Deep Learning for Solving and Estimating Dynamic Macro-Finance Models

Benjamin Fan[1,†], Edward Qiao[2,†], Anran Jiao[3], Zhouzhou Gu[4], Wenhao Li[4], and Lu Lu[3,*]

[1]University High School, Irvine, CA
[2]The Bishop's School, La Jolla, CA
[3]Department of Chemical and Biomolecular Engineering, University of Pennsylvania, Philadelphia, PA 19104, USA
[4]Marshall School of Business, University of Southern California, Los Angeles, CA
[†]These authors contributed equally to this work.
[*]Corresponding author. Email: lulu1@seas.upenn.edu

Abstract

Deep learning has been shown to be an effective method for solving partial differential equations (PDEs) by embedding the PDE residual into the neural network loss function. In this paper, we design a methodology that utilizes deep learning to simultaneously solve and estimate canonical continuous-time general equilibrium models in financial economics, including (1) industrial dynamics of firms and (2) macroeconomic models with financial frictions. Through these applications, we illustrate the advantages of our method: generality, simultaneous solution and estimation, leveraging the state-of-art machine-learning techniques, and handling large state space.

Keywords: Dynamic macro-finance models; Industrial dynamics of firms; Macroeconomic models with financial frictions; Partial differential equations; Deep learning; Parameter estimation

Contents

1 Introduction **3**

2 Machine Learning for PDEs **4**
2.1 Deep neural networks . . . 4
2.2 PINNs for solving forward PDEs . . . 4
2.3 PINNs for solving inverse PDEs . . . 5
2.4 Implementation . . . 5

3 A Model of Industrial Dynamics with Financial Frictions **6**
3.1 Problem setup . . . 6
3.2 Solving the Model . . . 9
3.3 Simultaneously Solving and Estimating the Model . . . 10

4 A Macroeconomic Model with the Financial Sector **11**
4.1 Problem setup . . . 12
4.2 Solving the Model . . . 13
4.2.1 Change of variable to deal with singularity . . . 13
4.2.2 Explicitly solving for ψ . . . 13
4.2.3 Other technical details . . . 14
4.3 Simultaneously Solving and Estimating the Model . . . 16
4.3.1 Estimation of one parameter . . . 16
4.3.2 Estimation of two parameters . . . 18

5 Conclusion **20**

Acknowledgements **21**

References **21**

Appendices **23**

A Setup and Derivations for the Firm Industrial Dynamic Model **23**

B Setup and Derivations for the Macroeconomic Model with a Financial Sector **23**

C Order Reduction in Forward Equation **25**

1 Introduction

Dynamic equilibrium models are the cornerstones of the fast-growing macro-finance literature that tries to understand how financial frictions and asset prices influence economic dynamics, in addition to addressing important policy questions including the design and impact of financial regulation, industrial policy, and monetary policy (He and Krishnamurthy (2013), Brunnermeier and Sannikov (2014), Drechsler et al. (2014), Bianchi and Bigio (2018), Gertler and Kiyotaki (2015), Huang (2018), Di Tella (2019), Krishnamurthy and Li (2020), Maxted (2020)). These models feature high degrees of nonlinearity originating from either agents' financial constraints or preferences, which make the linearization methods widely used in the macro literature infeasible.

The literature has thus far mostly focused on highly tractable models with a small number of state variables (typically one or two). Furthermore, since solving these models is quite time-consuming, model parameters are usually picked by calibration, which involves intensive model evaluation. Matching moments involves solving the model, simulating the model for a long period and calculating the moment value, and repeating the same procedure for a large number of parameter combinations. Although simulated methods of moment have been applied to corporate-finance models (Gomes et al. (2003), Whited and Wu (2006), Hennessy and Whited (2007), Matvos and Seru (2014)), those more explicit and standard methods applied to dynamic equilibrium models are restricted by the curse of dimensionality. Additionally, taking expectations is typical in dynamic problems, but it incurs a significant computational burden. Finally, for different problems, researchers typically need to tailor their numerical methods, which limits the accessibility of the literature, and these methods do not automatically take advantage of the rapidly evolving computational tools.

In this paper, we apply deep-learning methods for solving partial differential equations (Lu et al. (2021a)) to economic settings, simultaneously solving and estimating model parameters and allowing for heterogeneous agents. Our methods have the following advantages: (1) the model is solved globally; (2) the method allows for higher dimensionality; (3) deep learning is a proven and effective method in solving PDEs; (4) estimation and solution are in one step; (5) differentiation is handled analytically which increases numerical accuracy; (6) the underlying package automatically applies the state-of-art machine-learning algorithm, so the method keeps improving itself; (7) the method is versatile and can be applied to a vast variety of problems.

Machine learning (ML) models have already been used in economics and finance, but mainly for the purpose of better predicting economic and financial outcomes in markets such as stocks, insurance, corporate bankruptcy, cryptocurrency, etc. Recent papers, including Fernandez-Villaverde et al. (2020), Chen et al. (2021), Duarte (2018), use deep learning methods to solve economic models. Our paper contributes to this growing literature by combining equilibrium-model estimation with solution and building on the rapidly-improving deep-learning-based solution methods for PDEs in physics and science.

The machine-learning method in Lu et al. (2021a) is a type of physics-informed neural networks (PINNs). PINNs work by embedding PDE residual into the loss function of the neural network via automatic differentiation (Yu et al., 2021). As such, approximating the PDE is no more than minimizing the loss function, which can be done with gradient descent techniques. This method of solving PDEs is mesh-free and simple, and it can be applied to a wide variety of PDE types. In addition to solving forward problems, PINNs can be easily implemented to solve inverse PDE problems. These involve predicting the values of parameters given a set of measurements of the function. PINNs have achieved success with both forward and inverse problems in a diverse range of fields, including optics (Chen et al., 2020), systems biology (Yazdani et al., 2020), fluid mechanics (Raissi et al., 2020; Tartakovsky et al., 2020), and biomedicine (Kissas et al., 2020; Sahli Costabal et al., 2020). However, there have been fewer applications to problems in economics, which this paper will explore.

In this paper, we will consider two models. First, we solve a model of industrial dynamics with financial frictions, which considers an industrial equilibrium of a banking sector that that takes deposits, makes loans, and uses labor input to manage deposits and loans. Then, we consider a macroeconomic model with the financial sector featuring binding constraints, non-linear financial amplifications, and boundary singularity. We demonstrate the advantage of our methodology in providing a general framework for solving PDEs, as opposed to traditional methods, which require designing different algorithms for different problems. We also display the ease of solving inverse problems, in which we assume parameters are unknown and impose moment conditions. Thus, we are able to simultaneously solve and estimate parameters.

In Section 2, we introduce the physics-informed neural network method for solving both forward and inverse partial differential equations, along with the Python library DeepXDE. Following this introduction, in Sections 3 and 4 we apply the PINN method to solve several equations in economics.

2 Machine Learning for PDEs

We first describe general deep neural networks, allowing us to present the framework of physics-informed neural networks (PINNs), which will be leveraged to solve forward and inverse PDEs (i.e., model solution and model estimation). Afterward, overview the hyperparameters used in the following sections.

2.1 Deep neural networks

Although there are several different types of deep neural networks, throughout this paper, we use the feed-forward neural network (FNN). We define an L-layer FNN as a function $\mathcal{N}^L(\mathbf{x})\colon \mathbb{R}^{d_{\text{in}}} \to \mathbb{R}^{d_{\text{out}}}$ and say that there are $L-1$ hidden layers such that the ℓ-th layer has N_ℓ neurons. Clearly, $N_0 = d_{\text{in}}$ and $N_L = d_{\text{out}}$. Furthermore, for each $1 \leq \ell \leq L$, we define a weight matrix $\mathbf{W}^\ell \in \mathbb{R}^{N_\ell \times N_{\ell-1}}$ and bias vector $\mathbf{b}^\ell \in \mathbb{R}^{N_\ell}$. Then, letting $T^\ell(\mathbf{x}) = \mathbf{W}^\ell \mathbf{x} + \mathbf{b}^\ell$ be the affine transformation in the ℓ-th layer, for some non-linear activation function σ, we say

$$\mathcal{N}^L(\mathbf{x}) = T^L \circ \sigma \circ T^{L-1} \circ \ldots \circ \sigma \circ T^1(\mathbf{x}).$$

A network with $L = 4$ can be visualized in Fig. 1. There are several possible activation functions σ, and in this paper, we use both the hyperbolic tangent (tanh) and the swish (Ramachandran et al., 2017) activation functions. The former is the most typical activation function in machine learning, while the latter can better deal with problems with steep gradients, which are features of the problems we will discuss.

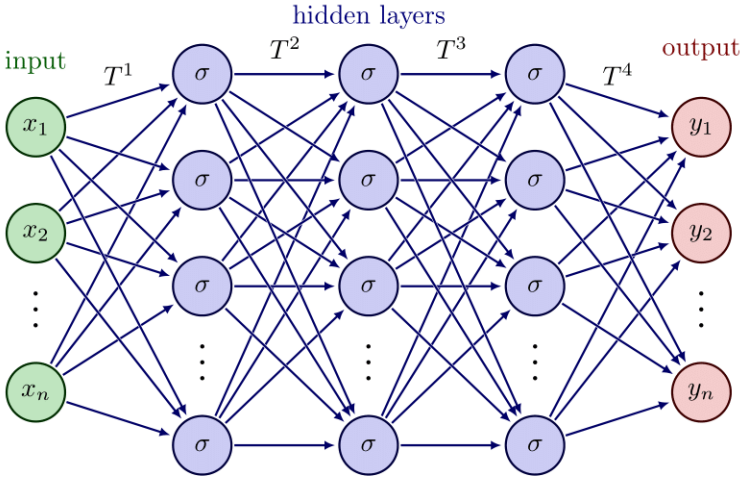

Figure 1: **Visualization of a deep neural network described in Section 2.1.** In this diagram, the number of layers L is 4.

