

Valaitis, Vytautas; Villa, Alessandro T.

Article

A machine learning projection method for macro-finance models

Quantitative Economics

Provided in Cooperation with:

The Econometric Society

Suggested Citation: Valaitis, Vytautas; Villa, Alessandro T. (2024) : A machine learning projection method for macro-finance models, Quantitative Economics, ISSN 1759-7331, The Econometric Society, New Haven, CT, Vol. 15, Iss. 1, pp. 145-173, <https://doi.org/10.3982/QE1403>

This Version is available at:

<https://hdl.handle.net/10419/296356>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



<https://creativecommons.org/licenses/by-nc/4.0/>

A machine learning projection method for macro-finance models

VYTAUTAS VALAITIS

School of Economics, University of Surrey

ALESSANDRO T. VILLA

Economic Research Department, Federal Reserve Bank of Chicago

We use supervised machine learning to approximate the expectations typically contained in the optimality conditions of an economic model in the spirit of the parameterized expectations algorithm (PEA) with stochastic simulation. When the set of state variables is generated by a stochastic simulation, it is likely to suffer from multicollinearity. We show that a neural network-based expectations algorithm can deal efficiently with multicollinearity by extending the optimal debt management problem studied by Faraglia, Marcet, Oikonomou, and Scott (2019) to four maturities. We find that the optimal policy prescribes an active role for the newly added medium-term maturities, enabling the planner to raise financial income without increasing its total borrowing in response to expenditure shocks. Through this mechanism, the government effectively subsidizes the private sector during recessions.

KEYWORDS. Machine learning, incomplete markets, projection methods, optimal fiscal policy, maturity management.

JEL CLASSIFICATION. C63, D52, E32, E37, E62, G12.

1. INTRODUCTION

In this paper, we exploit the computational gains that derive from the robustness to multicollinearity of neural networks to extend the optimal debt management problem studied by Faraglia et al. (2019) to four maturities. The hedging benefits provided by the additional maturities allow the government to respond to expenditure shocks by raising financial income without increasing the total outstanding debt. Through this mechanism, the government effectively subsidizes the private sector in recessions.

Vytautas Valaitis: v.valaitis@surrey.ac.uk

Alessandro T. Villa: alessandro.villa@chi.frb.org

We are thankful to Andrea Lanteri, Lukas Schmid, and Matthias Kehrig for their encouragement. The paper received the Student Award of the Society of Computational Economics and benefited from comments by Albert Marcet, Serguei Maliar, Lilia Maliar, Swapnil Singh, and seminar participants at the Duke Macro Breakfast, the 2018 Baltic Economic Conference, the Society for Computational Economics 24th International Conference, the 2018 Econometric Society Summer European Meeting, and the 2019 Econometric Society African meeting. Disclaimer: The views expressed in this paper do not represent the views of the Federal Reserve Bank of Chicago or the Federal Reserve System. Declaration of conflicts of interest: none.

We use a neural network (NN) in a supervised machine learning fashion to approximate the expectation terms typically contained in the optimality conditions of an economic model, in the spirit of the Parameterized Expectations Algorithm (PEA) with stochastic simulation, introduced by [den Haan and Marcet \(1990\)](#) and in a similar fashion to [Duffy and McNelis \(2001\)](#). On the one hand, stochastic simulation methods allow us to tackle problems with a high number of state variables, since they calculate solutions only in the states that are visited in equilibrium (i.e., the ergodic set). On the other hand, when the set of state variables is generated by a stochastic simulation, it is likely to suffer from multicollinearity. In this context, this paper makes two contributions. First, we show that an NN-based expectations algorithm can deal efficiently with multicollinearity by extending the optimal debt management problem studied by [Faraglia et al. \(2019\)](#) to four maturities. Second, we show that the optimal debt management policy prescribes an active role for the medium-term maturities, enabling the planner to raise financial income without increasing its total borrowing in response to expenditure shocks. We consider this problem a particularly interesting economic application that also poses significant computational challenges for four reasons.

First, the number of state variables increases in the number and length of maturities available. Second, this class of problems includes forward-looking constraints, and the problem can be made recursive at the cost of adding even more state variables. Following [Marcet and Marimon \(2019\)](#), we formulate the recursive Lagrangian to solve for the time-inconsistent optimal policy under full commitment with multiple maturities. When markets are incomplete, the Ramsey planner needs to keep track of all promises made in the previous periods. Because of these reasons, optimal maturity management problems suffer from the curse of dimensionality (see [Bellman \(1961\)](#)). For example, the optimal debt management problem with four maturities considered in Section 4 features 46 state variables. Third, because of the maturities, many of these state variables are multicollinear when the model is solved by using a stochastic simulation approach.¹ Fourth, this class of problems does not have a stochastic steady state, as documented in [Aiyagari, Marcet, Sargent, and Sappala \(2002\)](#), and tends to frequently hit the borrowing and lending constraints. Such properties render the model particularly hard to solve using perturbation methods around a particular point.²

In Section 4, we use the methodology to study the optimal government debt management policy when the Ramsey planner can issue an increasing number of debt in-

¹The state space includes lagged values of the same variables (e.g., lagged values of outstanding bonds and Lagrange multipliers). Multicollinearity in the state space might prevent standard regression-based algorithms from converging because the estimated regression coefficients may never stabilize due to high estimation variance and because misspecification of the true policy function under multicollinearity may lead to severe prediction bias, as we show in Section 2.5. Alternatively, people have used the stochastic simulation based on regularization (see [Judd, Maliar, and Maliar \(2011\)](#)) or have extended the PEA algorithm to condensed PEA; see [Faraglia et al. \(2019\)](#). In Section 5, we discuss how the NN-based expectations algorithm improves upon these methods.

²[Bhandari, Evans, Golosov, and Sargent \(2017b\)](#) propose a method that allows one to approximate a system around a current level of government debt, and [Lustig, Sleet, and Yeltekin \(2008\)](#) on the other hand, solve the optimal fiscal policy problem in incomplete markets with seven maturities up to 7 periods using value function iteration on a sparse grid.

struments with different maturities. Intuitively, the prices of longer maturities are typically more responsive to shocks than prices of shorter maturities. This differential response creates opportunities for hedging by borrowing in long-term and saving in short-term bonds. In this case, the value of liabilities falls by more than the value of assets in response to negative shocks (see [Angeletos \(2002\)](#), [Buera and Nicolini \(2004\)](#) and [Faraglia et al. \(2019\)](#)). Additionally, the fact that short bond prices are not as responsive to shocks allows the planner to smooth the price of new debt issuance by rebalancing the portfolio toward the longer maturities in economic booms and toward the shorter maturities in recessions. We find that the planner actively uses the additional medium-term maturities to exploit both the hedging and the price smoothing benefits. The government holds leveraged positions in all bonds and rebalances the portfolio with more emphasis on the shorter maturities in recessions. We find that, when the number of available maturities increases from two to three (and four), the total amount of outstanding debt becomes procyclical. The additional maturities allow the government to respond to expenditure shocks by raising financial income without increasing the total outstanding debt. Through this mechanism, the government effectively subsidizes the private sector in recessions, resulting in higher leisure and less volatile labor taxes.

1.0.0.1 Literature review This paper contributes to two strands of literature: (i) numerical methods in economics and (ii) optimal fiscal policy.

In terms of methods, this paper builds on the seminal work of [den Haan and Marcet \(1990\)](#), who introduced PEA. The idea of using neural networks to parameterize decision rules in a similar fashion to PEA goes back to [Duffy and McNelis \(2001\)](#). Our paper contributes to this literature by showing that an NN-based expectations algorithm can deal efficiently with multicollinearity by extending the optimal debt management problem studied by [Faraglia et al. \(2019\)](#) to more than two maturities. In particular, we exploit the computational gains to study the optimal government debt management problem of [Faraglia et al. \(2019\)](#) with three and four maturities, which yields new economic insights. Note that PEA has been extended more recently (see [Faraglia, Marcet, Oikonomou, and Scott \(2014\)](#) and [Faraglia et al. \(2019\)](#)) to deal with multicollinearity (condensed PEA) and overidentification (Forward-States PEA). Our methodology builds on condensed PEA and Forward-States PEA, in the context of optimal fiscal policy, allowing for machine learning to reduce the state space endogenously and handling multicollinearity effectively when a stochastic simulation approach is adopted. In contrast, condensed PEA achieves this result by introducing an external loop that tests a subset of the state space as a candidate to solve the model.

In contemporaneous work, [Maliar, Maliar, and Winant \(2021\)](#) and [Maliar and Maliar \(2022\)](#) discuss how neural networks can handle multicollinearity. In particular, they do so in the context of the [Krusell and Smith \(1998\)](#) model. We complement their work by demonstrating the robustness to multicollinearity in the context of optimal fiscal policy. Additionally, we show that the interaction between the capability of a neural network to deal with multicollinearity and its flexibility in approximating generic policy functions plays an important role in generating unbiased predictions. In Section 2.5, we show that if a researcher precommits to approximate the policy functions with polynomials that

are misspecified then, under multicollinearity among the state variables, the predictions will be biased. Thanks to the flexible nonparametric nature of a neural network, which does not require making *ex ante* assumptions about the functional form of the policy functions, this problem disappears.

PEA can potentially be used in combination with other standard econometric techniques that tackle the problem of multicollinearity, as in [Judd, Maliar, and Maliar \(2011\)](#). Similar to our paper, [Judd, Maliar, and Maliar \(2011\)](#) adopt a stochastic simulation approach and show how already established methods in econometrics can be used to alleviate the multicollinearity problem using a multicountry neoclassical growth model. We discuss the relation between our method and the methods of [Faraglia et al. \(2019\)](#) and [Judd, Maliar, and Maliar \(2011\)](#) in greater detail in Section 5.

Other papers that use machine learning to solve economic models include [Scheidegger and Bilionis \(2019\)](#), [Azinovic, Gaegauf, and Scheidegger \(2022\)](#), [Fernández-Villaverde, Hurtado, and Nuño \(2023\)](#), and [Duarte, Duarte, and Silva \(2023\)](#). [Fernández-Villaverde, Hurtado, and Nuño \(2023\)](#) use deep neural networks to approximate the aggregate laws of motion in a heterogeneous agents model featuring strong nonlinearities and aggregate shocks. [Duarte, Duarte, and Silva \(2023\)](#) casts the economic model in continuous time and uses neural networks to approximate the Bellman equation. [Maliar, Maliar, and Winant \(2021\)](#) and [Azinovic, Gaegauf, and Scheidegger \(2022\)](#) approximate all the model equilibrium conditions using neural networks and use the simulated data to train them. [Azinovic, Gaegauf, and Scheidegger \(2022\)](#) solve the life-cycle model with borrowing constraints, aggregate shocks, and financial frictions using unsupervised machine learning. The main difference of our paper is to leverage on supervised machine learning to deal effectively with the problem of multicollinearity typical of stochastic simulation approaches. In this context, we show how our algorithm can alleviate the curse of dimensionality, allowing us to explore the problem of the optimal maturity structure of government debt in a more realistic environment.

