980.

Hommes, C., Sonnemans, J., Tuinstra, J., and van de Velden, H. (2008). Expectations and bubbles in asset pricing experiments. *Journal of Economic Behavior & Organization*, 67(1):116–133.

Horváth, R. and Podpiera, A. (2012). Heterogeneity in bank pricing policies: The Czech evidence. *Economic Systems*, 36(1):87–108.

Mankiw, N. G., Reis, R., and Wolfers, J. (2003). Disagreement about inflation expectations. *NBER Macroeconomics Annual*, 18:209–248.

Marimon, R. and Sunder, S. (1995). Does a constant money growth rule help stabilize inflation? Experimental evidence. In *Carnegie-Rochester Conference Series on Public Policy*, volume 43, pages 111–156. Elsevier.

McAdam, P. and Mestre, R. (2008). Evaluating macro-economic models in the frequency domain: A note. *Economic Modelling*, 25(6):1137–1143.

McAdam, P. and Morgan, J. (2003). *Analysing monetary policy transmission at the euro area level using structural macroeconomic models*, page 75–90. Cambridge University Press.

Meijers, H., Nomaler, Ö., and Verspagen, B. (2018). Demand, credit and macroeconomic dynamics. A micro simulation model. *Journal of Evolutionary Economics*, pages 1–28.

Modigliani, F. and Brumberg, R. (1954). Utility analysis and the consumption function: An interpretation of cross- section data. *Franco Modigliani*, 1:388–436.

Osborne, M., Fuertes, A.-M., and Milne, A. (2017). In good times and in bad: Bank capital ratios and lending rates. *International Review of Financial Analysis*, 51:102–112.

Palagi, E., Napoletano, M., Roventini, A., and Gaffard, J.-L. (2017). Inequality, redistributive policies and multiplier dynamics in an agent-based model with credit rationing. *Italian Economic Journal*, 3(3):367–387.

Papadopoulos, G. (2019). Income inequality, consumption, credit and credit risk in a data-driven agent-based model. *Journal of Economic Dynamics and Control*, 104:39 – 73.

Pfajfar, D. and Santoro, E. (2010). Heterogeneity learning and information stickiness in inflation expectations. *Journal of Economic Behavior & Organization*, 75(3):426–444.

Platt, D. (2020). A comparison of economic agent-based model calibration methods. *Journal of Economic Dynamics and Control*, pages 103–859.

Poledna, S., Miess, M., and Hommes, C. H. (2020). Economic forecasting with an agent-based model. *IIASA Working Paper*, (20-001).

Popoyan, L., Napoletano, M., and Roventini, A. (2020). Winter is possibly not coming: Mitigating financial instability in an agent-based model with interbank market. *Journal of Economic Dynamics and Control*, page 103937.

Reissl, S. (2020a). Heterogeneous expectations, forecasting behaviour and policy experiments in a hybrid agent-based stock-flow-consistent model. *Journal of Evolutionary Economics*, pages 1–49.

Reissl, S. (2020b). Minsky from the bottom up–formalising the two-price model of investment in a simple agent-based framework. *Journal of Economic Behavior & Organization*, 177:109–142.

Riccetti, L., Russo, A., and Gallegati, M. (2013). Unemployment benefits and financial leverage in an agent based macroeconomic model. *Economics: The Open-Access, Open-Assessment E-Journal*, 7(2013-42):1–44.

Riccetti, L., Russo, A., and Gallegati, M. (2015). An agent based decentralized matching macroeconomic model. *Journal of Economic Interaction and Coordination*, 10(2):305–332.

Russo, A., Riccetti, L., and Gallegati, M. (2016). Increasing inequality, consumer credit and financial fragility in an agent based macroeconomic model. *Journal of Evolutionary Economics*, 26(1):25–47.

Teglio, A., Mazzocchetti, A., Ponta, L., Raberto, M., and Cincotti, S. (2019). Budgetary rigour with stimulus in lean times: Policy advices from an agent-based model. *Journal of Economic Behavior & Organization*, 157:59–83.

Valverde, S. C. and Fernández, F. R. (2007). The determinants of bank margins in European banking. *Journal of Banking & Finance*, 31(7):2043–2063.

A Interest rate setting in the ABM literature

Table 7 presents in a compact way the key aspects of the interest rate setting mechanism as implemented in many, recent macroeconomic ABMs. It is a certainly non-exhaustive review of the literature but provides an overview of the features considered in numerous studies. Regarding the bank-specific component of the lending rate, it should be noted that several models include regulatory constrains which influence the volume of the supplied credit and hence, indirectly affect the interest rate.[37] Nevertheless, since they are not directly implemented in the behavioural rule, the respective column marks them as "*None direct*".