2.2 PINNs for solving forward PDEs

We first discuss the use of PINNs to solve forward PDEs. Consider the following PDE parameterized by $\boldsymbol{\lambda}$ with solution $u(\mathbf{x}, t)$ for $\mathbf{x} = (x_1, \ldots, x_d)$ over the domain $\Omega \subset \mathbb{R}^d$:

$$f\left(\mathbf{x}; \frac{\partial u}{\partial x_1}, \ldots, \frac{\partial u}{\partial x_d}; \frac{\partial^2 u}{\partial x_1 \partial x_1}, \ldots, \frac{\partial^2 u}{\partial x_1 \partial x_d}; \ldots; \boldsymbol{\lambda}\right) = 0, \quad \mathbf{x} \in \Omega \tag{2.1}$$

with boundary conditions

$$\mathcal{B}(u, \mathbf{x}) = 0 \quad \text{on} \quad \partial\Omega.$$

To find the solution, we create a neural network $\hat{u}(\mathbf{x};\boldsymbol{\theta})$ with trainable parameters $\boldsymbol{\theta}$. When training the network, we use $\mathcal{T}_f$ points inside the domain and $\mathcal{T}_b$ points on the boundary. Then, the loss function is defined as

$$\mathcal{L}(\boldsymbol{\theta},\mathcal{T}) = w_f\mathcal{L}_f(\boldsymbol{\theta};\mathcal{T}_f) + w_b\mathcal{L}_b(\boldsymbol{\theta};\mathcal{T}_b) \tag{2.2}$$

with

$$\mathcal{L}_f(\boldsymbol{\theta};\mathcal{T}_f) = \frac{1}{|\mathcal{T}_f|}\sum_{\mathbf{x}\in\mathcal{T}_f}\left\| f\left(\mathbf{x};\frac{\partial\hat{u}}{\partial x_1},\ldots,\frac{\partial\hat{u}}{\partial x_d};\frac{\partial^2\hat{u}}{\partial x_1\partial x_1},\ldots,\frac{\partial^2\hat{u}}{\partial x_1\partial x_d};\ldots;\boldsymbol{\lambda}\right)\right\|_2^2,$$
$$\mathcal{L}_b(\boldsymbol{\theta},\mathcal{T}_b) = \frac{1}{|\mathcal{T}_b|}\sum_{\mathbf{x}\in\mathcal{T}_b}\|\mathcal{B}(\hat{u},\mathbf{x})\|_2^2,$$

for weights w_f and w_b. Now that we have defined a loss function, where the two losses are the L^2 norms of the residuals, we can train like any normal neural network using gradient-based optimizers such as Adam (Kingma and Ba, 2014) and L-BFGS (Byrd et al., 1995). Training the network produces an approximation of $u(\mathbf{x},t)$.

2.3 PINNs for solving inverse PDEs

Next, we discuss using PINNs to solve inverse PDEs. Inverse problems involve some unknown parameters λ in Eq. (2.1) to be solved for, but we are given some extra information about some points $\mathcal{T}_i \in \Omega$ besides the PDE and boundary conditions (Raissi et al., 2019; Lu et al., 2021a):

$$\mathcal{I}(u,\mathbf{x}) = 0 \quad \text{for} \quad \mathbf{x}\in\mathcal{T}_i.$$

Training the PINN in the inverse method is almost identical to training the forward PINN, except that the loss function Eq. (2.2) has an extra term:

$$\mathcal{L}(\boldsymbol{\theta},\boldsymbol{\lambda};\mathcal{T}) = w_f\mathcal{L}_f(\boldsymbol{\theta},\boldsymbol{\lambda};\mathcal{T}_f) + w_b\mathcal{L}_b(\boldsymbol{\theta},\boldsymbol{\lambda};\mathcal{T}_b) + w_i\mathcal{L}_i(\boldsymbol{\theta},\boldsymbol{\lambda};\mathcal{T}_i),$$

where

$$\mathcal{L}_i(\boldsymbol{\theta},\boldsymbol{\lambda};\mathcal{T}_i) = \frac{1}{|\mathcal{T}_i|}\sum_{\boldsymbol{x}\in\mathcal{T}_i}\|\mathcal{I}(\hat{u},\boldsymbol{x})\|_2^2,$$

the L^2 norm of the residual. Optimizing $\boldsymbol{\theta}$ and $\boldsymbol{\lambda}$ together, our solution is $\boldsymbol{\theta}^*,\boldsymbol{\lambda}^* = \arg\min_{\boldsymbol{\theta},\boldsymbol{\lambda}}\mathcal{L}(\boldsymbol{\theta},\boldsymbol{\lambda};\mathcal{T})$.

The above approach incorporates model solution (solving function u) and model estimation (solving parameter λ) together in a consistent way. Furthermore, we can easily generalize the above to a PDE system with multiple functions to be solved.

2.4 Implementation

We will apply PINNs to solve several forward and inverse PDE problems in economics, and we implement these with the Python library DeepXDE (Lu et al., 2021a). In all examples, we use the Glorot uniform initializer. In addition, we occasionally enforce boundary conditions automatically through an output transform (Lu et al., 2021b). Throughout the paper, we use a learning rate decay with the formula

$$\gamma_n = \frac{\gamma_0}{1+\alpha n/S},$$

where n is the number of iterations, γ_n is the learning rate after n iterations, γ_0 is the initial learning rate, α is the decay rate, and S is the decay step. The other hyperparameters are listed in Table 1.

Table 1: **Hyperparameters used for each problem.**

Section	Depth	Width	Activation Function	Optimizer	Learning Rate	#Iterations
3.2	4	64	tanh	Adam	0.005	150000
3.3	4	64	tanh	Adam	0.0001	150000
4.2	7	128	swish	Adam	0.001, 0.00001	300000
4.3.1	6	64	swish	Adam	0.0001	200000
4.3.2	6	64	swish	Adam	0.0005	150000

3 A Model of Industrial Dynamics with Financial Frictions

In this section, we consider an industrial equilibrium of a banking sector that takes deposits, makes loans, and uses labor input to manage deposits and loans. Labor productivity determines the number of loans and deposits that each bank can take. Both loan rate and deposit rates are endogenously determined via the competitive market equilibrium that features bank entry, exit, and equity payouts. The challenges of this problem are twofold: first, banks endogenously determine whether or not to enter or exit the market, so the problem features endogenous entry and exit boundaries; second, we need to track the entire distribution of banks in order to clear the market.

3.1 Problem setup

Time evolves continuously. All bank assets and debts are modeled as short-term. Banks can borrow via deposits at a rate r^d, and via capital market at a rate r (think of this as the policy rate, e.g., FFR). Banks can lend at loan rate r^l, or the capital market rate r. Banks need to hire workers to serve their deposits and loans. For banks with productivity z, the number of loans that l units of labor can serve is

$$f(z,l) = zl^\alpha, \quad \alpha \in (0,1)$$

which features decreasing return to scale. A rationale for this assumption is that as banks get bigger, it is increasingly difficult to find new depositors and new borrowers on which they can earn a profit. Similarly, the amount of deposits that d units of labor can serve is

$$f(z,d) = zd^\alpha.$$

Because banks can borrow and lend freely in the capital market at a rate r, they are not constrained by lending and deposit-taking choices, as long as they are nonnegative.

The stochastic process for z_t is given exogenously as

$$dz_t = \mu(z_t)dt + \sigma(z_t)dB_t,$$

with two reflecting boundaries $\underline{z}$ and $\bar{z}$.

Denote bank equity as e. We impose the financial friction as follows:

$$f(z,l) \le \phi e,$$

$$f(z,d) \le \phi e.$$

Banks incur a fixed operating cost c_f, and thus the instantaneous profit is

$$\pi(e_t, z_t, l_t, d_t) = \underbrace{r_t^l z_t l_t^\alpha}_{\text{lending revenue}} - \underbrace{r_t^d z_t d_t^\alpha}_{\text{deposits interest expense}} + \underbrace{(z_t d_t^\alpha + e_t - z_t l_t^\alpha) r_t}_{\text{net capital market lending}} - \underbrace{w_t \cdot (l_t + d_t)}_{\text{labor cost}} - \underbrace{c_f}_{\text{fixed cost}}.$$

To keep the problem simple, we assume full symmetry between the deposit market and the loan market. Both loan demand and deposit demand functions are in the same form:

$$r^l - r = \beta(L + D_0)^{-\varepsilon},$$

$$r - r^d = \beta(D + D_0)^{-\varepsilon},$$

where D and L are the aggregate amount of deposits and loans, respectively. Due to the full symmetry, the loan spread is equal to the deposit spread, $r^l - r = r - r^d$. In what follows, we will only use the notations on the loan side.