Our application also contributes to the strand of literature on optimal fiscal policy. In particular, it is relevant to the literature on the optimal maturity structure of government debt.³ [Lustig, Sleet, and Yeltekin \(2008\)](#) find that the optimal policy prescribes an almost exclusive role to the longest maturity in a model with no-lending constraints and a New Keynesian model where bonds are nominal. In our setting, we allow for government lending and study the hedging benefits of a choice between multiple maturities of real bonds. [Bhandari et al. \(2017b\)](#) study the optimal maturity structure in an open economy with two maturities, and [Bigio, Nuño, and Passadore \(2023\)](#) allow for a finite number of maturities in an economy with liquidity costs of issuing debt, where liquidity costs differ by maturity. [Faraglia et al. \(2019\)](#) is the closest paper to ours and studies the role of frictions in a closed economy with two types of bonds. Solving the Ramsey problem considered in this paper is particularly challenging, as the dimension of its state space increases significantly in function of the length of the maturities and the number of bonds. Moreover, this class of problem includes forward-looking constraints,

³[Aiyagari et al. \(2002\)](#), [Angeletos \(2002\)](#), [Buera and Nicolini \(2004\)](#), [Lustig, Sleet, and Yeltekin \(2008\)](#), [Faraglia et al. \(2019\)](#), [Bhandari et al. \(2017b\)](#), and [Bigio, Nuño, and Passadore \(2023\)](#).

so the commonly used recursive representation can not be adopted. [Marcet and Marimon \(2019\)](#) provide an alternative formulation to solve for the time-inconsistent optimal contract under full commitment: a recursive Lagrangian or saddle-point functional equation. The solution involves adding even more state variables to the original problem. These additional state variables, necessary to recursify the problem, create history dependence. In this context, we use our methodology to extend the literature to study optimal debt management with three and four maturities in a closed economy. We find that the optimal policy prescribes an active role for the medium-term bonds. The additional maturities enable the planner to raise financial revenue without increasing the total outstanding debt, in response to a positive expenditure shock. We show that, through this mechanism, the government uses the additional maturities to effectively subsidize the private sector in recessions, resulting in more leisure and less volatile labor taxes.

The paper is organized as follows. Section 2 is a user guide that introduces the reader to PEA, machine learning, and how to combine them in a simple Neoclassical Investment Model example. Section 3 introduces the reader to the problem of multicollinearity using a one-bond economy studied in [Aiyagari et al. \(2002\)](#) and describes the details of the NN-based expectations algorithm using a general model with N maturities. Section 4 presents and discusses the calibration and the quantitative results for the extended model with three and four maturities. Section 5 discusses and compares the NN-based expectations algorithm to other state-of-the-art methods. Section 6 concludes.

2. USER GUIDE: MACHINE LEARNING AND PEA

This section serves as an introduction to supervised machine learning. Specifically, it focuses on how to use it to solve a dynamic economic model in a similar fashion to PEA with stochastic simulation.⁴ Hence, the purpose of this section is solely to introduce the methodology in a simple environment. The method allows us to investigate more realistic models of increased complexity. Its benefits are highlighted in the application presented in Section 3 and arise from the ability of the algorithm to approximate nonlinear policy functions in the presence of a large and multicollinear state space.

2.1 Environment

The typical dynamic model contains intertemporal Euler equations, intratemporal Euler equations, and laws of motion

$$\begin{aligned} f^{\text{Inter}}(c_t, X_t) &= \beta \mathbb{E}[g(c_{t+1}, X_{t+1}) | X_t], \\ f^{\text{Intra}}(c_t, X_t) &= 0, \\ X_{t+1} &= h(X_t, c_t, \xi_{t+1}), \end{aligned}$$

⁴For a general introduction to machine learning, the reader can refer to [Hastie, Tibshirani, and Friedman \(2009\)](#). For a course tailored to economists, the reader can refer to the lecture notes by Jesús Fernández-Villaverde available here: <https://www.sas.upenn.edu/~jesusfv/teaching.html>.

where $c_t \in \mathbb{R}^C$ is a vector of C controls (with E dynamic choices and $C - E$ static choices), $X_t \in \mathbb{R}^S$ is a vector of endogenous and exogenous state variables, $\beta \in (0, 1)$ is a time-discount factor, $f^{\text{Inter}} : \mathbb{R}^C \times \mathbb{R}^S \rightarrow \mathbb{R}^E$, $f^{\text{Intra}} : \mathbb{R}^C \times \mathbb{R}^S \rightarrow \mathbb{R}^{C-E}$, $g : \mathbb{R}^C \times \mathbb{R}^S \rightarrow \mathbb{R}^E$, and $\xi_{t+1} \in \mathbb{R}^I$ is a vector of innovation shocks. For example, in the stochastic neoclassical investment model, c_t corresponds to consumption, f^{Inter} corresponds to the marginal utility of consumption, f^{Intra} does not apply if the model does not include intratemporal choices (e.g., labor), $g \equiv f(c_{t+1})(z_{t+1}\alpha K_{t+1}^{\alpha-1} + 1 - \delta)$, $X_t \equiv \{K_t, z_t\}$ is a vector that contains capital stock and TFP, and $h(X_t, c_t, \xi_{t+1})$ is a function that describes the laws of motion for capital stock, given by the resource constraint $K_{t+1} = (1 - \delta)K_t - c_t + z_t K_t^\alpha$ and the TFP Markov process, that is,

$$X_{t+1} = h(X_t, c_t, \xi_{t+1}) = \begin{pmatrix} K_{t+1} \\ \log z_{t+1} \end{pmatrix} = \begin{pmatrix} (1 - \delta)K_t - c_t + z_t K_t^\alpha \\ \rho \log z_t + \xi_{t+1} \end{pmatrix}.$$

The typical PEA approximates the conditional expectations in the intertemporal Euler equations as polynomial functions of the state space X_t ,

$$\forall e \in [1, E]: \quad \mathbb{E}[g_e(c_{t+1}, X_{t+1})|X_t] \simeq P_n(X_t; \eta_e).$$

The polynomial typically used in the PEA is

$$P_n(X_t; \eta_e) = \exp\left(\eta_{e,0} + \sum_{p=1}^P \sum_{s=1}^S [\eta_{e,p,s} \cdot (\ln X_{s,t})^p]\right),$$

where $\eta_e = [\eta_{e,0}, \eta_{e,1,1}, \dots, \eta_{e,1,S}, \dots]$. For a given sequence of exogenous aggregate shocks $\{\xi_t\}_{t=1}^T$, an initial guess of the polynomials' parameters η^1 , the standard stochastic PEA (described in Algorithm 1) aims to find parameters $\eta^n = \{\eta_1^n, \dots, \eta_E^n\}$ that solve all Euler equations and all laws of motion.

When $X \equiv \{X_t\}_{t=T_0}^{T-1}$ is generated by a stochastic simulation as in Algorithm 1, the matrix $X^T X$ is often ill-conditioned.⁵ Hence, with a finite-precision computer, the inverse of $X^T X$ cannot be computed reliably and it is challenging to compute the linear regression in line 9 of Algorithm 1. This problem potentially leads to jumps in the regression coefficients and failure to converge.

Moreover, in the simple illustrative case of the neoclassical investment model, a first-order polynomial ($P = 1$) is enough to approximate the expectation term in the Euler equation. Generically speaking, richer models that feature a larger state space and nonlinearities require the use of higher-order approximation ($P \gg 1$) and/or cross-state terms. These circumstances further aggravate the multicollinearity problem as the matrix $\hat{X}^T \hat{X}$, with $\hat{X} \equiv \{X_t, X_t^2, \dots\}_{t=0}^T$, is even more ill-conditioned.

2.2 Supervised machine learning

In this paper, we use machine learning as a tool to learn how to represent the function that maps from the set of simulated state variables $\{X_t\}_{t=0}^T$ to the set of simulated

⁵Let $\lambda = \lambda_1, \dots, \lambda_S$ be the vector of eigenvalues of the matrix $X^T X$, such that $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_S \geq 0$. Ill-conditioning refers to the fact that the ratio λ_1/λ_n is large, implying the matrix is close to being singular.

Algorithm 1 Stochastic (simulations) PEA.

Precondition: initial state X_0 , sequence $\{\xi_t\}_{t=0}^T$, initial guess η_n^1 , and dampening $0 < w < 1$

- 1: **while** η_n^i converges **do**
- 2: **for** $t \leftarrow 0$ to T **do** ▷ Generate $X \equiv \{X_t\}_{t=0}^T$
- 3: $c_t \leftarrow$ Solve $f^{\text{Intra}}(c_t, X_t) = 0$ and $f^{\text{Inter}}(c_t, X_t) = \beta P_n(X_t; \eta_n)$
- 4: $X_{t+1} \leftarrow h(X_t, c_t, \xi_{t+1})$
- 5: **end for**
- 6: **for** $t \leftarrow 0$ to $T - 1$ **do** ▷ Generate $Y \equiv \{y_t\}_{t=0}^{T-1}$
- 7: $y_t \leftarrow g(c_{t+1}, X_{t+1})$
- 8: **end for**
- 9: $\hat{\eta}_n^i \leftarrow (X^T X)^{-1} X^T Y$ ▷ Regress to find new weights
- 10: $\eta_n^{i+1} \leftarrow w \cdot \hat{\eta}_n^i + (1 - w) \cdot \eta_n^i$ ▷ Update with dampening
- 11: **end while**

terms $\{y_t\}_{t=0}^{T-1}$. For example, in the neoclassical investment model with log-utility over consumption, this would serve the purpose of representing the function

$$P(K_t, z_t) = \mathbb{E}[c_{t+1}^{-1}(z_{t+1}\alpha K_{t+1}^{\alpha-1} + 1 - \delta)|K_t, z_t].$$

Machine learning proposes a flexible structure for the function P and infers a function from the generated data $\{X_t\}_{t=T_0}^{T-1}$ (which we label *training data*) to the set of generated examples $\{y_t\}_{t=T_0}^{T-1}$ (which we label *training examples*). This particular task of using machine learning to learn a function that maps from inputs to outputs based on training data and examples is referred in the literature as *supervised learning*. And neural networks are a powerful class of universal approximators able to deal with strong nonlinearities.