Table 7: Interest rate setting mechanism implementation in the ABM literature.

Reference(s)	**Potential borrowers**	**Debt maturity**	**Collateral**	**Interest rate mechanism**	**Baseline rate**	**Borrower-specific IR component**	**Bank-specific IR component**	**Short description**
Dosi et al. (2013)	Firms	Implicit	No	$r_L = r(1+\psi_L)$	Central bank rate (r)	None	None direct	Interest rate is calculated as a fixed mark-up on the baseline rate
Ashraf et al. (2016, 2017); Popoyan et al. (2020)	Firms	Implicit	Inventory and fixed capital	$r_L = r_w + s/48$	Interest rate on government bonds set by the central bank (r_w)	None	None direct	Interest rate is determined by a fixed spread on the baseline rate
D'Orazio and Giulioni (2017); D'Orazio (2019)	Households	Implicit	No	$r_L = constant$	None	None	None direct	Constant interest rate
Palagi et al. (2017)	Households	Single-period	No	$r^b = r(1+\mu^b)$	Central bank rate (r)	None	None direct	Interest rate is calculated as a constant mark-up on the baseline rate
Poledna et al. (2020)	Firms	Implicit	No	$r_t = \bar{r} + \mu$	Central bank rate ($\bar{r}$)	None	None direct	Interest rate is calculated as a fixed spread on the baseline rate

Continues in next page

[37]Insofar as credit-rationing borrowers implies the absence of an applied interest rate altogether.

Table 7 (continued)

Reference(s)	**Potential borrowers**	**Debt maturity**	**Collateral**	**Interest rate mechanism**	**Baseline rate**	**Borrower-specific IR component**	**Bank-specific IR component**	**Short description**
Gaffeo et al. (2008); Delli Gatti et al. (2011)	Firms	Single-period	No	$r_{i,t} = \bar{r}\big(1 + \phi_{n,t} \cdot \mu(\lambda_{i,t})\big)$	Rate set by the monetary authority ($\bar{r}$)	Borrower i's total debt to net worth ratio ($\lambda_{i,t}$)	None direct	Interest rate is calculated as a mark-up on the baseline rate, increasing with borrowers' leverage
Cincotti et al. (2012); Teglio et al. (2019)	Firms	24	No	$r^{b,f} = r^{CB} + \gamma^b \pi^f$	Central bank rate (r^{CB})	Borrower f's probability of default (π^f), linked to their leverage	None direct	Interest rate is calculated as a mark-up on the baseline rate, linked to borrowers' credit risk
Dosi et al. (2015)	Firms	Implicit	No	$r_{j,t} = r_t^{deb}\big(1 + (q-1)k_{const}\big)$	Mark-up on central bank rate ($r_t^{deb} = (1+\mu)r_t$)	Borrower j's credit ranking (q) based on their past net worth to sales ratio	None direct	Interest rate is calculated as a mark-up on the baseline rate, depending on the borrowers' position in credit ranking
Assenza et al. (2015)	Firms	Implicit	No	$r_{f,t} = \mu\big[\frac{1+r/\theta}{\Xi(\theta, T_{f,t})} - \theta\big]$	Risk-free rate (r)	Borrower f's time to default ($T_{f,t}$), inversely related to their probability to default and a convex function of their leverage	None direct	Interest rate is increasing with the risk free rate and decreasing with borrowers' time to default (i.e. increasing with their leverage)

Continues in next page

Table 7 (continued)