Since the bank can freely adjust its labor input and scales of operations at each instance, the optimal decisions are,

$$l^* = d^* = \min\left\{\left(\frac{(r^l - r)z\alpha}{w}\right)^{\frac{1}{1-\alpha}}, \left(\frac{\phi e}{z}\right)^{1/\alpha}\right\}$$

Therefore, the optimized profit function is

$$\pi^*(e, z) = 2\left(r^l - r\right) z \left(l^*\right)^{\alpha} + e \cdot r - 2wl^* - c_f.$$

Denote $v(e, z)$ as the value function for a bank with equity e and productivity z, with $\underline{v}(e) = e$ the reservation value if a bank exists. We assume that banks pay out equity when they are financially unconstrained, and the equity payout function is

$$\zeta(e, z) = \max\left\{\kappa\left(\phi e - f(z, l^*)\right), 0\right\}.$$

In other words, equity payout smoothly increases as the bank gets further away from the financial constraint. Then we can write the bank equity dynamics as

$$de_t = \underbrace{\left(\pi^*(e_t, z_t) - \mathbf{1}_{(e_t,z_t)\in\mathcal{C}^c} \cdot \zeta(e_t, z_z)\right)}_{\equiv\mu_e(e_t,z_t)} dt - \mathbf{1}_{v(e_t,z_t)<\underline{v}(e_t)} \cdot e_t,$$

where the last term reflects the immediate exit when the bank continuation value is smaller than the liquidation value.

Banks optimize the expected discounted cash flows at the rate r. The HJB equation is:

$$r \cdot v(e, z) = \max\left\{\pi^*(e, z)(1 + v'_e) + (1 - v'_e)\zeta(e, z)1_{(e,z)\in\mathscr{C}^c} + v'_z\mu(z) + \frac{1}{2}v''_{zz}\sigma(z)^2, r\underline{v}(e)\right\},$$

where the set $\mathscr{C}$ is the region where the bank is financially constrained, and the unconstrained indicator $1_{(e,z)\in\mathscr{C}^c}$ is

$$1_{(e,z)\in\mathscr{C}c} = 1\left\{\left(\frac{\left(r^l - r\right) z\alpha}{w}\right)^{\frac{1}{1-\alpha}} < \left(\frac{\phi e}{z}\right)^{1/\alpha}\right\}$$

At the reflecting barriers, we have

$$\partial_z v(e, \underline{z}) = \partial_z v(e, \bar{z}) = 0, \quad \text{for any } e.$$

We also have the boundary condition

$$v(e, z) = \underline{v}(e) = 0.01 \quad \text{when } e = 0.01$$

Banks will exit the market with zero equity, because $de_t \leq 0$ for $e_t = 0$, and $\pi(0) = -c_f < 0$, i.e., zero equity is an absorbing state, and continuing the operations when equity is zero will guarantee negative profit and thus is worse than exiting the market.

Finally, we describe the entry dynamics. Banks first decide whether or not to enter, and then draw their productivity from the distribution $\psi(z)$. We assume that entry incurs a one-time cost c_e, and each entrant has the same initial equity of e_0. The mass of firms entering the market is determined by

$$m = \bar{m} \exp\left(\beta_M\left(\iint_{e,z} v(e, z)\psi(e, z)dedz - c_e\right)\right)$$

The above is a softer version of the free-entry condition. When $\beta_M \to \infty$, entry incentive with respect to entry benefit is going to infinity, so the present value of entry must be zero and we arrive at the free-entry condition

$$\iint_{e,z} v(e,z)\psi(e,z)dedz - c_e = 0$$

The main role of using a softer version is to guarantee numerical stability.

Solve for Invariant Distribution and Estimation. Denote the stationary bank distribution as $g(e,z)$. This distribution does not include banks that exit the market, so we have

$$g(e,z) = 1_{v(e,z)>\underline{v}(e)}g(e,z),$$

where the assumption is that when banks are indifferent between staying or exiting the market, they choose to exit the market.

The Kolmogorov Forward Equation for the stationary distribution in banking industrial dynamic model is:

$$0 = -\frac{\partial}{\partial z}\left(\mu_z(z)g(e,z)\right) - \frac{\partial}{\partial e}\left(\mu_e(e,z)g(e,z)\right) + \frac{1}{2}\frac{\partial^2}{\partial z^2}\left(\sigma(z)^2 g(e,z)\right) + m\psi(e,z),\ v(e,z) > \underline{v}(e)$$

With the stationary distribution, we can get the aggregate loan

$$L = \iint g(e,z)f(z,l^*(e,z))dzde$$

and thus the equilibrium loan spread is determined by the household loan demand function,

$$r^l - r = \beta(L + D_0)^{-\varepsilon}$$

and D is defined as: $\iint_{e,z} d(e,z)g(e,z)dedz$ (here g is not normalized), where $d = z(l^*)^\alpha, l^* = \min\left\{\left(\frac{(r^l-r)z\alpha}{w}\right)^{\frac{1}{1-\alpha}}, \left(\frac{\phi e}{z}\right)^1\right.$

Specification: Model parameter specifications are shown in Table 2, and additional problem setup details can be found in Appendix A.

Table 2: **Parameter specification of model in Section 3.1.**

Description	Value
Bank equity payout rate	$\kappa = 0.005$
Share of labor	$\alpha = 0.3$
Leverage constraint parameter	$\phi = 10$
Benchmark interest rate	$r = 0.03$
Fixed operating cost	$c_f = 0.03$
Entry cost	$c_e = 0.1$
Boundary and mean of productivity	$\underline{z} = 0.2, \overline{z} = 10, z_m = 5$
Lower and upper bound of state space	$e_{\min} = 0.01, e_{\max} = 1.2$
Drift and volatility of z	$\mu(z) = -0.005(z - z_m), \sigma(z) = 0.08$
Deposit supply function's constant	$D_0 = 1.0$
Entrance distribution parameters	$\sigma_\psi = \frac{\overline{z}-\underline{z}}{4}, \bar{e} = 0.15, \bar{m} = 0.1.$
Entrance elasticity	$\beta_M = 1 \times 10^3$

We assume that $\psi(e,z)$ is a truncated normal distribution, ($\Phi(\cdot)$ is the c.d.f. of normal distribution):

$$\psi(e,z) = \frac{1}{\bar{e} \times (\Phi(\frac{\overline{z}-z_m}{\sigma_\psi}) - \Phi(\frac{\underline{z}-z_m}{\sigma_\psi}))} \frac{1}{\sqrt{2\pi\delta_\psi^2}} \exp\left(-\frac{(z-z_m)^2}{2\delta_\psi^2}\right) \times \mathbf{1}_{e<\bar{e}} \times \mathbf{1}_{z\in[\underline{z},\overline{z}]}$$

Boundary Conditions: (1) Banks exit at $\underline{v}(e) = e$ (absorbing boundary), which means: $g(e,z) = 0$, when $v(e,z) = e$. (2) reflecting boundary for stochastic productivity z: $-\mu_z(z)g(e,z) + \frac{1}{2}\frac{\partial}{\partial z}(\sigma(z)^2 g(e,z)) = 0$, when $z = \underline{z}, \overline{z}$.

Moment Targets: define the density $g'(e,z)$ as $\max\{g(e,z), 0\}$ discounted by normalization factor[1]: $\frac{1}{\iint_{e,z} \max\{g(e,z),0\} dedz}$. In estimation, we adjust unknown parameter to match the average equity:

$$e_{\text{target}} = \iint_{e,z} e g'(e,z) dedz$$

Additional Targets: (1) average leverage: $\ell_{\text{target}} = \iint_{e,z} \frac{f(z,l^*(e,z))}{e} g'(e,z) dedz$; (2) average productivity: $z_{\text{target}} = \iint_{e,z} g'(e,z) z dedz$.

3.2 Solving the Model

First, we solve the forward problem with no unknown parameters. We set r^l to its true value of 0.04334. In the next section, we simultaneously solve the model and estimate r^l.

When training the PDE, we used loss weights of 10^3 for the PDE residual and 10^0 for the Neumann boundary conditions. We automatically enforce the Dirichlet boundary condition through the output transform

$$v(e,z) = (e - 0.01) \cdot \mathcal{N}(e,z) + 0.01,$$

where $\mathcal{N}(e,z)$ is the network output. Clearly, we then have $v(0.01, z) = 0.01$.

Furthermore, we trained with 2^{14} training points sampled inside the domain, 2^9 training points sampled on the boundary, and 10000 points sampled inside the domain for testing. Lastly, we used a learning rate scheduler, with an initial learning rate of 0.005, decay rate 1.0, and decay step 5000.

The training results are displayed in Figure 2. The training loss decreases steadily (Fig. 2A) and the L^2 relative error decreases to 0.623% (Fig. 2B). Furthermore, as shown in Figs. 2C and D, the PINN accurately predicts the solution to the equation. The absolute difference between the PINN prediction and reference solution is displayed in the error plot in Fig. 2E.

[1]As KFE operator $\hat{L}^*$ is the Markov process's generator, the distribution is always positive when evolving over time. However, in numerical exercise, $g(e,z)$ at some point can be negative.