2.3 Fitting neural networks

In the NN-based expectations algorithm, the equivalent of the *regression phase* is called the *training phase*. As described in Supplemental Appendix C (Valaitis and Villa (2024)), a neural network is characterized by unknown weights $\{w, \psi\}$.⁶ Similar to a regression, the objective is to seek weights such that the neural network fits the samples $\{X_t, y_t\}_{t=0}^{T-1}$. More precisely, the problem is to find

$$\{w_{0,m}, w_m; m = 1, 2, \dots, M\}, \quad \{\psi_{0,e}, \psi_e; e = 1, 2, \dots, E\},$$

such that the sum of squares

$$R(w, \beta) = \sum_{e=1}^E \sum_{t=0}^{T-1} (y_{t,e} - \mathcal{F}_e(X_t; w, \psi))^2$$

⁶Appendix C contains details about the neural network structure used in this section.

is minimized. In a standard linear regression setting, typically (but not necessarily) this problem is solved analytically. This problem could also be solved using a gradient iterative procedure (e.g., gradient descent). This approach is typically more robust to multicollinearity since it does not require inverting the matrix $\hat{X}^T \hat{X}$. An iteration n of gradient descent updates the weights of the neural network according to

$$w_m^{(n+1)} = w_m^{(n)} - \gamma_r \sum_{k=1}^K \frac{\partial R_k(w)}{\partial w_m}, \quad (1)$$

$$\psi_e^{(n+1)} = \psi_e^{(n)} - \gamma_r \sum_{k=1}^K \frac{\partial R_k(w)}{\partial \psi_e}, \quad (2)$$

where the gradient can be derived using the chain rule for differentiation. More specifically, the partial derivatives $\frac{\partial R_k(w)}{\partial w_m}$ and $\frac{\partial R_k(w)}{\partial \beta_e}$ in equations (1) and (2) can be efficiently computed through a two-pass algorithm called backpropagation (Rumelhart, Hinton, and Williams (1986)). Backpropagation applies the chain-rule sequentially, iterating from the output layer to the input layer. Each neuron in the hidden layer receives and dispatches information only from and to neurons that are directly connected. For this reason, this process can be efficiently parallelized. When the backpropagation algorithm is applied to a single-layer neural network, it is known as the delta rule (Widrow and Hoff (1960)). One cycle through the full training samples is called a training epoch. In other words, completing a training epoch means that all training samples have had a chance to update the model parameters. Batch (or offline) learning builds the model digesting the entire training set at once, whereas online training allows the network to update the weights as new observations come in. The former is typically implemented by batch gradient descent, when the latter can typically handle larger training sets and is implemented by stochastic gradient descent. When the neural network weights are updated, the speed at which the model changes can be updated through the parameters γ_r in equations (1) and (2). The parameter γ_r is called the learning rate and it is similar in spirit to a dampening parameter. Intuitively, it represents how quickly the model “learns.” It can either be a constant (for batch learning) or optimized dynamically at each update by minimizing the error function.

Other aspects that can affect the fitting of the neural network are: (i) the initial weights, (ii) the problem of overfitting, (iii) inputs normalization, and (iv) the number of neurons. Initial neural network weights are chosen as near zero random values. Figure C.2 suggests that when the weight α is close to zero, the sigmoid approaches a linear function. This choice of initial weights allows the model to adapt to nonlinearities starting from the linear case.⁷ In practical terms, we solve the model by first initializing the neural network to a simplified version of the model. For example, before solving the neoclassical investment model as described in Algorithm 2, it is possible to solve the model analytically (in this particular case, by setting $\delta = 1$), simulate an equilibrium sequence

⁷Substantial research effort has been put into choosing the initial weights depending on the specific neural network architecture (e.g., see Glorot and Bengio (2010) for deep neural networks).

with the analytical solution, and use it to train the neural network. In a more complicated scenario, such as the optimal maturity management problem presented in this paper, we first solve the model without debt (i.e., the government only uses the income from taxes to finance government expenditure) to first initialize the neural network. The model should not overfit the data. Since stochastic simulation methods (such as PEA) only explore a subset of the ergodic set of state variables (i.e., those combinations of state variables simulated in the equilibrium), we optimize the model for out-of-sample predictions. We split the simulated data randomly in the training set (in-sample) and validation set (out-of-sample) with a 70–30 proportion, respectively.⁸ The number of epochs is determined by maximizing the neural network's performance on the validation set. All inputs are normalized to have mean zero and unitary standard deviation. This procedure ensures that all inputs have a comparable magnitude. If some inputs were of a bigger order of magnitude, the weights linked to those inputs would experience a faster update speed.⁹ This could potentially impair the learning process and lead to a slower convergence, or worse, mean squared prediction errors. The choice of the number of neurons in the hidden layer should be guided by the trade-off between in-sample fit and out-of-sample performance, as illustrated in Figure 1 (the figure refers to the neoclassical growth model that we present as an illustrative example in the next section). Increasing the number of hidden units tends to increase the in-sample fit but leads to overfitting. We select the number of units by minimizing the mean squared prediction error calculated on the validation set.

2.4 Example: Neural networks and PEA applied to the neoclassical investment model

This section describes the implementation of the NN-based expectations algorithm applied to the neoclassical investment model. We use Matlab and we leverage on the *Statistics and Machine Learning Toolbox*. The illustrative example code, together with the comparison with other methods and the procedure that selects the optimal number of neurons, is publicly available.¹⁰ The purpose of using Matlab and disseminating this application is to facilitate the adoption of machine learning in economics with well-known tools in an easy-to-adopt package. In this example, we use a single layer neural network with 12 neurons (this number of neurons minimizes the mean squared prediction error out-of-sample as shown in Figure 1).

We first calculate the steady state, which is particularly useful to build a guess to initialize the neural network weights. The command *feedforwardnet(12)* creates a neural network with one hidden layer that contains 12 neurons. By default, this neural network is trained (through the function *train*) with Levenberg–Marquardt backpropagation, and

⁸We did not find a significant difference in allocations and forecast errors when we changed the training set to be 50% and 90% of the sample.

⁹In the context of deep neural networks, the distribution of each layer's inputs varies during the training phase, since the weights of the previous layers change as well. Typically, this requires adopting lower learning rates and carefully choosing initial parameters. This problem is known as internal covariate shift. In this context, batch normalization can achieve higher a learning rate (see Ioffe and Szegedy (2015)) by reducing the problem of internal covariate shift.

¹⁰Downloadable from <https://www.alessandrotenzinvilla.com/research.html>

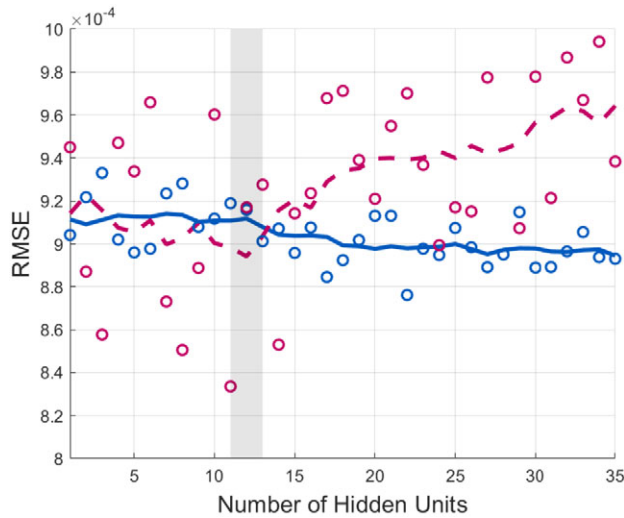


FIGURE 1. Root Mean Squared Error (RMSE) in function of the number of neurons in the hidden layer. *Note:* The figure shows the relation between the number of hidden units and neural network performance in the neoclassical growth model. Solid blue line—network performance on the training set. Dashed purple line—network performance on the validation set. Circles show network performance for a specific number of units. Lines represent the moving averages.

has a maximum number of epochs set to 1000. We generate an initial data set using the deterministic steady state and substituting the value of the shock.

We then proceed to solve the model using the equivalent of Algorithm 1, except we use the neural network to approximate the expectation contained in the optimality conditions of the model. We call this the NN-based expectations algorithm, and a detailed description of the code is laid out in Algorithm 2.

In a more complex environment with a large state space—where a stochastic simulation approach is desirable—the advantages of this method lie in the interaction between a satisfactory approximation of the model nonlinearities and the degree of multicollinearity among the simulated states. In PEA, the choice of polynomials is quite arbitrary as the policies' functional forms are ex ante unknown (one has to rely on an ex post accuracy test to make sure that the approximation is satisfactory). If the functional form of the chosen approximator cannot satisfactorily approximate the equilibrium policies, the presence of multicollinearity can lead to bias in the parameter estimates. The neural network does not suffer from this problem as it is a universal approximator. In the next section, we conclude the user guide illustrating this point.

2.5 Neural networks and multicollinearity

One general problem is that the functional form, not just the parameters, that links the state variables and the approximated terms is ex ante unknown. A standard practice is to make these approximations using polynomials with order and cross-terms typically chosen through trial and error. When the policies are correctly specified, multicollinearity leads to consistent, yet noisy parameter estimates. However, if the chosen functional

Algorithm 2 NN-based expectations algorithm applied to the neoclassical growth model.

Precondition: parameters $\beta, \alpha, \delta, \rho, \sigma^\epsilon, k_1$, and ϵ ; utility functions $u(c) = \log(c)$, $u_c(c) = c^{-1}$.

```

1: ▷ Simulate log AR(1) process and an initial guess for  $k$  and  $c$ 
2:  $\log(z_{t+1}) \leftarrow \rho \cdot \log(z_t) + \xi_{t+1}$ 
3:  $k_{I,t} \leftarrow [(1 - \beta(1 - \delta))/(\beta \cdot \alpha \cdot z_t)]^{1/(\alpha-1)}$ 
4:  $c_{I,t} \leftarrow z_t \cdot k_{I,t}^\alpha - \delta \cdot k_{I,t}$ 

5: ▷ Create and train the NN using the initial  $k_I$  and  $z$  alongside the RHS $t$ 
6: Net  $\leftarrow$  feedforwardnet(12)
7: RHS $t$   $\leftarrow u_c(c_{I,t+1})[\alpha \cdot z_{t+1} \cdot k_{I,t+1}^{\alpha-1} + 1 - \delta]$ 
8: Net  $\leftarrow$  train(Net, [ $k_{I,t}, z_t$ ], RHS $t$ )
9:  $k_{old} \leftarrow k_I$ 

10: ▷ Solve the model
11: while error >  $\epsilon$  do
12:   ▷ Generate  $\{c_t\}_{t=1}^T$  and  $\{k_t\}_{t=1}^T$ 
13:   for  $t \leftarrow 1$  to  $T$  do
14:     ▷ The output of the function Net is the expectation object, given the current
       state
15:      $\mathbb{E}_t[\text{RHS}_{t+1}] \leftarrow \text{Net}([k_t; z_t])$ 
16:      $c_t \leftarrow u_1^{-1}(\beta \cdot \mathbb{E}_t[\text{RHS}_{t+1}])$ 
17:      $k_{t+1} \leftarrow z_t \cdot k_t^\alpha + (1 - \delta) \cdot k_t - c_t$ 
18:   end for
19:   ▷ Train the NN using the new  $k$  and  $z$  alongside the above RHS
20:   RHS $t$   $\leftarrow u_c(c_{t+1}) \cdot (\alpha \cdot z_{t+1} \cdot k_{t+1}^{\alpha-1} + 1 - \delta)$ 
21:   Net  $\leftarrow$  train(Net, [ $k_t, z_t$ ], RHS $t$ )
22:   ▷ Checking convergence and updating  $k_{old}$ 
23:   error  $\leftarrow \max(|k_{old,t} - k_t|)$ 
24:    $k_{old,t} \leftarrow k_t$ 
25: end while

```

forms are not suitable to approximate the true policy functions, multicollinearity can potentially lead to severely biased and less precise predictions, as we show in the following simple example.