Reference(s)	Potential borrowers	Debt maturity	Collateral	Interest rate mechanism	Baseline rate	Borrower-specific IR component	Bank-specific IR component	Short description
Botta et al. (2019)	Households, Firms	Implicit	No	$r_{i,t} = r^B_{t-1} \cdot \big(1 + Debt_{i,t}/Y_{i,t}\big)$	Interest rate on risk-free government bonds (r^B_t)	Borrower i's leverage ratio (total debt to HH's monthly disposable income or firm's net profit)	None direct	Interest rate is calculated as a mark-up on the risk-free rate, determined based on each loan's risk, as reflected by the benchmark debt-service ratio
Giri et al. (2019)	Firms	Implicit	No	$r_{z,t} = i^{CB}_{t-1} + rp_{z,t}$	Central bank rate (i^{CB}_{t-1})	Borrower z's debt to net worth ratio ($rp_{z,t}$)	None direct	Interest rate is a mark-up on the policy rate, associated with an increasing function of borrowers' leverage
Reissl (2020a)	Firms	Implicit	No	$r^j_L = r_0 + \mu(iL^j + \frac{rep^j}{CF^j})$	Rate set by the monetary authority (r_0)	Borrower j's debt-service ($iL^j + rep^j$) over their cash flow (CF^j)	None direct	Interest rate is a spread, increasing with borrowers' debt service to cash flow ratio, added to the baseline rate
Caiani et al. (2016, 2019)	Firms	5 years	No	$r^l_{b,t} = r^{-l}_{b,t-1}(1 \pm FN)$	Market average interest rate in the previous period ($r^{-l}_{b,t-1}$)	None	Bank's capital ratio influences the interest rate by adding (subtracting) a stochastic mark-up (FN) when it is above (below) target	Interest rate is determined on the basis of the bank's current capital ratio and a common target capital ratio

Continues in next page

Table 7 (continued)

Reference(s)	Potential borrowers	Debt maturity	Collateral	Interest rate mechanism	Baseline rate	Borrower-specific IR component	Bank-specific IR component	Short description
Reissl (2020b)	Households, Firms	Implicit	Partial	$r^i_X = \theta_i(\overline{r_{cb}} + default_i)$	Mid-point interest rate of central bank's corridor ($\overline{r_{cb}}$)	None	Default rate of loans/mortgages in bank i's portfolio. Mark-up (θ_i) stochastically changes based on bank's interest revenues and its rate compared to the sector's average	Interest rate is a mark-up over the sum of the central bank rate and the current default rate of loans/mortgages in bank i's portfolio
Delli Gatti et al. (2010)	Firms	Implicit	No	$r^x_{z,t} = \alpha A^{-\alpha}_{z,t} + \alpha(\ell_{x,t})^\alpha$	None	Borrower x's leverage ratio ($\ell_{x,t}$)	Bank z's net worth ($A_{z,t}$)	Interest rate is decreasing with the financial soundness of the bank and increasing with the borrowers' leverage ratio
Riccetti et al. (2013); Russo et al. (2016)	Households, Firms	Implicit	No	$r_{b,x,t} = r^{CB}_t + \widehat{r_{b,t}} + \overline{r_{x,t}}$	Central bank rate (r^{CB})	Borrower x's leverage ratio	Bank's ability to lend all possible funds stochastically changes component $\widehat{r_{b,t}}$	Interest rate is the sum of three elements; the policy rate, a bank-specific component and a borrower specific one
Cardaci (2018)	Households	Single- or 120 periods	Yes	$r_{t,b,h} = \bar{r} + \widehat{r_{b,t}} + r_{t,h}$	Central bank rate ($\bar{r}$)	Borrower h's total debt service ratio ($r_{t,h} = \mu TDS$, with $\mu > 0$)	Bank b's leverage ratio ($\widehat{r_{b,t}} = \rho LB_{b,t}$, with $\rho > 0$)	Interest rate is the sum of three elements; the policy rate, a bank-specific component and a borrower specific one

Continues in next page

Table 7 (continued)

Reference(s)	**Potential borrowers**	**Debt maturity**	**Collateral**	**Interest rate mechanism**	**Baseline rate**	**Borrower-specific IR component**	**Bank-specific IR component**	**Short description**
Gurgone et al. (2018)	Firms	random, between 2 and 10 periods	No	$r_{t,h,j} = \frac{1+cf_{h,t}}{1-\rho_{j,t}} - 1$	None	The default probability ($\rho_{t,h,j}$) is an increasing function of borrower j's leverage rate , corrected for the financial vulnerability perceived by bank h in terms of its own expected shortfall	Bank's cost of funds ($cf_{j,t}$) i.e. interest paid on its liabilities	Each bank charges an interest rate, taking into account their counterparty risk and their own cost of funds
Cardaci and Saraceno (2019)	Households	Single-period	No	$r^L_{t,h} = \bar{r} + \widehat{r}_t + r_{t,h}$	Central bank rate ($\bar{r}$)	Borrower h's total debt service ratio ($r_{t,h} = \mu TDS$, with $\mu > 0$)	System-specific component; reflecting the sensitivity of the bank to the overall household debt-to-GDP ratio	Interest rate is the sum of three elements; the policy rate, an economy-wide component and a borrower-specific one
Dawid et al. (2019)	Firms	18	No	$r^b_{i,t} = r^c(1 + \lambda^B \cdot PD^b_{k,t} + \epsilon^b_t)$	Central bank rate (r^c)	Borrower k's probability of default ($PD^b_{k,t}$)	Stochatically varying component ($\epsilon^b_t \sim U[0,1]$) proxying operational costs	Interest rate is an increasing function of the credit risk reflecting the risk premium that the bank charges to more risky borrowers
Alexandre and Lima (2020)	Firms	10	No	$r_{i,t} = r^B\big(1 + 0.5(l^B_t)^\gamma + 0.5(l_{i,t})^\gamma\big)$	Central bank rate (r^B)	Borrower i's total debt to net worth ratio ($l_{i,t}$)	Bank's total loans over net worth (l^B_t)	Interest rate is a mark-up on the policy rate, associated with an increasing function of bank's and borrowers' leverage