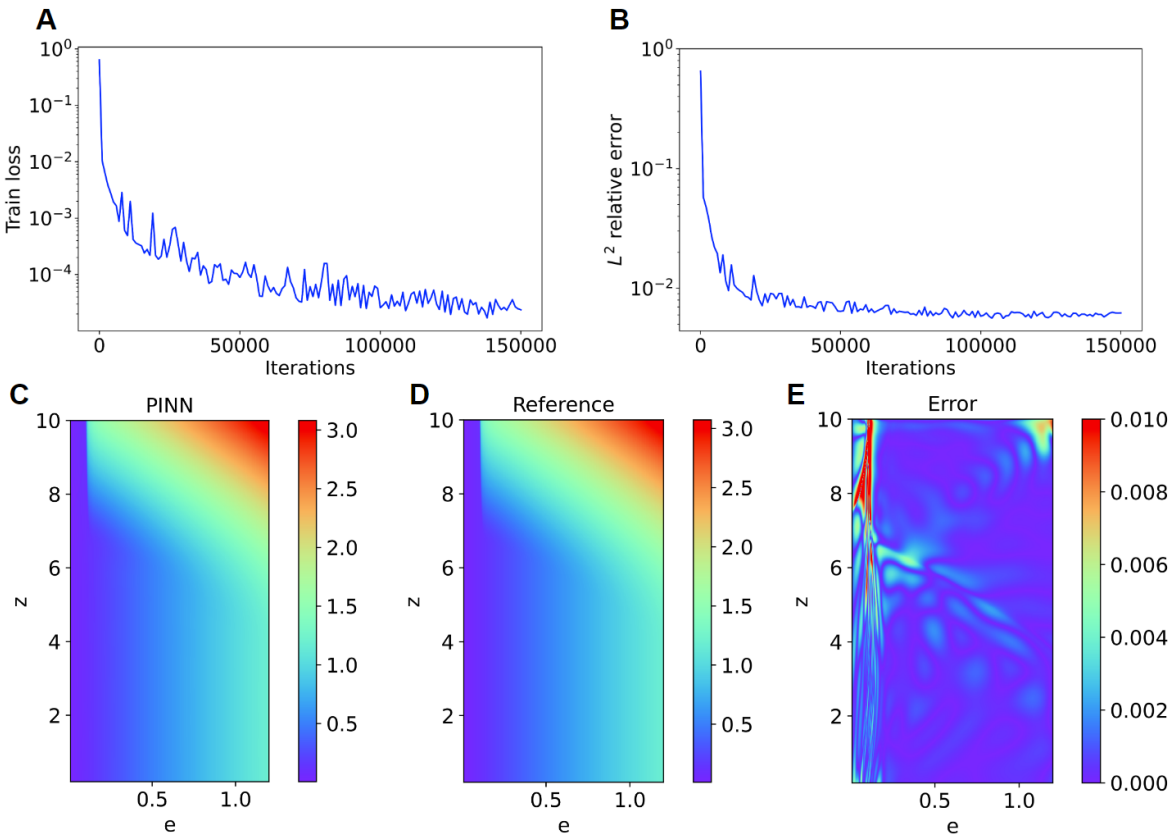

Figure 2: **Example in Section 3.2: PINN results for forward HJB equation.** (**A**) The trajectory of train loss throughout training. (**B**) The trajectory of the L^2 relative error throughout training. (**C**) 2D heat map for predicted solution for v using PINN. (**D**) 2D heat map for reference from Matlab data. (**E**) 2D heat map for the absolute difference between PINN prediction and Matlab reference data.

3.3 Simultaneously Solving and Estimating the Model

Next, we simultaneously solve the model and estimate r^l. Our additional information will come from a description of the entry dynamics. Banks first decide whether or not to enter, and then draw their productivity from the distribution $\psi(e, z)$. We assume that entry incurs a one-time cost c_e, and each entrant has the same initial equity of e_0. The mass of firms entering the market is determined by

$$m = \bar{m} \exp\left(\beta_M \left(\iint_{e,z} v(e,z)\psi(e,z)dedz - c_e\right)\right)$$

The above is a softer version of the free-entry condition. When $\beta_M \to \infty$, entry incentive with respect to entry benefit goes to infinity, so the present value of entry must be zero and we arrive at the free-entry condition:

$$\iint_{e,z} v(e,z)\psi(e,z)dedz - c_e = 0$$

The main role of using a softer version is to guarantee numerical stability.

We assume that $\psi(e, z)$ is a truncated normal distribution, ($\Phi(\cdot)$ is the cumulative distribution function of normal distribution):

$$\psi(e,z) = \frac{1}{\bar{e} \times \left(\Phi\left(\frac{\bar{z}-z_m}{\sigma_\psi}\right) - \Phi\left(\frac{\underline{z}-z_m}{\sigma_\psi}\right)\right)} \frac{1}{\sqrt{2\pi\delta_\psi^2}} \exp\left(-\frac{(z-z_m)^2}{2\delta_\psi^2}\right) \times \mathbf{1}_{e<\bar{e}} \times \mathbf{1}_{z\in[\underline{z},\bar{z}]}.$$

Using the free-entry condition, we simultaneously solve the model and estimate r^l. When training the PDE, we used loss weights of 10^4 for the PDE residual, 10^0 for the free-entry condition, and 10^1 for the

Neumann boundary conditions. We again enforce the Dirichlet boundary condition through an output transform. Furthermore, we trained with 2^{16} training points sampled inside the domain, 2^{10} training points sampled on the boundary, and 10000 points sampled inside the domain for testing. Lastly, we used a learning rate scheduler, with an initial learning rate of 0.0001, decay rate 1.0, and decay step 2000.

The training results are displayed in Figure 3 while the predicted value of r^l after training is displayed in Table 3. The training loss decreases steadily (Fig. 3A), and the L^2 relative error ends at 1.99% (Fig. 3B). As shown in Figs. 3D and E, the PINN predicts the solution to the equation relatively accurately. The absolute difference between the PINN prediction and reference solution is displayed in the error plot in Fig. 3F.

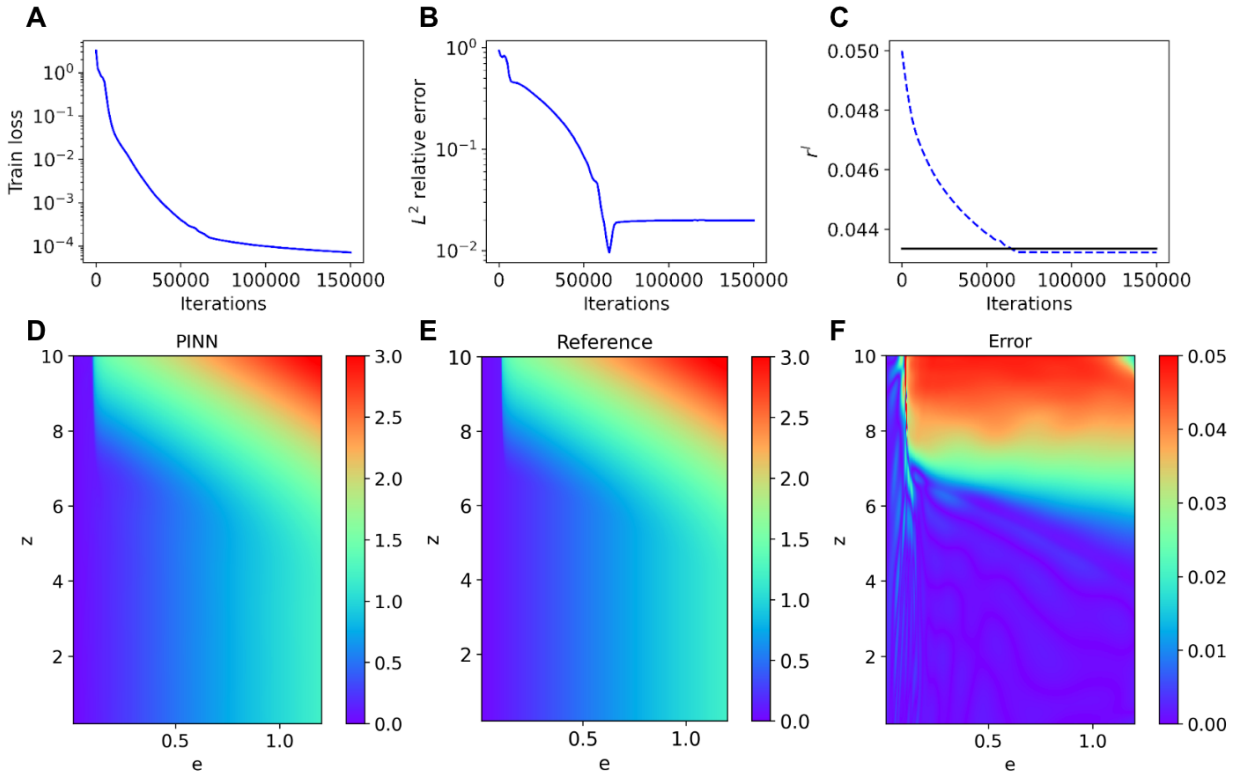

Figure 3: **Example in Section 3.3: PINN results for inverse HJB equation with r^l estimation.** (**A**) The trajectory of train loss throughout training. (**B**) The trajectory of the L^2 relative error throughout training. (**C**) The convergence of r^l throughout training. (**D**) 2D heat map for predicted solution for v using PINN. (**E**) 2D heat map for reference from Matlab data. (**F**) 2D heat map for the absolute difference between PINN prediction and Matlab reference data.

Table 3: **Example in Section 3.3: Predicted value of r^l after training.**

Variable	Ref	Prediction
r^l	0.04307	0.04334

4 A Macroeconomic Model with the Financial Sector

The next problem we discuss describes the model in Brunnermeier and Sannikov (2014). The problem features occasional binding constraints, non-linear financial amplifications, and singularity at the boundary. We describe the model only briefly – see Brunnermeier and Sannikov (2014) for the detailed model setup.

4.1 Problem setup

There are two types of agents, a continuum mass of bankers and a continuum mass of households, both with risk-neutral utility. There are two types of assets: productive capital and risk-free asset. Per unit of capital, banker productivity is a while household productivity is $\underline{a} < a$. Both bankers and households can borrow and lend at the risk-free rate. Banker discount rate δ is smaller than the household discount rate, which motivates them to borrow from households.