For simplicity, imagine that we would like to approximate the policy function of the neoclassical investment model with two TFP shocks, full depreciation $\delta = 1$, and log-utility $u(c) = \log c$. The true functional form of the policy function for the capital choice (in this simple case it can be solved analytically) is

$$k_{t+1}(k_t, z_{1,t}, z_{2,t}) = \beta \alpha \exp(z_{1,t}) \exp(z_{2,t}) k_t^\alpha, \quad (3)$$

where the true parameters are $\alpha = 0.36$, $\beta = 0.95$, and $z_{1,t}$ is a log-AR(1) process with persistence 0.8 and standard deviation of the innovation shock 0.0224. Moreover, z_2 is determined using the following formula:

$$z_2 = \lambda \hat{z}_2 + (1 - \lambda) z_1,$$

where \hat{z}_2 is a log-AR(1) process with persistence 0.8 and standard deviation of the innovation shock 0.0224.

We generate, through stochastic simulation, equilibrium sequences $\{X_t, y_t\}_{t=0}^T$ where $y_t = k_{t+1}$, and $X_t = [k_t \ z_{1,t} \ z_{2,t}]$ is a vector that contains the three state variables. The objective is to use $\{X_t, y_t\}_{t=0}^T$ in order to infer the functional form of equation (3). When $X \equiv \{X_t\}_{t=0}^T$ is generated by a stochastic simulation as in Algorithm 1, the matrix $X^T X$ is often ill-conditioned. We simulate different degrees of multicollinearity by randomly generating sequences $\{x_{1,t}\}_{t=0}^T$ and $\{x_{2,t}\}_{t=0}^T$ with different degrees of correlation (i.e., different values of $\lambda \in [0, 1]$), and we calculate the associated $\{y_t\}_{t=0}^T$ using equation (3). We evaluate the success of the prediction in function of different degrees of multicollinearity using a (i) linear polynomial and a (ii) neural network. Note that on purpose we incorrectly assume that the mapping between state variables and policy is linear $y_t = \beta_1 x_{1,t} + \beta_2 x_{2,t} + \beta_3 x_{3,t}$.¹¹ Also note that the neural network has a flexible non-parametric nature and, therefore, does not require making ex ante assumptions about the functional form of the policy functions. The success of the prediction is assessed using the mean squared prediction error (MSPE), which is the average prediction error at time t over many training samples. The error can be decomposed in bias and variance terms

$$\text{MSPE}_t = \mathbb{E}[(y_t - \hat{y}_t)^2] = \underbrace{\mathbb{E}[y_t - \mathbb{E}(\hat{y}_t)]^2}_{\text{Bias}_t^2} + \underbrace{\mathbb{E}[\hat{y}_t - \mathbb{E}(\hat{y}_t)]^2}_{\text{Variance}_t}. \quad (4)$$

Figure 2 reports the average MSPE for the entire validation set in function of the correlation between the two exogenous shocks $\{x_{2,t}\}_{t=0}^T$ and $\{x_{3,t}\}_{t=0}^T$. Note that the higher the correlation, the higher the multicollinearity between $\{x_{2,t}\}_{t=0}^T$ and $\{x_{3,t}\}_{t=0}^T$.

The higher the correlation between the state variables, the higher the inaccuracy of the polynomial regression model. Moreover, if we decompose the MSPE using equation (4), we find that most of the prediction error comes from the bias squared term as shown in Figures B.2 and B.3 in Appendix B.4. Because of its nonparametric nature, the neural network adapts to the shape of the function to approximate without having to guess the functional form ex ante. This experiment suggests that the nonparametric nature of a neural network is particularly handy in solving economic models characterized by policy functions with functional forms that are ex ante unknown and that potentially

¹¹We purposely choose a polynomial that cannot correctly approximate the true policy functions, since often the true functional form is ex ante unknown. We check that the results are robust to many types of misspecification. However, the purpose of the following example is simply to illustrate the possibility that misspecification under multicollinearity can lead to biased predictions. This problem would not arise with a universal approximator, such as neural networks, because they do not require prespecification of the functional form.




FIGURE 2. Mean squared prediction error with a neural network and a polynomial. *Note:* The figure shows the mean squared prediction error $1/n \sum_{t=1}^n [y_t - \mathbb{E}(\hat{y}_t)]^2$ in function of the correlation between x_2 and x_3 . Blue line with circles—NN, purple line with crosses—polynomial regression.

contain significant nonlinearities and whose domain presents multicollinear states.¹² In the next section, we illustrate the use of neural networks in a model that contains such features.

3. MODEL AND SOLUTION METHOD

The model we work with is an extension of the one-bond economy analyzed in [Aiya-gari et al. \(2002\)](#) and extended to two bonds in [Faraglia et al. \(2019\)](#). We work with this model for two reasons. First, it is a difficult computational problem that features a large multicollinear state-space with nonlinearities difficult to approximate with a parametric approach (i.e., borrowing and lending constraints).¹³ Second, extending [Faraglia et al. \(2019\)](#) to more than two maturities is a relevant economic problem since it helps in determining the optimal maturity structure of government debt. We start by introducing the reader to a one-bond economy with a single maturity of N periods. We then present our methodology in a general model with N maturities. The numerical advantages of our methodology allow us to explore the optimal maturity structure of government debt with three bonds. Quantitative results are presented in Section 4.

3.1 Illustrative model: One-bond economy

The economy is populated by a representative household with preferences over consumption c and leisure l . The representative household chooses sequences of consump-

¹²Note that another option would be to specify a rich polynomial structure with many higher order and cross-terms. One problem with such approach is that higher-order terms of the same variable are extremely multicollinear.

¹³The nonparametric nature of a neural network makes it suited to approximating policy functions with strong nonlinearities, which would be harder to capture with polynomials.

tion $\{c_t\}_{t=0}^\infty$ and leisure $\{l_t\}_{t=0}^\infty$ to maximize its time-0 expected lifetime utility

$$\mathbb{E}_0 \sum_{t=0}^\infty \beta^t [u(c_t) + v(l_t)],$$

subject to the budget constraint

$$p_t^N b_{t+1}^N + c_t = (1 - \tau_t)(1 - l_t) + p_t^{N-1} b_t^N,$$

where b_t^N indicates an N -periods maturity bond and p_t^N is its corresponding price.¹⁴ The only source of aggregate risk in the economy is an exogenous stream of government expenditures $\{g_t\}_{t=0}^\infty$. In each period, the government can finance g_t by: (i) levying a proportional labor tax τ_t and (ii) by issuing a nonstate contingent bond with maturity of N periods. Hence, the government's budget constraint is

$$g_t + p_t^{N-1} b_t^N = \tau_t(1 - l_t) + p_t^N b_{t+1}^N.$$

The aggregate resource constraint of the economy is $c_t + g_t = 1 - l_t$, where $1 - l_t$ is the period's GDP. We assume the government can buy back and reissue the entire stock of outstanding debt in each period. The government sets taxes and issues debt to solve a Ramsey taxation problem. We adopt the primal approach and assume the government's ability to borrow and lend is bounded. Under these conditions, the government's problem is

$$\max_{\{c_t\}_{t=0}^\infty, \{b_t\}_{t=0}^\infty} \mathbb{E}_0 \sum_t \beta^t [u(c_t) + v(1 - c_t - g_t)],$$

subject to a sequence of measurability constraints¹⁵

$$b_{t+1}^N \beta^N \mathbb{E}_t[u_{c,t+N}] - b_t^N \beta^{N-1} \mathbb{E}_t[u_{c,t-1+N}] - g_t u_{c,t} + (u_{c,t} - v_{l,t})(g_t + c_t) = 0,$$

with borrowing and lending limits¹⁶

$$\bar{M} \geq b_{t+1}^N, \quad \underline{M} \leq b_{t+1}^N.$$

The government's optimality conditions are

$$u_{c,t} - v_{l,t} + \mu_t (u_{cc,t} c_t + u_{c,t} + v_{ll,t} (c_t + g_t) - v_{l,t}) + u_{cc,t} (\mu_{t-N} - \mu_{t-N+1}) b_{t-N+1}^N = 0, \tag{5}$$

$$\mu_t = \mathbb{E}_t(u_{c,t+N})^{-1} \left[\mathbb{E}_t(u_{c,t+N} \mu_{t+1}) + \frac{\xi_{U,t}}{\beta^N} - \frac{\xi_{L,t}}{\beta^N} \right], \tag{6}$$

$$b_{t+1}^N \beta^N \mathbb{E}_t(u_{c,t+N}) = b_t^N \beta^{N-1} \mathbb{E}_t(u_{c,t+N-1}) - g_t u_{c,t} - (u_{c,t} - v_{l,t})(g_t + c_t), \tag{7}$$

¹⁴In principle, households are able to trade government securities in the secondary market. However, since we assume households are identical, there is no trade in equilibrium and, for ease of notation, we omit these trades from the household's budget constraint.

¹⁵See Aiyagari et al. (2002) for details on how to use the recursive Lagrangian approach in this context.

¹⁶ $\bar{M}_N \geq b_{t+1}^N$ is the government saving constraint, which is equivalent to a household's borrowing constraint.




FIGURE 3. Autocorrelation function of the equilibrium bond sequence. *Note:* The figure shows the autocorrelation function of b_t^N . The numbers are obtained after simulating the model equilibrium dynamics for $T = 5000$.

where μ_t is the Lagrange multiplier on the time t measurability constraint, and $\xi_{U,t}$ and $\xi_{L,t}$ are the Lagrange multipliers on the upper and the lower bounds, respectively. By issuing debt at time t , the government commits to increasing taxes and/or to reissuing debt at time $t + N$. When the government sets taxes between time t and time $t + N$, it needs to take into account its past actions in the form of all lags of the state variables up to N . More formally, the Ramsey planner’s state space X_t is

$$X_t = \{g_t, \{\mu_{t-i}\}_{i=1}^N, \{b_{t-i}^N\}_{i=0}^{N-1}\}.$$

The state space contains $2N + 1$ variables, with many lags of the same state variable (e.g., μ), which tend to be highly correlated with each other. Moreover, equation (6) reveals that the Lagrange multiplier on the implementability constraint μ_t follows a random walk, creating an additional source of multicollinearity between the state variables. We solve the model with maturity $N = 10$, and we report in Figure 3 the autocorrelation function of the simulated equilibrium bond’s sequence $\{b_t^N\}$. It is clear that the previous 10 lags of the same variable, which are all part of the state space, are highly correlated with each other in the simulated sequence.