B Expectations formation under the Heuristics Switching Model

The Heuristics Switching Model (HSM) is described in detail in the works of Anufriev and Hommes (2012a,b) and the subsequent literature. However, a brief description is provided here. It consists of 3 parts; a pool of forecasting heuristics, a measure of their performance and a heuristic selection rule.

For a generic variable y_t, the forecasting heuristics are the following linear, adaptive or trend-extrapolating rules:

A weak trend following rule (WTR):

$$y^e_{wtr,t} = y_{t-1} + \omega_{wtr} \cdot (y_{t-1} - y_{t-2}) \tag{12}$$

A strong trend following rule (STR):

$$y^e_{str,t} = y_{t-1} + \omega_{str} \cdot (y_{t-1} - y_{t-2}) \tag{13}$$

An adaptive expectations rule (ADA):

$$y^e_{ada,t} = y^e_{t-1} + \omega_{ada} \cdot (y_{t-1} - y^e_{t-1}) \tag{14}$$

A learning anchoring and adjustment rule (LAA):

$$y^e_{laa,t} = \frac{y^{av}_{t-1} + y_{t-1}}{2} + (y_{t-1} - y_{t-2}) \tag{15}$$

, where $y^{av}_{t-1} = \frac{1}{t}\sum_{j=0}^{j=t-1} y_j$

The first, second and fourth heuristics (WTR, STR and LAA) are trend-extrapolation rules. The first two vary only in the strength of the trend-following behaviour ($\omega_{str} > \omega_{wtr}$), while the last one is a bit more complicated since expectations are extrapolated from an *anchor* which takes into account past realized values (Anufriev and Hommes, 2012a,b). The third heuristic is a simple adaptive rule, where past observations and expectations are combined to form the current expectation of y_t.

The performance of each heuristic i is evaluated according to the following measure:

$$U_{i,t-1} = -(y_{t-1} - y^e_{i,t-1})^2 + \eta U_{i,t-2} \tag{16}$$

The parameter $\eta \in [0, 1]$ controls how strongly past errors affect households' choices, i.e. their memory. If $\eta = 0$, past performance is forgotten, whereas for every other value of $\eta \in (0, 1]$ all past prediction errors are considered, albeit in a decaying manner. Indeed, literature on household memory and expectations suggests that when households make financial decisions their memories do persist (Ampudia and Ehrmann, 2017) and can go back around 10-15 years or more before fading away (Ehrmann and Tzamourani, 2012).

The final element of the HSM is the selection rule which determines the probability $n_{i,t}$ of heuristic i being selected in each period t:

$$n_{i,t} = \delta n_{i,t-1} + (1-\delta)\frac{exp(\beta U_{i,t-1})}{\sum_{i=1}^{4} exp(\beta U_{i,t-1})} \quad (17)$$

The stochastic choice among the forecasting heuristics depends on their performance and two behavioural parameters, namely δ and β. The first parameter, $\delta \in [0, 1]$, represents inertia in heuristic selection, i.e. households' tendency to stick to their previous choice. The second parameter, $\beta \geq 0$, known as "intensity of choice", reflects how fast households will switch to the best performing heuristic. The values of the HSM-related parameters are fixed to the ones identified in the original studies of Anufriev and Hommes (2012a,b) and have been since shown to present excellent fit properties in various settings (Assenza et al., 2014).

C Calibration of initial DSTI and its evolution

One of the key series that controls the supply of credit is households' maximum debt-service-to-income ratio (DSTI). For each credit request it receives, the bank calculates the potential borrower's DSTI as the sum of interest payments plus amortisations (assuming that the demanded loan is granted) over their monthly income. If this figure is below the maximum value accepted by the bank, then the requested amount is granted. Otherwise, only a fraction of it is offered.