To solve for the equilibrium, we need to solve the value of capital $q(\omega)$ and the endogenous marginal value of wealth $\theta(\omega)$ for the banker, where ω is the fraction of the banker wealth among total wealth and is the state variable that determines asset prices and allocations. There is an endogenous boundary η^* where banks pay dividends. The function $q(\eta)$ is increasing and $\theta(\eta)$ is decreasing over $[0, \eta^*]$, with boundary conditions

$$q(0) = \underline{q}, \quad \theta(\eta^*) = 1, \quad q'(\eta^*) = 0, \quad \theta'(\eta^*) = 0, \quad \text{and} \quad \lim_{\eta \to 0} \theta(\eta) = \infty,$$

where the boundary value $\underline{q}$ is determined by

$$\underline{q} = \max_{q} \frac{\underline{a} - \iota(x)}{r - (\Phi(x) - \underline{\delta})}. \tag{1}$$

Note that the maximum possible value of $q(\eta)$ is determined by the following equation:

$$\frac{a - \iota(\bar{q})}{\bar{q}} + \Phi(\bar{q}) - \delta - r = 0,$$

where $\iota(\cdot)$ and $\Phi(\cdot)$ are given by

$$\Phi(x) = \frac{x - 1}{\kappa} \quad \text{and} \quad \iota(x) = \Phi(x) + \frac{1}{2}\kappa\Phi(x)^2.$$

Functions $q(\eta)$ and $\theta(\eta)$ are second-order ODEs. Thus, we need two boundary conditions for $q(\eta)$ and two boundary conditions for $\theta(\eta)$. The above are five conditions, but we also have an unknown boundary η^*.

To solve the model, we first find $\psi \in (\eta, \eta + q(\eta)/q'(\eta))$, such that

$$\frac{a - \underline{a}}{q(\eta)} + \underline{\delta} - \delta + (\sigma + \sigma^q(\eta))\sigma^\theta(\eta) = 0, \tag{2}$$

where

$$\sigma^\eta(\eta)\eta = \frac{(\psi - \eta)\sigma}{1 - (\psi - \eta)q'(\eta)/q(\eta)}, \quad \sigma^q(\eta) = \frac{q'(\eta)}{q(\eta)}\sigma^\eta(\eta)\eta, \quad \text{and} \quad \sigma^\theta(\eta) = \frac{\theta'(\eta)}{\theta(\eta)}\sigma^\eta(\eta)\eta.$$

We can easily prove that the left side of equation (2) is monotonic in ψ and has a unique solution. If the above solution indicates $\psi > 1$, then we set $\psi = 1$ and recalculate $\sigma^\eta(\eta)$, $\sigma^q(\eta)$, and $\sigma^\theta(\eta)$.

Next, we compute the second-order derivatives, via

$$q''(\eta) = \frac{2\left(\mu^q(\eta)q(\eta) - q'(\eta)\mu^\eta(\eta)\eta\right)}{\sigma^\eta(\eta)^2\eta^2}, \tag{3}$$

$$\theta''(\eta) = \frac{2\left(\mu^\theta(\eta)\theta(\eta) - \theta'(\eta)\mu^\eta(\eta)\eta\right)}{\sigma^\eta(\eta)^2\eta^2} \tag{4}$$

where

$$\mu^\eta(\eta) = -\frac{(\psi - \eta)(\sigma + \sigma^q(\eta))(\sigma + \sigma^q(\eta) + \sigma^\theta(\eta))}{\eta} + \frac{a - \iota(q(\eta))}{q(\eta)} + (1 - \psi)(\underline{\delta} - \delta),$$

$$\mu^q(\eta) = r - \frac{a - \iota(q(\eta))}{q(\eta)} - \Phi(q(\eta)) + \delta - \sigma\sigma^q(\eta) - \sigma^\theta(\eta) \cdot (\sigma + \sigma^q(\eta)), \quad \text{and} \quad \mu^\theta = \rho - r.$$

Specification: Model parameter specifications are shown in Table 4 and additional problem setup details and derivations can be found in Appendix B.

Table 4: **Parameter specification of model in Section 4.1**

Description	Value
Capital productivity (experts)	$a = 0.11$
Capital productivity (households)	$\underline{a} = 0.05$
Discount rate (experts)	$\rho = 0.06$
Discount rate (households)	$r = 0.06$
Unit volatility of capital	$\sigma = 0.025$
Capital depreciation rate (experts)	$\delta = 0.03$
Capital depreciation rate (households)	$\underline{\delta} = 0.08$
Loss factor in investment	$\kappa = 10$

4.2 Solving the Model

Our goal is to solve for the functions q, θ, and ψ. To do this, we use an assortment of techniques.

4.2.1 Change of variable to deal with singularity

First, when training the model, we perform a change of variable for θ. We define the function $\hat{\theta}(\eta) = \frac{1}{\theta(\eta)}$. The motivation behind this is to deal with the singularity for θ at $\eta = 0$. With this change of variable, we rewrite all the boundary conditions:

$$\begin{aligned}\lim_{\eta\to 0}\theta(\eta) = \infty \longrightarrow \hat{\theta}(0) = 0,\\ \theta(\eta^*) = 1 \longrightarrow \hat{\theta}(\eta^*) = 1,\\ \theta'(\eta^*) = 0 \longrightarrow \hat{\theta}'(\eta^*) = 0.\end{aligned}$$

We now rewrite the ODEs in the problem setup with the new variable $\hat{\theta}$. We may compute

$$\sigma^\theta(\eta)\hat{\theta}(\eta) = -\hat{\theta}(\eta)'\sigma^\eta(\eta)\eta.$$

This allows us to rewrite (2) as

$$\left(\frac{a-\underline{a}}{q} + \underline{\delta} - \delta\right)\hat{\theta}(\eta) + (\sigma + \sigma_q)\sigma^\theta(\eta)\hat{\theta}(\eta) = 0.$$

We also rewrite these definitions:

$$\begin{aligned}\mu^\eta(\eta)\eta\hat{\theta}(\eta) &= -(\psi-\eta)(\sigma+\sigma^q)(\hat{\theta}(\eta)(\sigma+\sigma^q) + \sigma^\theta(\eta)\hat{\theta}(\eta)) + \eta\hat{\theta}(\eta)\left(\frac{a-\iota}{q} + (1-\psi)(\underline{\delta}-\delta)\right),\\ \mu^q(\eta)\hat{\theta}(\eta) &= \hat{\theta}(\eta)\left(r - \frac{a-\iota(q(\eta))}{q} - \Phi(q(\eta)) + \delta - \sigma\sigma^q\right) - \sigma^\theta(\eta)\hat{\theta}(\eta)\cdot(\sigma+\sigma^q(\eta)).\end{aligned}$$

Lastly, (3) and (4) are now rewritten as

$$\begin{aligned}q''(\eta)\sigma^\eta(\eta)^2\eta^2\hat{\theta}(\eta) &= 2(\mu^q(\eta)\hat{\theta}(\eta)q(\eta) - q'(\eta)\mu^\eta(\eta)\eta\hat{\theta}(\eta),\\ \sigma^\eta(\eta)^2\eta^2\cdot(2\hat{\theta}'(\eta)^2 - \hat{\theta}(\eta)\hat{\theta}''(\eta)) &= 2(\mu^\theta\hat{\theta}(\eta)^2 + \hat{\theta}'(\eta)\mu^\eta(\eta)\hat{\theta}(\eta).\end{aligned}$$

4.2.2 Explicitly solving for ψ

While training, we simultaneously solve for q and $\hat{\theta}$. We train ψ as an auxiliary function that updates every 1000 iterations, given that q and $\hat{\theta}$ are increasing. Our initial guess for ψ is $\psi(\eta) = 1$. This guess for ψ

is flexible since it is updated once q and $\hat{\theta}$ are increasing, which generally happens quickly. To compute ψ when training, we solve for the function explicitly. We have

$$\begin{aligned}\sigma^q(\eta) &= \frac{q'(\eta)}{q(\eta)}\sigma^\eta(\eta)\eta \\ &= \frac{q'(\eta)(\psi-\eta)\sigma}{q(\eta)-(\psi-\eta)q'(\eta)}.\end{aligned}$$

Then,

$$\begin{aligned}\sigma^q(\eta)+\sigma &= \frac{q'(\eta)(\psi-\eta)\sigma}{q(\eta)-(\psi-\eta)q'(\eta)}+\sigma \\ &= \frac{q'(\eta)(\psi-\eta)\sigma+\sigma q(\eta)-\sigma(\psi-\eta)q'(\eta)}{q(\eta)-(\psi-\eta)q'(\eta)} \\ &= \frac{\sigma q(\eta)}{q(\eta)-(\psi-\eta)q'(\eta)}.\end{aligned}$$

Furthermore,

$$\begin{aligned}(\sigma+\sigma^q(\eta))\sigma^\theta(\eta) &= \left(\frac{\sigma q(\eta)}{q(\eta)-(\psi-\eta)q'(\eta)}\right)\frac{\theta'(\eta)(\psi-\eta)\sigma}{\theta(\eta)(1-(\psi-\eta)q'(\eta)/q(\eta))} \\ &= \left(\frac{\sigma q(\eta)}{q(\eta)-(\psi-\eta)q'(\eta)}\right)\frac{q(\eta)\theta'(\eta)(\psi-\eta)\sigma}{\theta(\eta)(q(\eta)-(\psi-\eta)q'(\eta))} \\ &= \frac{\theta'(\eta)}{\theta(\eta)}\left(\frac{(\sigma q(\eta))^2(\psi-\eta)}{(q(\eta)-(\psi-\eta)q'(\eta))^2}\right).\end{aligned}$$

The condition on ψ becomes

$$\begin{aligned}&\frac{a-\underline{a}}{q(\eta)}+\underline{\delta}-\delta+(\sigma+\sigma^q(\eta))\sigma^\theta(\eta)=0 \\ \implies &\frac{a-\underline{a}}{q(\eta)}+\underline{\delta}-\delta+\frac{\theta'(\eta)}{\theta(\eta)}\left(\frac{(\sigma q(\eta))^2(\psi-\eta)}{(q(\eta)-(\psi-\eta)q'(\eta))^2}\right)=0 \\ \implies &\theta(\eta)(q(\eta)-(\psi-\eta)q'(\eta))^2(a-\underline{a}+q(\eta)(\underline{\delta}-\delta))+\sigma^2q(\eta)^3(\psi-\eta)\theta'(\eta)=0,\end{aligned}$$

from which we may solve for ψ using the quadratic formula, after performing the change of variable $\hat{\theta}(\eta)=\frac{1}{\theta(\eta)}$.