For this reason, the model is hardly solvable using PEA (Algorithm 1). In the literature, this problem has been tackled by an algorithm called condensed PEA. Condensed PEA proposes to approximate the expected values in equations ((5), (6), and (7)) using functions of a subset X_t^C of the state space X_t (X_t^C is also called the core set). These approximations are $\mathbb{E}_t(u_{c,t+N}) \simeq P_1(X_t^C; \eta_1)$, $\mathbb{E}_t(u_{c,t+N-1}) \simeq P_2(X_t^C; \eta_2)$ and $\mathbb{E}_t(u_{c,t+N}\mu_{t+1}) \simeq P_3(X_t^C; \eta_3)$, where both the functions and the core set (including its cardinality) are ex ante unknown. The subset X^C of the information set X is selected through an iterative procedure called condensed PEA. In essence, this method adds an additional loop to PEA and keeps extracting orthogonal components from the state space, similar to the Principle Component Analysis (PCA), but the number of factors does not have to be chosen ex ante. A more detailed description of the procedure can be

found in Section 5, Algorithm 4, where we compare our methodology to existing ones in the literature. In the next section, we present our methodology in a model with N maturities. Due to the presence of multiple lagged bonds, the multicollinearity problem is further accentuated.

3.2 Optimal maturity management with N bonds

The economy is populated by a representative household with preferences over consumption c and leisure l . The representative household chooses sequences of consumption $\{c_t\}_{t=0}^{\infty}$ and leisure $\{l_t\}_{t=0}^{\infty}$ to maximize its time-0 expected lifetime utility:

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t [u(c_t) + v(l_t)],$$

subject to the budget constraint:

$$\sum_{i=1}^N p_t^i b_{t+1}^i + c_t = (1 - \tau_t)(1 - l_t) + \sum_{i=1}^N p_t^{i-1} b_t^i,$$

where b_t^i indicates an i -periods maturity bond and p_t^i is its corresponding price. The only source of aggregate risk in the economy is an exogenous stream of government expenditures $\{g_t\}_{t=0}^{\infty}$. In each period, the government can finance g_t by: (i) levying a proportional labor tax τ_t and (ii) by issuing nonstate contingent bonds with maturity $1, \dots, N$. The government's budget constraint reads

$$\sum_{i=1}^N p_t^{i-1} b_t^i = \tau_t h_t - g_t + \sum_{i=1}^N p_t^i b_{t+1}^i.$$

3.2.0.1 Sequential formulation of the Ramsey problem Combining the technology constraint, $c_t + g_t = h_t$, with the household's labor optimality condition, $1 - \tau_t = v_{l,t}/u_{c,t}$, yields an expression for surplus

$$s_t \equiv \tau_t h_t - g_t = c_t - (1 - \tau_t)h_t = c_t - \frac{v_{l,t}}{u_{c,t}}(c_t + g_t).$$

Substitute bonds prices $p_{i,t}$, pinned down by the household's Euler equations, to get

$$\sum_{i=1}^N b_t^i \mathbb{E}_t \left[\beta^{i-1} \frac{u_{c,t+i-1}}{u_{c,t}} \right] = s_t + \sum_{i=1}^N b_{t+1}^i \mathbb{E}_t \left[\beta^i \frac{u_{c,t+i}}{u_{c,t}} \right],$$

with borrowing and lending limits¹⁷

$$\forall i: \quad \bar{M} \geq b_{t+1}^i, \quad \underline{M} \leq b_{t+1}^i, \quad \bar{M}_{\text{total}} \geq \sum_{i=1}^N b_{t+1}^i, \quad \underline{M}_{\text{total}} \leq \sum_{i=1}^N b_{t+1}^i.$$

¹⁷ $\bar{M}_N \geq b_{t+1}^N$ is the government saving constraint, which is equivalent to a household's borrowing constraint.

The optimality conditions are

$$\begin{aligned}
 c_t : \quad & u_{c,t} - v_{l,t} + \mu_t [u_{c,t} - v_{l,t} + u_{cc,t}c + v_{ll,t}(c_t + g_t)] \\
 & + \sum_{i=1}^N (\mu_{t-i} - \mu_{t-i+1}) b_{t-i+1}^i u_{cc,t} = 0, \\
 \forall i, b_{t+1}^i : \quad & \mu_t = [\mathbb{E}_t u_{c,t+i}]^{-1} \left[\mathbb{E}_t \mu_{t+1} u_{c,t+i} + \frac{\xi_{U,t}^i}{\beta^i} - \frac{\xi_{L,t}^i}{\beta^i} + \frac{\xi_{U,t}^{\text{Total}}}{\beta^i} - \frac{\xi_{L,t}^{\text{Total}}}{\beta^i} \right], \\
 \mu_t : \quad & \sum_{i=1}^N b_t^i \mathbb{E}_t \left[\beta^{i-1} \frac{u_{c,t+i-1}}{u_{c,t}} \right] = s_t + \sum_{i=1}^N b_{t+1}^i \mathbb{E}_t \left[\beta^i \frac{u_{c,t+i}}{u_{c,t}} \right],
 \end{aligned}$$

where $\xi_{U,t}$ and $\xi_{L,t}$ are the Lagrange multipliers on the upper and the lower bounds, respectively, and $\xi_{U,t}^{\text{Total}}$ and $\xi_{L,t}^{\text{Total}}$ are the Lagrange multipliers on the upper and the lower bounds on the total bond portfolio. In the following section, we describe our computational strategy in detail. Details on the implementation and results using Epstein–Zin preferences can be found in Appendix A.

3.3 NN-based expectations algorithm

In this section, we describe the main algorithm, which is an extension of the basic idea illustrated in Section 2.4, applied to an optimal fiscal policy model with incomplete markets and multiple maturities. Here, we present the key steps, while implementation details can be found in Appendix B.1. There are N bonds available with maturities from 1 to N periods. The state space at time t is $\mathcal{I}_t = \{g_t, \{b_{t-k}^i\}_{k=0}^{N-1}\}_{i=1}^N, \{\mu_{t-k}\}_{k=1}^N\}$. The neural network needs to approximate $\mathbb{E}_t[u_{c,t+i}]$, $\mathbb{E}_t[\mu_{t+i}u_{c,t+i}]$, and $\mathbb{E}_t[u_{c,t+i-1}]$ in function of \mathcal{I}_t . We model these relationships using one single-layer neural network $\mathcal{ANN}(\mathcal{I}_t)$. In particular, if the long maturity is $N > 1$, then the terms to approximate are

$$\begin{aligned}
 \mathcal{ANN}_1^i(\mathcal{I}_t) &= \mathbb{E}[u_{c,t+i}|\mathcal{I}_t] \quad \text{for } i = [1, \dots, N], \\
 \mathcal{ANN}_2^i(\mathcal{I}_t) &= \mathbb{E}[\mu_{t+1}u_{c,t+i}|\mathcal{I}_t] \quad \text{for } i = [1, \dots, N], \\
 \mathcal{ANN}_3^i(\mathcal{I}_t) &= \mathbb{E}[u_{c,t+i-1}|\mathcal{I}_t] \quad \text{for } i = [1, \dots, N].
 \end{aligned}$$

For example, in the two-bond case there are six terms to approximate and, if the short bond has 1 period maturity, they reduce to five.¹⁸ Given starting values for μ_t and $\{b_t^i\}_{i=1}^N$ and initial weights for \mathcal{ANN} , simulate a sequence of $\{c_t\}_{t=1}^T$, $\{\mu_t\}_{t=1}^T$ and $\{b_{t+1}^i\}_{i=1}^N\}_{t=1}^T$ as follows:¹⁹

¹⁸Use S and N to denote short- and long-bond maturities, respectively. The six terms are $\mathbb{E}_t(u_{c,t+N})$, $\mathbb{E}_t(u_{c,t+N-1})$, $\mathbb{E}_t(u_{c,t+N-1}\mu_{t+1})$, $\mathbb{E}_t(u_{c,t+S})$, $\mathbb{E}_t(u_{c,t+S-1})$, and $\mathbb{E}_t(u_{c,t+S}\mu_{t+1})$. The term that does not require approximation in the latter case is $\mathbb{E}_t(u_{c,t+S-1})$, which becomes just $u_{c,t}$ when $S = 1$.

¹⁹The network can be initially trained using an educated guess for $\{b_{t+1}^i\}_{i=1}^N$, c_t , μ_t . It is important that the initial training sequence is not constant. More details can be found in Appendix B.1.

- As suggested by [Maliar and Maliar \(2003\)](#), we initially restrict the solution artificially within tight bounds on all debt instruments, and refine the solution gradually while we open the bounds slowly. These bounds are particularly important and initially need to be tight and open slowly, since the neural network at the beginning can only make accurate predictions around zero debt, that is, our initialization point. Additionally, we use penalty functions instead of the ξ -terms to avoid out of bound solutions.²⁰ Since μ_t is identified by the first-order condition for b_t^i , it is overidentified if the number of available maturities is greater than one:

$$\forall i: \quad \mu_t = \mathcal{ANN}_1^i(\mathcal{I}_t)^{-1} \left[\mathcal{ANN}_2^i(\mathcal{I}_t) + \frac{\xi_{U,t}^i}{\beta^i} - \frac{\xi_{L,t}^i}{\beta^i} + \frac{\xi_{U,t}^{\text{Total}}}{\beta^i} - \frac{\xi_{L,t}^{\text{Total}}}{\beta^i} \right].$$

We tackle this problem by using the forward-states approach described in [Farglia et al. \(2019\)](#). This involves approximating the expected value terms at time $t + i$ with functions of the state variables that are relevant at $t + 1$ instead of t and invoking the law of iterated expectations, such that we calculate $\mathbb{E}_t \mathcal{ANN}^i(\mathcal{I}_{t+1})$ instead of $\mathcal{ANN}^i(\mathcal{I}_t)$. This is done in two steps. First, we replace the $\mathcal{ANN}^i(\mathcal{I}_t)$ terms in the optimality conditions with $\mathbb{E}_t \mathcal{ANN}^i(\mathcal{I}_{t+1})$ and, instead of approximating $\mathbb{E}_t(u_{c,t+i})$, $\mathbb{E}_t(u_{c,t+i-1})$, and $\mathbb{E}_t(u_{c,t+i}\mu_{t+1})$, we use the information set \mathcal{I}_{t+1} to approximate $\mathbb{E}_{t+1}(u_{c,t+i})$, $\mathbb{E}_{t+1}(u_{c,t+i-1})$, and $\mathbb{E}_{t+1}(u_{c,t+i}\mu_{t+1})$. Then we use Gaussian quadrature to calculate the conditional expectations of the neural network evaluated at \mathcal{I}_{t+1} .