The maximum DSTI is assumed to be set by the bank, reflecting its risk tolerance and is exogenously injected into the model. However, publicly available data on DSTI are very scarce. The only available source for such information is the BIS's debt service ratios statistics. The database, as of September 2019, covers 32 countries and provides an estimate of DSTI series at an aggregate level, in quarterly frequency.

This database has two main limitations; it doesn't cover all countries examined in this study and, "*...it does not necessarily accurately measure [DSTI's] level relative to what one could obtain from the correct micro data*" (Drehmann et al., 2015, p. 91). The former limitation extends not only to a country being included, but also to the fact that the reported DSTI series at sector level[38] might not represent its sub-components in a sufficient manner. One can reasonably assume that the two components of total household DSTI, mortgage and consumer credit, follow different paths.[39] Therefore, the calibration procedure aims to overcome these limitations.

The first one, is the most difficult to address in the absence of more detailed data. Thus, it is circumvented in an axiomatic way; it is assumed that the actual DSTI series for consumer credit in any given country (CY, SI and the UK) must have roughly similar dynamics with some of those reported in the database. The second one is approached in a brute-force way; a range of initial DSTI figures is applied on a number of country-specific DSTI growth rates and the resulting path is fed as an input in the model. The pair of {intitial DSTI, DSTI growth rate path} which yields the lowest distance and highest synchronisation with the historical series of consumer credit is the one used throughout the study.

The aforementioned procedure is applied on every country scenario examined in the study. However, only the results for CY are discussed below. The results for SI and the UK verified the ones identified in Papadopoulos (2019) and are not shown here due to space considerations. It is worth mentioning that during the calibration procedure the interest rate setting mechanisms are

[38]Subscript H indicates that the series reflects the household sector, while P the private, non-financial sector.

[39]In (Papadopoulos, 2019, p. 61) the author clearly shows that, for the case of the US, housing and consumer credit DSTIs have indeed very different dynamics.

switched off and the historical series of lending rates are used instead. This is done to ensure that every variable which could affect the evolution of credit is as close as possible to the realised data. Figure 10 presents the evolution of the raw data (left) as well as their interpolated, monthly growth rates (right) for several country cases.

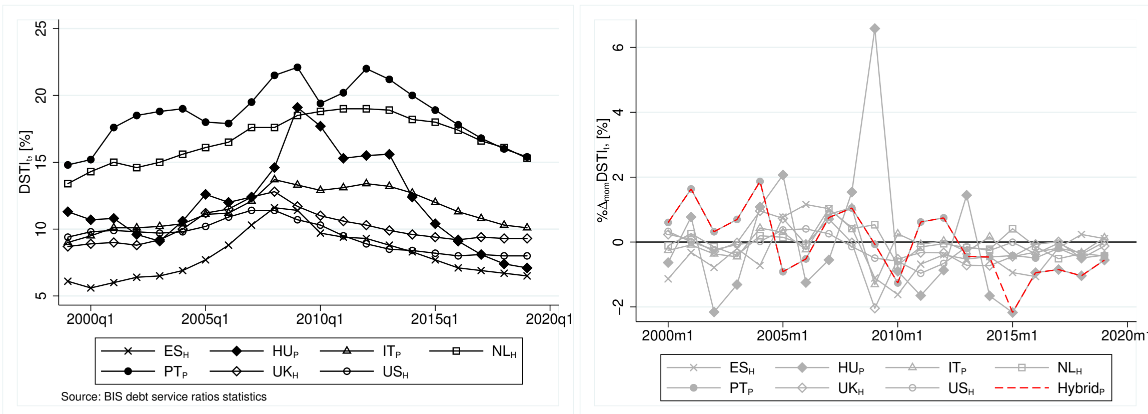

Figure 10: Scenarios examined for DSTI's trajectory calibration in the CY-specific scenario.

The selection of the individual series to be examined in the calibration exercise is based on their presumed adequacy to reflect the unknown, real-world DSTI in CY. The economy was deeply affected by the GFC. In fact, the impact was so strong that the government resorted to international assistance in 2012 and 2013 and the country's banks were bailed-in in early 2013. Thus, it is likely that DSTI followed an increasing pattern until about 2008, then exhibited a plateau and subsequently plummeted after 2013 as both demand and supply froze. An additional element is that in 2013[40], policy makers, in response to the financial crisis, imposed limits on the maximum DSTIs which in all probability must have resulted in a further decline (Central Bank of Cyprus, 2013).