4.2.3 Other technical details

We explicitly calculate $\underline{q}$. Solving (1), we find that $\underline{q}=0.4862$.

We employ several additional techniques when constructing the model to improve its results. To satisfy the boundary conditions, we enforce hard constraints via an output transform in addition to utilizing soft constraints (Lu et al., 2021b). We construct the solution

$$\begin{aligned}q(\eta) &= (\eta-\eta^*)^2\mathcal{N}_q(\eta)+C, \\ \hat{\theta}(\eta) &= \eta(\eta-\eta^*)^2\mathcal{N}_{\hat{\theta}}(\eta)-(\eta/\eta^*)^2+2\eta/\eta^*,\end{aligned}$$

where C is a trainable variable that represents $q(\eta^*)$ while $\mathcal{N}_q(\eta)$ and $\mathcal{N}_{\hat{\theta}}(\eta)$ are the first and second components of the network output, respectively. Note that this automatically satisfies the constraints $q'(\eta^*)=0$, $\hat{\theta}(0)=0$, $\hat{\theta}(\eta^*)=1$, and $\hat{\theta}'(\eta^*)=0$. For the final boundary condition $q(0)=\underline{q}$, we use a soft constraint and incorporate it into the loss function.

As discussed, we train two variables $C=q(\eta^*)$ and η^*, and we set initial guesses of $C=1.0$ and $\eta^*=0.4$. While the value of C is flexible, the initial guess for η^* comes from experimentation.

Lastly, initial testing suggested that $q(\eta)$ features a steep gradient for small η. As a result, we employ two techniques to capture this gradient better. We first add a feature layer via an input transform (Yazdani et al., 2020). The feature layer used enables us to stretch the input domain making it easier to capture the sharp increase for small η. Another technique used involves adding an additional ten anchor training points in the interval $[0, 10^{-4}]$.

Next, we want to ensure that q is increasing and θ is decreasing. By the change of variable, θ decreasing is equivalent to $\hat{\theta}$ increasing, so to enforce q increasing and $\hat{\theta}$ increasing, we add loss terms

$$\mathcal{L}_q = \int_0^{\eta^*} (\min(q'(\eta), 0))^2 d\eta \quad \text{and} \quad \mathcal{L}_{\hat{\theta}} = \int_0^{\eta^*} (\min(\hat{\theta}'(\eta), 0))^2 d\eta,$$

representing loss from the mean squared value of $\min(q'(\eta), 0)$ and $\min(\hat{\theta}'(\eta), 0)$. This pushes the functions to have positive derivatives, which would mean they are increasing. We used loss weights of 10^6 for the ODE system, q increasing, and $\hat{\theta}$ increasing losses, and a loss weight of 1 for the boundary condition $q(0) = \underline{q}$.

While training, we use two different learning rate schedulers. The first learning rate scheduler has an initial learning rate of 0.001 with a decay rate of 0.5 and decay step of 1500. We use this scheduler for the first 100000 iterations. Then, we switch to a learning rate scheduler with initial learning rate 0.00001, decay rate 0.5, and decay step 1500 for another 200000 iterations.

The training results are displayed in Figure 4 while the predicted value of η^* is displayed in Table 5. When calculating the L^2 error for these results, we only consider the portion of the graph where $\eta < \eta^*$. This is an accurate measurement since the final result features a prediction for η^* that is very close to the reference value. The prediction for q (Fig. 4A) has an L^2 relative error of 1.14%. The prediction for θ (Fig. 4B) has an L^2 relative error of 0.258%. Finally, the prediction for ψ (Fig. 4C) has an L^2 relative error of 0.273%.

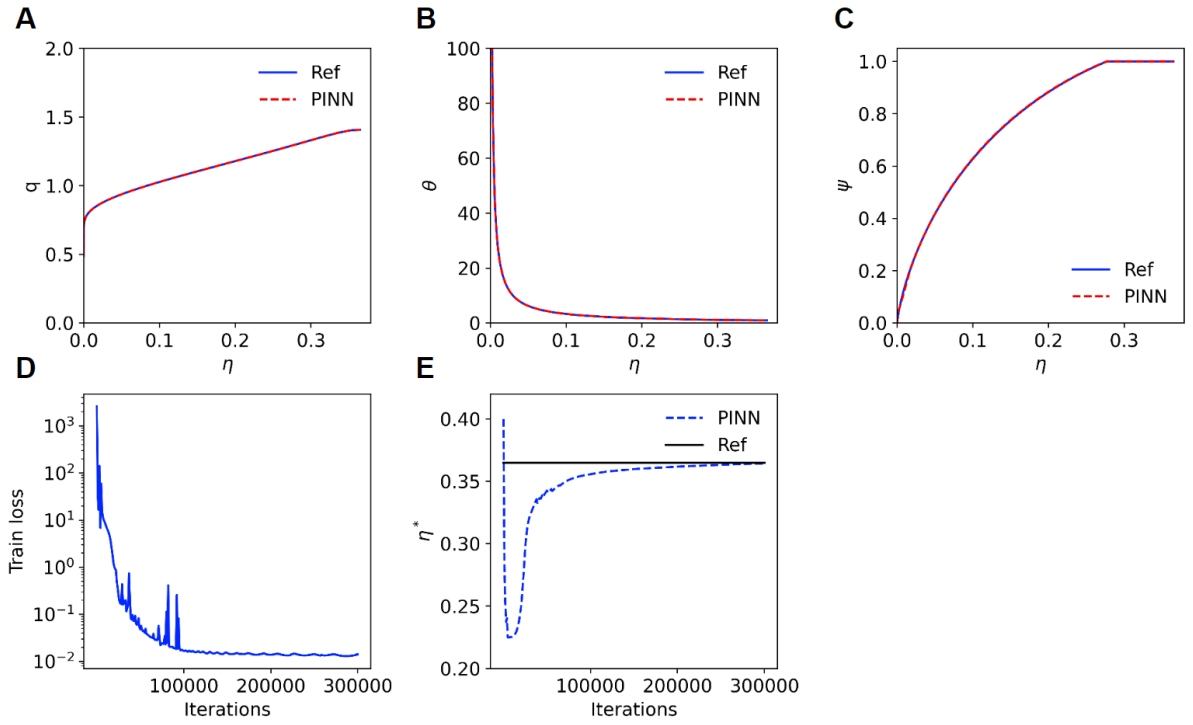

Figure 4: **Example in Section 4.1: PINN results for forward B&S model.** (**A**) The reference value and prediction for q after training. (**B**) The reference value and prediction for θ after training. (**C**) The reference value and prediction for ψ after training. (**D**) The trajectory of train loss throughout training. (**E**) The convergence of η^* throughout training.

Table 5: **Example in Section 4.1: Predicted value of η^* after training.**

Variable	Ref	Prediction
η^*	0.3648	0.3644

4.3 Simultaneously Solving and Estimating the Model

In this section, we simultaneously solve the model and also estimate one or two parameters of the model.

4.3.1 Estimation of one parameter

First, we estimate only one parameter. Specifically, we assume a is unknown and impose a moment condition. There are two ways to match the moment condition.

The first method is simulation. We simply simulate η_t according to its law of motion,

$$\frac{d\eta_t}{\eta_t} = \mu^\eta(\eta_t)dt + \sigma^\eta(\eta_t)dB_t,$$

and then take the average moment value to match the target.

A more efficient method is through the Kolmogorov forward equation. We can solve for the Kolmogorov forward equation that describes the ODE of the stationary density function of η_t, $f(\eta)$. The equation is

$$-\frac{\partial}{\partial\eta}\left(\mu^\eta(\eta)\eta f(\eta)\right) + \frac{\partial^2}{\partial\eta^2}\left(\frac{1}{2}\sigma^\eta(\eta)^2\eta^2 f(\eta)\right) = \frac{\partial}{\partial\eta}J(\eta) = 0,$$

with boundary condition $-\mu^\eta\eta f(\eta) + \frac{1}{2}\frac{\partial}{\partial\eta}\left(\sigma^\eta(\eta)^2 f(\eta)\right)|_{\eta=\eta^*} = 0$. For simplicity, we define Q function as

$$Q(\eta) = \frac{2\mu^\eta(\eta)\eta - \frac{\partial}{\partial\eta}(\sigma^\eta(\eta)^2\eta^2)}{\sigma^\eta(\eta)^2\eta^2}.$$

Then, according to Appendix C, the density function can be written as

$$f(\eta) = A\exp\left(-\int_\eta^{\eta^*} Q(\eta')d\eta'\right),$$

where A is the normalization factor:

$$\frac{1}{A} = \int_0^{\eta^*}\exp\left(-\int_{\eta'}^{\eta^*} Q(\eta)d\eta\right)d\eta'.$$

We match the moment target

$$a_{\text{target}} = \int_0^{\eta^*}[\psi(\eta)a + (1-\psi(\eta))\underline{a}]f(\eta)d\eta$$

where we define $a_{\text{target}} = 0.1095$. In this paper, we use method 2, employing the Kolmogorov forward equation.