- To perform the stochastic simulation, choose T big enough and find $\{c_t\}_{t=1}^T$, $\{\mu_t\}_{t=1}^T$ and $\{b_{t+1}^i\}_{i=1}^N\}_{t=1}^T$ that solve the following system of $(N + 2)T$ equations:

$$\left\{ \begin{aligned} & \mu_t = \left[\mathbb{E}_t \mathcal{ANN}_1^i(\mathcal{I}_{t+1}) \right]^{-1} \\ & \quad \times \left[\mathbb{E}_t \mathcal{ANN}_2^i(\mathcal{I}_{t+1}) + \frac{\xi_{U,t}^i}{\beta^i} - \frac{\xi_{L,t}^i}{\beta^i} + \frac{\xi_{U,t}^{\text{Total}}}{\beta^i} - \frac{\xi_{L,t}^{\text{Total}}}{\beta^i} \right], \quad \forall i, \\ & u_{c,t} - v_{l,t} + \mu_t [u_{c,t} - v_{l,t} + u_{cc,t}c + v_{ll,t}(c_t + g_t)] \\ & \quad + \sum_{i=1}^N (\mu_{t-i} - \mu_{t-i+1}) b_{t-i+1}^i u_{cc,t} = 0, \\ & \sum_{i=1}^N b_t^i \beta^{i-1} \mathbb{E}_t \mathcal{ANN}_3^i(\mathcal{I}_{t+1}) = u_{c,t} s_t + \sum_{i=1}^N b_{t+1}^i \beta^i \mathbb{E}_t \mathcal{ANN}_1^i(\mathcal{I}_{t+1}). \end{aligned} \right. \tag{8}$$

The system of equations (8) contains multiple Lagrange multipliers (arising from the inequality constraints). This poses a significant computational challenge. Ideally, one would numerically solve the unconstrained model and then verify that the constraints do not bind and if, for example, M^N binds, set $b_{t+1}^N = \bar{M}^N$ and find the associated values for consumption and leisure. In a multiple-bond model, this is challenging because after setting $b_{t+1}^N = \bar{M}^N$, one needs to check if other constraints do not bind in the recomputed solution, and if they do, enforce them and

²⁰We also find that including ξ terms explicitly in the training set improves prediction accuracy. More details can be found in Appendix B.1.

recalculate the solution again, and so on. To overcome this challenge, we augment the objective function with the following differentiable penalty function:

$$\forall i: \Xi(b_{t+1}^i) = \begin{cases} \frac{\phi}{2} \cdot (b_{t+1}^i - \bar{M}^i)^2 \\ \quad + \int \log(1 + \phi \cdot (b_{t+1}^i - \bar{M}^i)) db_{t+1}^i & \text{if } b_{t+1}^i > \bar{M}^i, \\ 0 & \text{if } \underline{M}^i \leq b_{t+1}^i \leq \bar{M}^i, \\ \frac{\phi}{2} \cdot (\underline{M}^i - b_{t+1}^i)^2 \\ \quad + \int \log(1 + \phi \cdot (\bar{M}^i - b_{t+1}^i)) db_{t+1}^i & \text{if } b_{t+1}^i < \underline{M}^i, \end{cases}$$

where ϕ controls the severity of the penalty. More details can be found in Appendix B. The system of equations (8) can be rederived after including the aforementioned penalty function. We solve the system of equations (8) using the Levenberg–Marquardt algorithm. Since this is a local solver, there is no guarantee that the system is solved globally given a particular initial guess. In our implementation, we attempt to solve the system for at most *maxrep* number of different starting points. If the solution errors are below our specified threshold, the algorithm proceeds with the solution and moves to the next period t . If the solution errors are not below our specified threshold, we pick the solution with the lowest error.

3. If the solution error in the stochastic simulation is large, or a reliable solution could not be found, the algorithm automatically restores the previous period neural network and performs the stochastic simulation with a reduced bound. More specifically, if an unreliable solution has been detected in iteration i , the algorithm restores the iteration $i - 1$'s environment and performs the stochastic simulation with

$$\text{Bound}_{i-1} = \alpha \cdot \text{Bound}_{i-1} + (1 - \alpha) \cdot \text{Bound}_{i-2}.$$

4. If the solution calculated shrinking the bound at iteration $i - 1$ is still not satisfactory, the algorithm does not go back another iteration but uses the same neural network and tries to lower the Bound_{i-1} again toward Bound_{i-2} . Once a reliable solution is found, the algorithm proceeds to calculate the solution for iteration i again, but with

$$\text{Bound}_i = \text{Bound}_{i-1} + (\text{Bound}_{i-1} - \text{Bound}_{i-2}).$$

In this way, if an error is detected multiple times we guarantee that both Bound_i and Bound_{i-1} keep shrinking toward Bound_{i-2} , and there should exist a point close enough to Bound_{i-2} such that the system can be reliably solved with both Bound_{i-1} and Bound_i .

5. If the solution found at iteration i is satisfactory, the neural network enters the learning phase supervised by the implied model dynamics, the bounds are increased, and a new iteration starts.

We repeat this procedure until the neural network predictions converge and the simulated sequences of $\{b_t^i\}_{i=1}^N$ and c_t do not change.²¹ Algorithm 3 describes the algorithm in greater detail and Appendix B.1 contains more details.

4. NUMERICAL RESULTS

In this section, we exploit the computational gains that derive from the robustness to multicollinearity of the NN-based expectations algorithm to study the optimal maturity management problem of Section 3 with four maturities of 1, 5, 10, and 15 periods. Specifically, we are interested in the effects on policy and allocations arising from the additional hedging opportunities with respect to a portfolio with only a short and a long maturity. We first present the calibration and then our numerical results.

4.1 Calibration

We calibrate the model following the strategy of Faraglia et al. (2019). Specifically, we use additively separable utility in consumption and leisure

$$u(c) = \frac{c^{1-\gamma}}{1-\gamma}, \quad v(l) = \chi \frac{l^{1-\eta_l}}{1-\eta_l}$$

with $\gamma = 1.5$ and $\eta_l = 1.8$, respectively. We calibrate χ such that households spend on average 2/3 of their time endowment on leisure in the steady state, which gives a value of 2.87.

We set β to 0.96 and for the sake of comparison, we follow the calibration strategy for g_t from Faraglia et al. (2019). We assume that g_t follows an AR(1) process $g_t = \mu_g + \rho_g g_{t-1} + \epsilon_t$, $\epsilon_t \sim N(0, \sigma_g^2)$ with ρ_g equal to 0.95. Then we look for the value of μ_g such that government expenditure is on average equal to 25% of GDP. This gives a value of 0.0042. Lastly, we set the value for σ_g such that g_t is always at least 15% and at most 35% of GDP in a simulated sample of ten thousand periods, which gives a value of 0.0031. Note that such parameterization is also broadly aligned with the estimates from the data.²²

The government has four debt instruments at its disposal. We set maturities to 1, 5, 10, and 15 years and denote b^1 , b^5 , b^{10} , b^{15} as short, medium, and long and very long bonds, respectively. In addition to debt limits on individual bonds, we introduce a total debt limit of $\pm 100\%$ of GDP both in our benchmark model with only short and long bonds and in our calibration with four bonds. A fixed limit on total debt allows us to make a fair comparison and isolate the effects of the hedging benefits of the additional bond on the household's welfare. Table 1 summarizes the parameter values.

Before proceeding, it is worth noting that we tested our methodology with the two-bond case. Our results in a two-bond model confirms the findings of Angeletos (2002) and Faraglia et al. (2019), where the optimal debt portfolio includes a negative short bond position and a positive long bond position, as shown in Table 2. Moreover, as also

²¹There is no need to check μ_t , which can be backed out analytically from the first-order condition for c_t .

²²We obtain very similar estimates using the sum of government consumption and gross investment from the NIPA tables.

Algorithm 3 NN-based expectations algorithm applied to optimal maturity management.

Precondition: parameters from Table 1; utility functions $u(c) = \frac{c^{1-\gamma}}{1-\gamma}$, $u_c(c) = c^{-\gamma}$, $i = 0$, Bound(0) = 0.

```

1: ▷ Simulate AR(1) process
2:  $g_{t+1} \leftarrow \mu_g + \rho_g \cdot g_t + \epsilon_{t+1}$ 

3: ▷ Create and train the NN using initial conditions
4: Net  $\leftarrow$  feedforwardnet(Num. Neurons)

5: ▷ Solve the model
6: while Bound( $i$ ) <  $B_{max}$  OR OutofBoundIter < NumOutofBound do
7:   ▷ Generate  $\{c_t\}_{t=1}^T$ ,  $\{\mu_t\}_{t=1}^T$ , and  $\{b_{t+1}^i\}_{i=1}^N\}_{t=1}^T$ 
8:   for  $t \leftarrow 1$  to  $T$  do
9:     for  $r \leftarrow 1$  to maxrep do
10:       $x_g \leftarrow \{c(r)_{guess}, b(r)_{guess}^1, \dots, b(r)_{guess}^N\}$ 
11:       $\{c_t(r), \mu_t(r), \{b_{t+1}^i(r)\}_{i=1}^N, residuals(r)\} \leftarrow$  Solve (8)  $\{ANN(\mathcal{I}_{t+1}), Bound(i), x_g\}$ 
12:     end for
13:      $r^* \leftarrow \min_r residuals(r)$ 
14:      $\{c_t, \mu_t, \{b_{t+1}^i\}_{i=1}^N\} \leftarrow \{c_t(r^*), \mu_t(r^*), \{b_{t+1}^i(r^*)\}_{i=1}^N\}$ 
15:     end for
16:     if residuals( $r^*$ ) > threshold then Restart from line 7 with a smaller bound
17:     end if
18:     ▷ Train the NN using the new simulated sequences
19:      $\mathcal{I}_t \leftarrow \{g_t, \{\{b_{t-k}^i\}_{k=0}^{N-1}\}_{i=1}^N, \{\mu_{t-k}\}_{k=1}^N\}$ 
20:      $RHS_{1,t}^i \leftarrow u_{c,t+i}$  for  $i = [1 \dots N]$ 
21:      $RHS_{2,t}^i \leftarrow \mu_{t+i} u_{c,t+i}$  for  $i = [1 \dots N]$ 
22:      $RHS_{3,t}^i \leftarrow u_{c,t+i-1}$  for  $i = [1 \dots N]$ 
23:     Net  $\leftarrow$  train(Net,  $\mathcal{I}_{t+1}$ ,  $RHS_t$ )
24:     ▷ Checking convergence and updating  $\{b_{old,t}^i\}_{i=1}^N$  and  $c_{old,t}$ 
25:     error $_b \leftarrow \max(|\{b_{old,t}^i\}_{i=1}^N - \{b_t^i\}_{i=1}^N|)$ 
26:     error $_c \leftarrow \max(|c_{old,t} - c_t|)$ 
27:     if max(error $_b$ , error $_c$ ) <  $\epsilon$  then Break
28:     end if
29:      $\{b_{old,t}^i\}_{i=1}^N \leftarrow \{b_t^i\}_{i=1}^N$ 
30:      $c_{old,t} \leftarrow c_t$ 
31:     Bound( $i$ )  $\leftarrow$  Bound( $i$ ) + BoundStep
32:     if Bound( $i$ ) >  $\bar{M}$  then
33:       Bound( $i$ )  $\leftarrow \bar{M}$ 
34:       OutofBoundIter  $\leftarrow$  OutofBoundIter + 1
35:     end if
36:      $i \leftarrow i + 1$ 
37: end while

```

TABLE 1. Calibrated parameters.