It should be noted that while every scenario shows the previously discussed general pattern, two stand apart; PT_P displays the closest pattern to the expected one before 2014, while HU_P is the one with the sharpest drop after 2014. Thus, a hybrid scenario is created by combining the two monthly growth rates in January 2015. Figure 11 shows the output of the grid-search procedure described above, for the case of the CY-specific scenario.

[40]Effective as of 31 March 2014 and lifted on 1 April 2016.

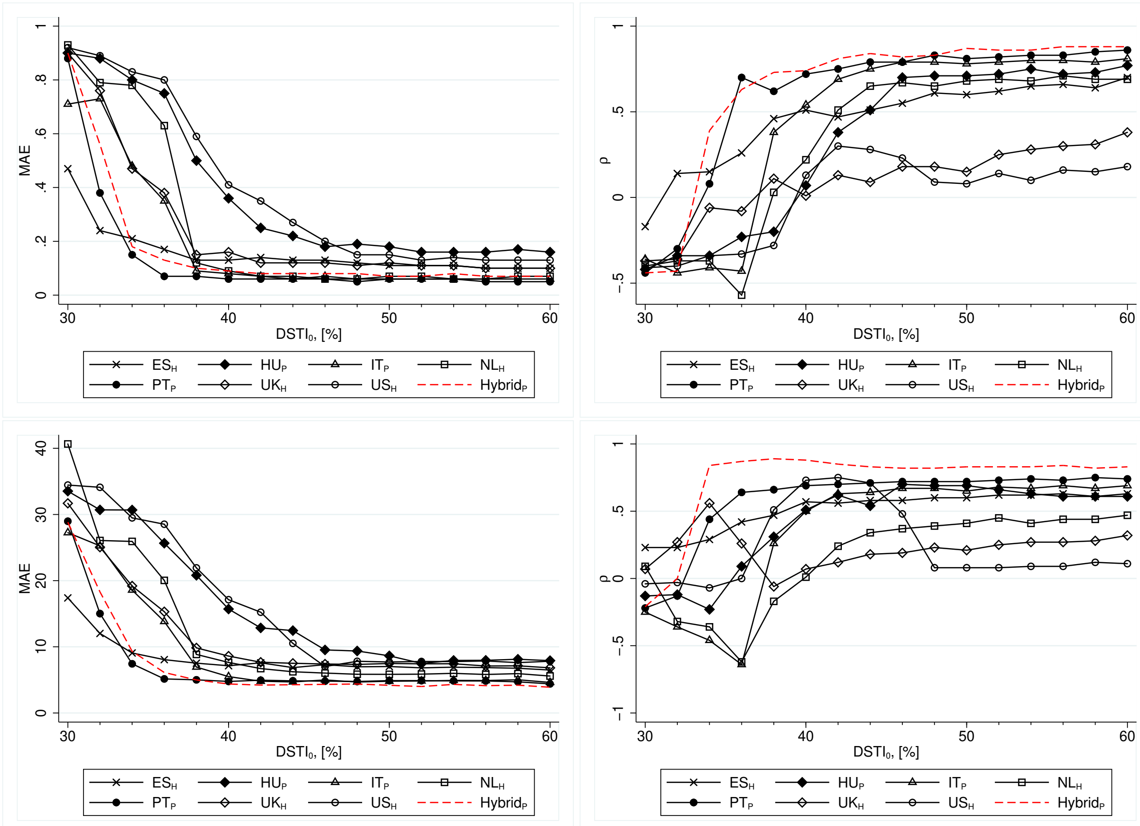

Figure 11: Distance and synchronisation measures between artificial and historical series of consumer credit per initial DSTI values for the CY-specific scenario.

The top row in Figure 11 displays the distance and synchronisation between the model-generated and observed data of consumer credit in (rescaled) levels, while the bottom row in growth rates. Eight scenarios are examined with initial values of DSTI ranging between 30% to 60% per scenario.

The most prominent feature of Figure 11 is the improvement in performance across the board for initial DSTI values above 40%. Practically, each scenario's fit reaches its maximum and does not increase above that level of $DSTI_0$. Nevertheless, three scenarios outperform the rest. In particular PT_P (filled circles) and IT_p (hollow triangles) and the hybrid one (red, dashed line) seem to result in a better fit than the rest. Among the three, the hybrid one exhibits slightly better performance hence, it is the one used in the simulation.