The detailed training results for one trial is displayed in Figure 5. The figure shows that the PINN achieves accurate estimations of q, θ, ψ, and a. Furthermore, running several trials, we find that the solution consistently converges to the global minimum no matter the initial guess of a. To demonstrate this consistent convergence, we run ten trials in which we randomly select a from the range $(0.06, 0.16)$. The parameter a is accurately estimated, and the final value is $a = 0.1099 \pm 0.0007$ (Table 6) with L^2 errors of less than 1% for q, θ, and ψ.

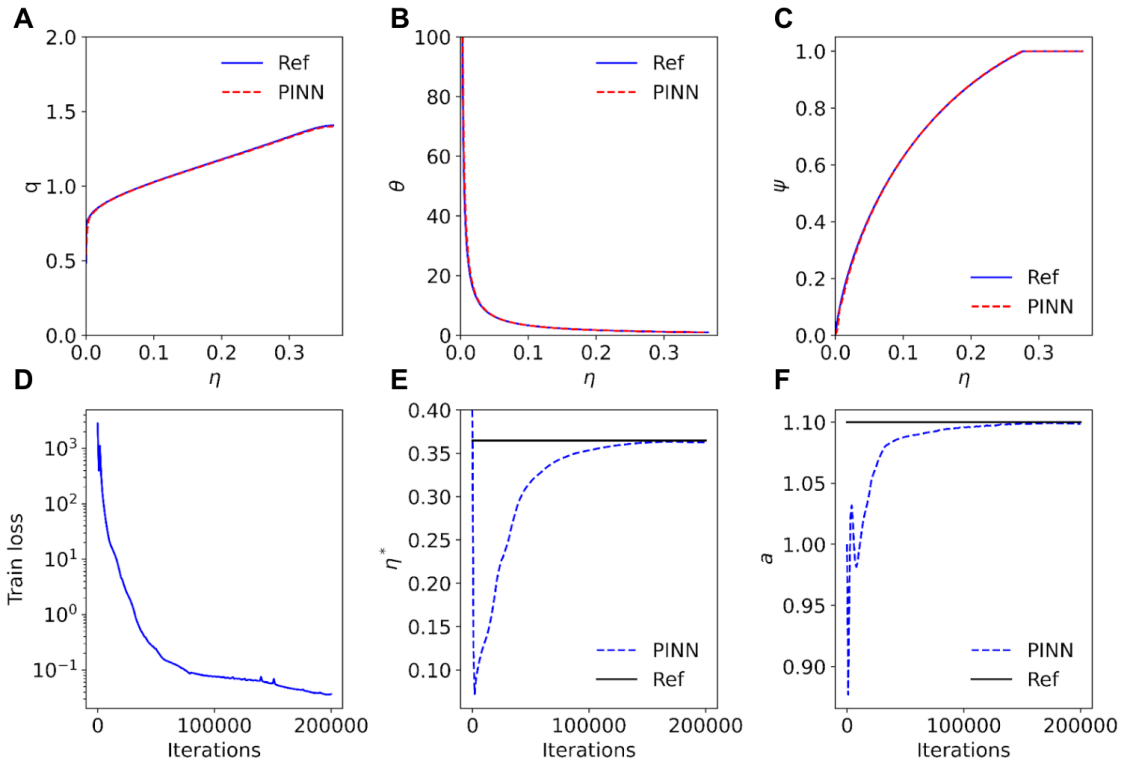

Figure 5: **Example in Section 4.3.1: PINN results for inverse B&S model with one unknown parameter.** (**A**) The reference value and prediction for q. (**B**) The reference value and prediction for θ. (**C**) The reference value and prediction for ψ. (**D**) The trajectory of train loss throughout training. (**E**) The convergence of η^* throughout training. (**F**) The convergence of a throughout training.

Table 6: **Example in Section 4.3.1: Results from several trials of inverse B&S model with one unknown parameter.** The choices of a are random in the range $(0.06, 0.16)$, and the table depicts the final results for each trial.

Trial	a		a_{target}		L^2 relative errors		
	Initial	Final	Value	Error	q	θ	ψ
1	0.1500	0.1098	0.1093	0.16%	0.17%	0.53%	0.06%
2	0.0921	0.1098	0.1093	0.22%	0.22%	0.69%	0.08%
3	0.1472	0.1100	0.1095	0.01%	0.01%	0.02%	0.00%
4	0.1230	0.1099	0.1094	0.11%	0.12%	0.36%	0.04%
5	0.1236	0.1098	0.1094	0.14%	0.14%	0.45%	0.05%
6	0.1005	0.1098	0.1093	0.22%	0.23%	0.71%	0.09%
7	0.1509	0.1099	0.1094	0.08%	0.08%	0.26%	0.03%
8	0.1415	0.1099	0.1094	0.07%	0.07%	0.21%	0.03%
9	0.0889	0.1099	0.1094	0.09%	0.09%	0.29%	0.03%
10	0.0645	0.1098	0.1093	0.15%	0.16%	0.49%	0.06%

4.3.2 Estimation of two parameters

Now, we estimate parameters a and σ. We define our moment conditions:

$$a_{\text{target}} = \int_0^{\eta^*} [\psi(\eta)a + (1-\psi(\eta))\underline{a}]f(\eta)d\eta \quad \text{and} \quad \text{Vol}_{\text{target}} = \frac{\sqrt{\int_0^{\eta^*}(q-\bar{q})^2 f(\eta)d\eta}}{\bar{q}},$$

where $\bar{q}$ is defined as $\int_0^{\eta^*} q(\eta)f(\eta)d\eta$, $a_{\text{target}} = 0.1095$, and $\text{Vol}_{\text{target}} = 0.04126$. Again, we employ the Kolmogorov forward equation.

To implement the integrals in the moment conditions, we use numerical integration. When creating the model, we use the same change of variable for θ. However, instead of explicitly solving for ψ, noting that the left hand side of equation (2) is monotonic, we utilize a basic bisection algorithm provided by the Python library scipy to solve for ψ. While training, we scale a up by 10 and σ up by 100, and we scale them back down for the results. We also use loss weights: 10^6 for the ODE on q'', $2 \cdot 10^5$ for the ODE on θ'', 10^6 for q increasing, 10^4 for θ decreasing, 10^5 for the a_{target} moment condition, $5 \cdot 10^3$ for the $\text{Vol}_{\text{target}}$ moment condition, and 10^1 for the $q(0) = \underline{q}$ boundary condition. Lastly, with the learning rate scheduler, we have $5 \cdot 10^{-4}$ as the initial learning rate, 1.0 as the decay rate, and 1500 as the decay step. We train this model for $1.5 \cdot 10^5$ iterations.

With our model, we find that the convergence and final values are sensitive to the initial guess of σ but not to the initial guess of a. Furthermore, we find that the prediction of a is consistently accurate, regardless of the initial guess. Although the predictions of σ are less accurate, there are initial guesses that result in accurate predictions of σ. An example is demonstrated in Figure 6. In this example, the train loss decreases steadily (Fig. 6A). Furthermore, with the final predicted values for a and σ, the predictions of a_{target} and $\text{Vol}_{\text{target}}$ are very close (Figs. 6D and E). Similarly, the graphs of q, θ, and ψ are very close to their true values (Figs. 6F, G, and H).

To illustrate this convergence, the trajectories of five trials that converge to the global minimum are displayed in Figure 7A. Furthermore, the trajectories of five trials that converge to a local minimum instead are displayed in Fig. 7B.

In practice, trials that converge to a local minimum can be identified because a_{target} and $\text{Vol}_{\text{target}}$ are farther from the true values (Table 7). We consider the results with three runs of ten random trials each, in which we select a and σ at random in the ranges $(0.6, 1.6)$ and $(0.01, 0.04)$. We generally receive satisfactory results (Fig. 8A). For the next runs, we run five random trials and then apply Bayesian optimization to determine the next five initial guesses. Bayesian optimization is a method that uses Bayes' theorem to maximize a continuous function without information about derivatives (Snoek et al., 2012). Through this, we find that Bayesian optimization consistently produces better results than random trials (Fig. 8B).