		Parameter	Value
Preferences	Discount factor	β	0.96
	Risk aversion	γ	1.5
	Labor disutility	χ	2.87
	Leisure curvature	η_l	1.8
Government	Average g_t	μ_g	0.0042
	Volatility of g_t	σ_g	0.0031
	Autocorr. of g_t	ρ_g	0.95
	Debt limits	$\bar{M}, \underline{M}, \bar{M}_{total}, \underline{M}_{total}$	$\pm 100\%$ of GDP

shown in Table 2, the bond portfolio positions are large and volatile as in Buera and Nicolini (2004).

4.2 Optimal debt management with three and four bonds

Tables 2 and 3 summarize the equilibrium outstanding debt-to-GDP ratio for each maturity and for each model with an increasing number of bonds. Moments are calculated given a sequence of government expenditure shocks with persistence and volatility specified in Table 1.

As shown in Tables 2 and 3, the optimal policy includes an active use of all available maturities. Table 2 shows that the average position of each maturity is significantly different from zero and that bond positions are volatile, suggesting their active use responding to expenditure shocks. Table 3 shows the correlations of all the maturities with government expenditure and among themselves. First, it shows the position of short maturity is positively correlated with expenditure shocks while the other maturities are negatively correlated. Second, short maturity is negatively correlated with all other maturities. These two together suggest that, in addition to holding a leveraged portfolio on average, it is optimal to rebalance the portfolio toward shorter maturities. As shown in

TABLE 2. Selected bond moments: means and variances.

Model	$\mathbb{E}(b^1/GDP)$	$\mathbb{E}(b^5/GDP)$	$\mathbb{E}(b^{10}/GDP)$	$\mathbb{E}(b^{15}/GDP)$
1 Bond	0.017	-	-	-
2 Bonds	-0.03	-	0.343	-
3 Bonds	-0.555	0.704	0.632	-
4 Bonds	-0.63	0.884	0.908	-0.173
	$\sigma(b^1/GDP)$	$\sigma(b^5/GDP)$	$\sigma(b^{10}/GDP)$	$\sigma(b^{15}/GDP)$
1 Bond	0.243	-	-	-
2 Bonds	0.1	-	0.122	-
3 Bonds	0.591	0.34	0.533	-
4 Bonds	0.218	0.266	0.27	0.374

Note: The table shows the average outstanding debt for each maturity. Moreover, the table also reports the standard deviations of each outstanding position.

TABLE 3. Selected bond moments: correlations.

Model	$\rho(g_t, b_t^1)$	$\rho(g_t, b_t^5)$	$\rho(g_t, b_t^{10})$	$\rho(g_t, b_t^{15})$		
1 Bond	0.549	-	-	-		
2 Bonds	0.707	-	-0.482	-		
3 Bonds	0.35	-0.181	-0.302	-		
4 Bonds	0.762	-0.094	-0.212	-0.22		
	$\rho(b_t^1, b_t^5)$	$\rho(b_t^1, b_t^{10})$	$\rho(b_t^{10}, b_t^5)$	$\rho(b_t^{15}, b_t^1)$	$\rho(b_t^{15}, b_t^5)$	$\rho(b_t^{15}, b_t^{10})$
1 Bond	-	-	-	-	-	-
2 Bonds	-	-0.796	-	-	-	-
3 Bonds	-0.944	-0.985	0.931	-	-	-
4 Bonds	-0.458	-0.565	0.918	0.047	-0.877	-0.82

Note: The table shows the correlations between each maturity of outstanding debt and government expenditure. Moreover, the table also reports the cross-correlations among the bonds.

Table 4, the additional hedging benefits of the additional maturities are reflected in a higher average leisure and a lower consumption volatility, while the economy sustains a lower average consumption. Labor tax volatility and autocorrelation also decrease significantly, while the average level rises.

Next, we inspect the economic mechanism of how hedging benefits provided by the additional maturities affect household allocations and taxes. As known since Angeletos (2002), differences in long and short bond prices provide a tool to hedge against shocks by borrowing in long bonds and accumulating assets in the short term. Since long prices are more volatile than short prices, when a negative shock hits, the value of government liabilities falls more than the value of government assets, thus providing insurance against negative shocks. In addition to decreasing the government’s liabilities, the differential response of long and short prices also affects the terms of issuing new debt. Since long prices fall more than shorter ones, it becomes cheaper for the planner to obtain funds by issuing shorter debt. This is why we observe portfolio rebalancing and a negative correlation between the long and short bonds.

Table 5 shows how optimal debt management affects government finances as we increase the number of debt instruments.

To inspect how this rebalancing matters for the government’s budget, we decompose government income into labor tax income and net financial income, which is the inflow from issuing new bonds minus the outflow due to outstanding debt. Most im-

TABLE 4. Allocations and policies.

Model	$\mathbb{E}(c_t)$	$\sigma(\ln(c_t))$	$\mathbb{E}(l_t)$	$\sigma(\ln(l_t))$	$\mathbb{E}(\tau_t)$	$\sigma(\ln(\tau_t))$	$\rho(\ln(\tau_t), \ln(\tau_{t-1}))$
1 Bond	0.252	0.029	0.666	0.006	0.247	0.121	0.971
2 Bonds	0.250	0.029	0.668	0.004	0.255	0.106	0.929
3 Bonds	0.248	0.028	0.670	0.005	0.27	0.10	0.914
4 Bonds	0.247	0.027	0.671	0.006	0.274	0.091	0.841

Note: The table shows the effects of the optimal policy on consumption and leisure as the number of bonds increases.

TABLE 5. Government income and borrowing.

Description	Moment	1 Bond	2 Bonds	3 Bonds	4 Bonds
Corr. Debt/GDP and g_t	$\rho(\frac{\sum_t b_t^i}{y_t}, g_t)$	0.547	0.136	-0.079	-0.131
Corr. Net Financial Income and g_t	$\rho(\sum_i (p_t^i b_{t+1}^i - p_t^{i-1} b_t^i), g_t)$	0.186	0.405	0.416	0.511
Corr. Net Financial Income (constant price) and g_t	$\rho(\sum_i (\mathbb{E}(p_t^i) b_{t+1}^i - \mathbb{E}(p_t^{i-1}) b_t^i), g_t)$	0.078	0.11	-0.103	0.019
Av. Net Financial Income (%)	$\mathbb{E}(\frac{\sum_i p_t^i b_{t+1}^i - p_t^{i-1} b_t^i}{y_t})$	-0.142	-0.845	-2.213	-2.569
Av. Labor Tax Income (%)	$\mathbb{E}(\frac{\tau_t(1-l_t)}{y_t})$	24.7	25.5	27.0	27.4

Note: The table shows selected moments from the models with one, two, three, and four maturities. The first row shows the correlation between the outstanding debt/GDP ratio and expenditure shocks. Rows two and three show the correlation between government financial income and expenditure shocks. The last two rows show the average net financial income and the average labor tax income. Net Financial Income is defined as the inflow from issuing new debt at the net of the cost of buying back the outstanding debt. Net Financial Income (constant price) is the counterfactual and corresponds to Net Financial Income holding bond prices fixed at their average values.

portantly, as the number of maturities increases, the correlation between total debt and government expenditures changes sign, as shown in the first row of Table 5.

In the one- and two-bond economy, the government borrows from the private sector to finance expenditure shocks. In the three- and four-bond economy, the government reduces its total debt to subsidize the private sector and smooth its consumption. At the same time, net financial income becomes even more positively correlated with g_t and allows for smoother labor taxes, despite a falling total debt in bad times. The reduction of total debt together with rising financial income is achieved precisely because the planner holds leveraged positions and responds to expenditure shocks by substituting to short bonds.

As further evidence of this mechanism, we construct a counterfactual measure of net financial income assuming that bond prices were fixed at their mean values. The counterfactual correlation is reported in the third row of Table 5. The low correlation here suggests that the comovement between net financial income and government expenditures is achieved by exploiting the differential response of short, medium, and long prices. This indicates that if prices were constant, portfolio rebalancing would have little effect on the cyclicity of financial income and the government's budget.

Looking at the averages in rows four and five, we see that as the number of maturities increases, the government becomes a net payer to the private sector and collects a larger share of its income in labor taxes. This happens because the increase in labor taxes outweighs the decrease in average labor supply. Although average household labor income falls, the household is compensated for holding government debt.

5. COMPARISON WITH ALTERNATIVE METHODS

There are other simulation-based numerical methods designed to address the issue of multicollinearity among state variables. In this section, we discuss and compare our method to the two most prominent ones: the Condensed PEA (C. PEA) used in Faraglia

et al. (2019) and the generalized stochastic simulation algorithm (GSSA) described in Judd, Maliar, and Maliar (2011).

5.1 Relation to condensed PEA

This method extracts orthogonal components from the information set. The method is similar to the Principle Component Analysis (PCA), except that the number of factors does not have to be chosen ex ante. Algorithm 4 reports the pseudocode of condensed PEA and highlights with colors the part of condensed PEA that changes when the NN-based Expectations Algorithm (NN EA) is implemented. In this section, we solve the model with a short one-period bond and a long ten-period bond using both algorithms. The model is solved with individual bonds bounds \bar{M} , \underline{M} set at $\pm 100\%$. As reported in Table 6, the two algorithms reach a similar outcome, but the NN-based expectations algorithm is significantly faster. The speed gains mainly come from the removal of the external loop (lines 1, 13, 14, 15, 16, and 17 in condensed PEA) as the neural network digests the information set at once.

On the one hand, eliminating the external loop (line 1) reduces the complexity of the algorithm significantly, since condensed PEA requires testing an unknown number of combinations of core regressors. On the other hand, our algorithm requires substituting OLS (lines 10 and 11) with a neural network training algorithm, which has higher complexity. Note that the computation time of one entire simulation from 1 to T (lines 3–9) takes significantly more time under condensed PEA.²³ In total, the condensed

Algorithm 4 From condensed PEA to NN-based EA.

```

Precondition: initial state  $X_0$ , initial  $X^{C,1}$  and  $X^{\text{Out}}$ , sequence  $\{\xi_t\}_{t=0}^T$ , initial guess  $\eta_n^1$ ,
dampening  $0 < w < 1$ 
1: while core set  $X^{C,k}$  converges do
2:   while  $\eta_n^i$  converges do
3:     for  $t \leftarrow 0$  to  $T$  do                                     ▷ Generate  $X \equiv \{X_t\}_{t=0}^T$ 
4:        $c_t \leftarrow f^{-1}(P_n(X_t^{C,k}; \eta_n^i))$ 
5:        $X_{t+1} \leftarrow h(X_t, c_t, \xi_{t+1})$ 
6:     end for
7:     for  $t \leftarrow 0$  to  $T$  do                                     ▷ Generate  $Y \equiv \{y_{t+1}\}_{t=0}^T$ 
8:        $y_{t+1} \leftarrow g(c_{t+1}, X_{t+1})$ 
9:     end for
10:     $\hat{\eta}_n^i \leftarrow \text{argmin}(Y - P_n(X^{C,k}; \eta_n^i))^2$                     ▷ Regress to find new weights
11:     $\eta_n^{i+1} \leftarrow w \cdot \hat{\eta}_n^i + (1 - w) \cdot \eta_n^i$                 ▷ Update with dampening
12:  end while
13:   $\omega \leftarrow \text{argmin} X^{\text{Out}} - \omega X^{C,k}$                                ▷ Update core set
14:   $X^{\text{Res}} \leftarrow X^{\text{Out}} - \omega X^{C,k}$ 
15:   $\lambda \leftarrow \text{argmin} Y - P_n(X^{C,k}; \eta_n^i) - \lambda X^{\text{Res}}$ 
16:   $X^{C,k+1} \leftarrow X^{C,k} \cup \lambda X^{\text{Res}}$ 
17: end while

```

²³Lines 3–9 of Algorithm 4 take on average 34 s for condensed PEA, whereas the neural network training phase took around 15 s. Lines 13–16 of Algorithm 4 take on average 0.003 s for condensed PEA, whereas the neural network training phase took around 0.21 s.

TABLE 6. Moments condensed PEA versus NN EA.

Method	Time	Forecast Error	$\mathbb{E}(c_t)$	$\mathbb{E}(l_t)$	$\mathbb{E}(\tau_t)$	$\sigma(\ln(c_t))$	$\sigma(\ln(l_t))$	$\sigma(\ln(\tau_t))$	$\rho(\tilde{S}_t, \tilde{S}_{t-1})$
C. PEA	203,810 s	0.373	0.256	0.663	0.219	0.024	0.007	0.116	0.798
NN EA	23,744 s	0.391	0.256	0.664	0.222	0.024	0.007	0.123	0.869

Note: The table shows the equilibrium moments calculated with condensed PEA and the NN Expectations Algorithm (NN EA). The forecast error is calculated as $\sum_{t=1}^T \sum_{e=1}^E |y_{e,t} - \hat{\mathbb{E}}(y_{e,t})|$, where $y_{e,t}$ here is used to indicate the realized value correspondent to each expectation $\hat{\mathbb{E}}(y_t)$ predicted by either condensed PEA or NN EA. Bond bounds are set as \bar{M} , \underline{M} at $\pm 100\%$. We follow Faraglia et al. (2019) and calculate $\tilde{S}_t = \frac{\bar{b}_t^1}{\bar{b}_t^1 + \bar{b}_t^{10}}$ as the ratio between the market value of short-term debt \bar{b}_t^1 and the market value of the total outstanding debt $\bar{b}_t^1 + \bar{b}_t^{10}$. The computation times are obtained using MATLAB 2023a and a computer with an Intel(R) Core(TM) i7-8750H CPU (9 M Cache, up to 4.10 GHz) with RAM 16 GB.

PEA needed to cycle 4 times before the core set $X^{C,k}$ converged. This means that condensed PEA required approximately four times more iterations than NN EA.

Note that, as the number of maturities increases, the condensed PEA requires testing a much higher number of combinations of core regressors. In this sense, NN EA is a more scalable approach. On a related note, in Appendix B.2, we also present a comparison between condensed PEA and our algorithm based on time complexity. Our calculations show that NN EA has a lower time complexity than condensed PEA, if condensed PEA requires more than one loop on the core set of state variables to converge.

5.2 Relation to GSSA

Next, we compare our methodology to the GSSA method proposed by Judd, Maliar, and Maliar (2011). GSSA resolves the multicollinearity problem using standard econometric techniques (i.e., single value decomposition, principal components, and ridge regression), combined with stochastic simulation. We solve the government debt management problem with one maturity using GSSA.²⁴ In particular, we use ridge regression since, as noted in Judd, Maliar, and Maliar (2011), it works best under severe multicollinearity.²⁵ In our application, ridge regression combined with stochastic simulation works when bond constraints are loose, whereas when the constraints are tight, ridge regression coefficients fail to converge. As we illustrate in details in Appendix B.3, the reason lies in the fact that ridge regression requires choosing penalty parameters. This choice presents nontrivial challenges. In our application, the simulated data changes in each iteration since we are solving the model with an iterative procedure. In principle, this should require updating the penalty parameters in each iteration. However, in our experience, changing penalty parameters at each iteration also creates instability as the simulated data is endogenous to the penalty choice in the previous iterations. For this reason, we fix the penalty parameters (see Appendix B.3 for more details on how we choose the penalty parameters) during the whole procedure and, since the simulated debt sequence tends to change significantly with each iteration when the bond constraints are tight, the algorithm fails to converge. Recall that, as explained in Section 3.3,

²⁴We also attempted to solve the model with two maturities but we were not successful.

²⁵Our application features a very ill-conditioned matrix $X^T X$, especially when the number of maturities increases.

our algorithm initially restricts the solution artificially within tight bounds on all debt instruments, and refines the solution gradually while it opens the bounds slowly. To conclude, the challenges in choosing (or fixing) the penalty parameters, combined with the frequent changes in debt sequences in each iteration induced by initially tight bounds, render ridge regression with stochastic simulation hard to scale in our application. Further details can be found in Appendix B.3.

6. CONCLUSION

In this paper, we exploit the computational gains that derive from the robustness to multicollinearity of neural networks to extend the optimal debt management problem studied by Faraglia et al. (2019) to four maturities. The hedging benefits provided by the additional maturities allow the government to respond to positive expenditure shocks by raising financial income without increasing the total outstanding debt. We show that, through this mechanism, the government uses the additional maturities to effectively subsidize the private sector in recessions, resulting in more leisure and less volatile labor taxes.

REFERENCES

- Aiyagari, S. Rao, Albert Marcet, Thomas J. Sargent, and Juha Sappala (2002), “Optimal taxation without state-contingent debt.” *Journal of Political Economy*, 110 (6), 1220–1254. [146, 148, 149, 157, 158]
- Angeletos, George-Marios (2002), “Fiscal policy with noncontingent debt and the optimal maturity structure.” *Quarterly Journal of Economics*, 117 (3), 1105–1131. [147, 148, 164, 167]
- Azinovic, Marlon, Luca Gaegauf, and Simon Scheidegger (2022), “Deep equilibrium nets.” *International Economic Review*, 63 (4), 1471–1525. [148]
- Bellman, Richard Ernest (1961), *Adaptive Control Processes: A Guided Tour*. Princeton University Press. [146]
- Bhandari, Anmol, David Evans, Mikhail Golosov, and Thomas J. Sargent (2017b), “Fiscal policy and debt management with incomplete markets.” *The Quarterly Journal of Economics*, 132, 617–663. [146, 148]
- Bigio, Saki, Galo Nuño, and Juan Passadore (2023), “Debt-maturity management with liquidity costs.” *Journal of Political Economy Macroeconomics*, 1 (1), 119–190. [148]
- Buera, Francisco and Juan Pablo Nicolini (2004), “Optimal maturity of government debt without state contingent bonds.” *Journal of Monetary Economics*, 51, 531–554. [147, 148, 166]
- den Haan, Wouter and Albert Marcet (1990), “Solving the stochastic growth model by parameterizing expectations.” *Journal of Business and Economic Statistics*, 8 (1), 31–34. [146, 147]

Duarte, Diogo, Victor Duarte, and Dejanir H. Silva (2023), “Machine learning for continuous-time finance.” Working Paper. [148]

Duffy, John and Paul D. McNelis (2001), “Approximating and simulating the stochastic growth model: Parameterized expectations, neural networks, and the genetic algorithm.” *Journal of Economic Dynamics and Control*, 25 (9), 1273–1303. [146, 147]

Fraglia, Elisa, Albert Marcet, Rigas Oikonomou, and Andrew Scott (2014), “Optimal fiscal policy problems under complete and incomplete financial markets: A numerical toolkit.” Working Paper. [147]

Fraglia, Elisa, Albert Marcet, Rigas Oikonomou, and Andrew Scott (2019), “Government debt management: The long and the short of it.” *Review of Economic Studies*, 86, 2554–2604. [145, 146, 147, 148, 157, 162, 164, 168, 169, 170, 171]

Fernández-Villaverde, Jesús, Samuel Hurtado, and Galo Nuño (2023), “Financial frictions and the wealth distribution.” *Econometrica*, 91 (3), 869–901. [148]

Glorot, Xavier and Yoshua Bengio (2010), “Understanding the difficulty of training deep feedforward neural networks.” In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics* (Yee Whye Teh and Mike Titterton, eds.), Proceedings of Machine Learning Research, Vol. 9, 249–256, PMLR, Chia Laguna Resort, Sardinia, Italy. [152]

Hastie, Trevor, Robert Tibshirani, and Jerome Friedman (2009), *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, second edition. Springer. [149]

Ioffe, Sergey and Christian Szegedy (2015), “Batch normalization: Accelerating deep network training by reducing internal covariate shift.” CoRR. abs/1502.03167. [153]

Judd, Kenneth, Lilia Maliar, and Serguei Maliar (2011), “Numerically stable and accurate stochastic simulation approaches for solving dynamic economic models.” *Quantitative Economics*, 2, 173–210. [146, 148, 169, 170]

Krusell, Per and Anthony A. Jr. Smith (1998), “Income and wealth heterogeneity in the macroeconomy.” *Journal of Political Economy*, 106 (5). [147]

Lustig, Hanno, Christopher Sleet, and Sevin Yeltekin (2008), “Fiscal hedging with nominal assets.” *Journal of Monetary Economics*, 55, 710–727. [146, 148]

Maliar, Lilia and Serguei Maliar (2003), “Parameterized expectations algorithm and the moving bounds.” *Journal of Business and Economic Statistics*, 21 (1), 88–92. [162]

Maliar, Lilia and Serguei Maliar (2022), “Deep learning classification: Modeling discrete labor choice.” *Journal of Economic Dynamics and Control*, 135, 104295. [147]

Maliar, Lilia, Serguei Maliar, and Pablo Winant (2021), “Deep learning for solving dynamic economic models.” *Journal of Monetary Economics*, 122, 76–101. [147, 148]

Marcet, Albert and Ramon Marimon (2019), “Recursive contracts.” *Econometrica*, 87 (5), 1589–1631. [146, 149]

Rumelhart, David E., Geoffrey E. Hinton, and Ronald J. Williams (1986), “Learning representations by back-propagating errors.” *Nature*, 323, 533–536. [152]

Scheidegger, Simon and Ilias Bilionis (2019), “Machine learning for high-dimensional dynamic stochastic economies.” *Journal of Computational Science*, 33, 68–82. [148]

Valaitis, Vytautas, and Alessandro T. Villa (2024), “Supplement to ‘A machine learning projection method for macro-finance models’.” *Quantitative Economics Supplemental Material*, 15, <https://doi.org/10.3982/QE1403>. [151]

Widrow, Bernard and Marcian E. Jr. Hoff (1960), “Learning representations by back-propagating errors.” Adaptive switching circuits, IRE WESCON convention record, part 4, 96–104, IRE, New York. [152]

Co-editor Tao Zha handled this manuscript.

Manuscript received 6 August, 2019; final version accepted 11 September, 2023; available online 20 September, 2023.

The replication package for this paper is available at <https://doi.org/10.5281/zenodo.8326559>. The Journal checked the data and codes included in the package for their ability to reproduce the results in the paper and approved online appendices.