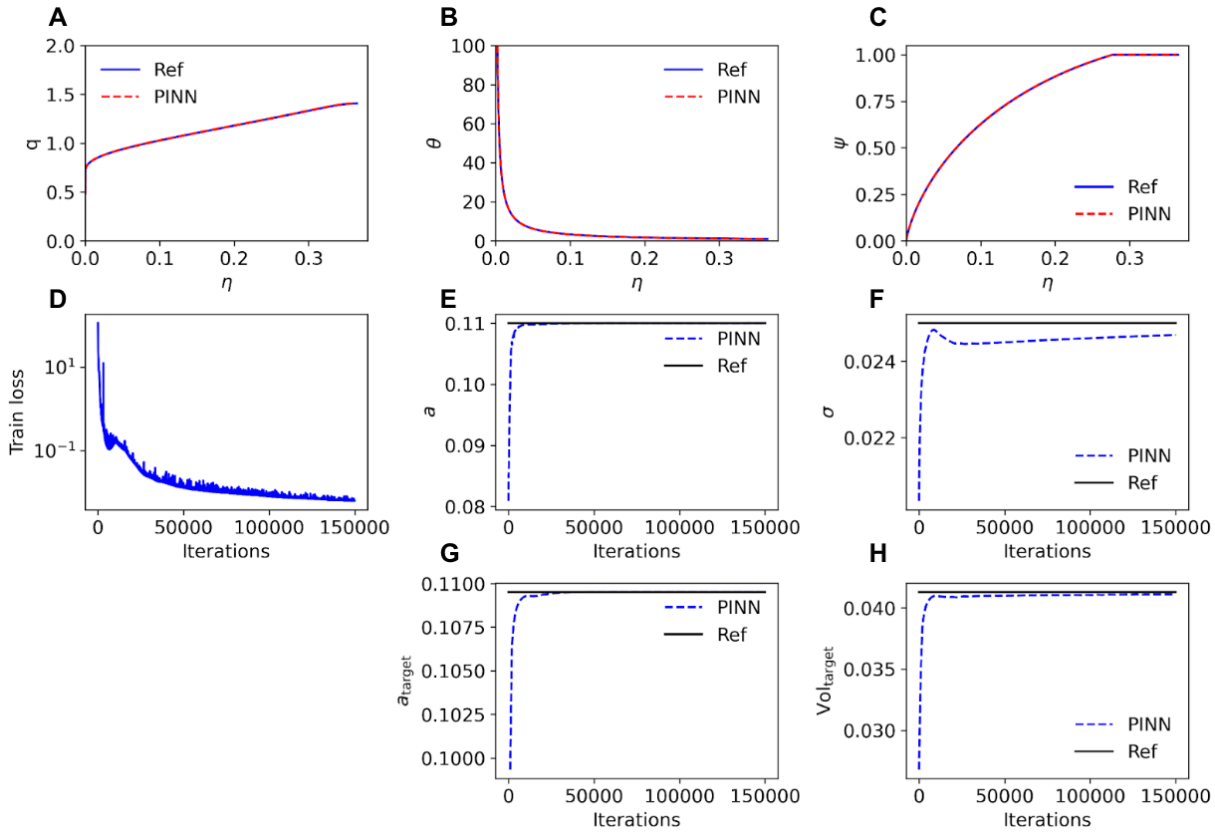

Figure 6: **Example in Section 4.3.2: PINN results for inverse B&S model with two unknown parameters.** (**A**) The reference value and prediction for q. (**B**) The reference value and prediction for θ. (**C**) The reference value and prediction for ψ. (**D**) The trajectory of train loss throughout training. (**E**) The convergence of a throughout training. (**F**) The convergence of σ throughout training. (**G**) The trajectory of a_{target} throughout training. (**H**) The trajectory of $\text{Vol}_{\text{target}}$ throughout training.

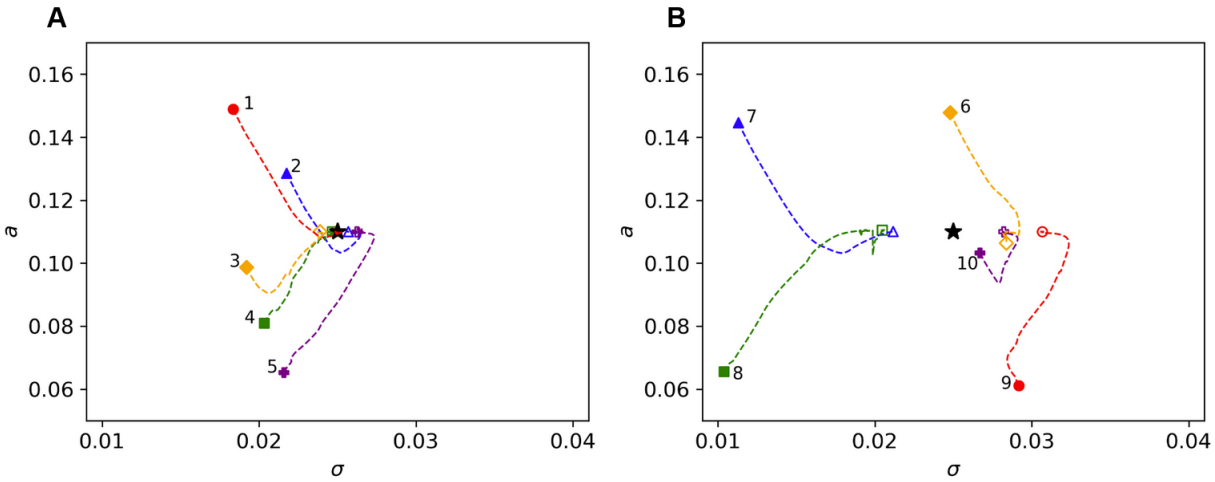

Figure 7: **Example in Section 4.3.2: Trajectories in the parameter space** (σ, a)**.** The black star represents the true value (σ, a), the shaded-in shape represents the initial guess, the empty shape represents the final estimation, and the path connecting them is the trajectory. (**A**) Five trajectories that converged to the global minimum. (**B**) Five trajectories that converged to a local minimum.

Table 7: **Example in Section 4.3.2: PINN parameter estimation from different initial values of a and σ.** Trials 1–5 converge to the global minimum. Trials 6–10 converge to a local minimum.

Trial	Initial values		Final values		a_{target}		$\text{Vol}_{\text{target}}$		L^2 relative errors		
	a	σ	a	σ	Value	Error	Value	Error	q	θ	ψ
1	0.1489	0.0183	0.1100	0.0249	0.1095	0.00%	0.0412	0.19%	0.06%	0.10%	0.05%
2	0.1286	0.0217	0.1100	0.0257	0.1095	0.01%	0.0416	0.88%	0.33%	0.56%	0.22%
3	0.0809	0.0203	0.1100	0.0247	0.1095	0.01%	0.0411	0.37%	0.15%	0.26%	0.10%
4	0.0987	0.0192	0.1100	0.0239	0.1095	0.01%	0.0407	1.43%	0.55%	0.87%	0.37%
5	0.0653	0.0216	0.1099	0.0262	0.1094	0.07%	0.0419	1.47%	0.65%	1.20%	0.38%
6	0.0612	0.0292	0.1100	0.0307	0.1096	0.04%	0.0392	5.00%	1.96%	2.79%	1.31%
7	0.1446	0.0113	0.1100	0.0212	0.1100	0.42%	0.0391	5.23%	2.71%	4.45%	1.41%
8	0.0655	0.0104	0.1104	0.0205	0.1059	3.32%	0.0408	1.23%	4.71%	12.35%	0.35%
9	0.1478	0.0248	0.1064	0.0284	0.1095	0.03%	0.0429	4.05%	1.50%	2.71%	1.03%
10	0.1033	0.0267	0.1100	0.0282	0.1094	0.06%	0.0442	7.09%	2.60%	4.97%	1.79%

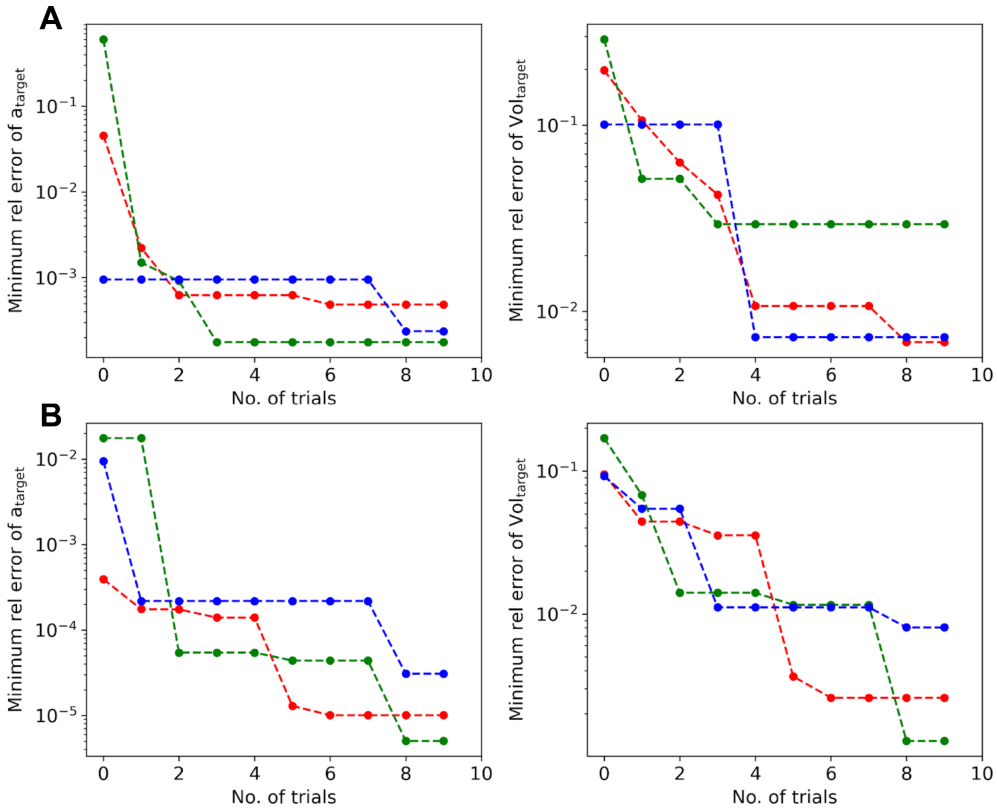

Figure 8: **Example in Section 4.3.2: PINN results from varying initial guesses using Bayesian optimization.** The results from three runs of ten initial guesses, plotting the minimum relative error of a_{target} and $\text{Vol}_{\text{target}}$ so far. (**A**) The initial guesses of a and σ are chosen randomly from the intervals $(0.06, 0.16)$ and $(0.01, 0.04)$, respectively. (**B**) The first five initial guesses of a and σ are chosen once again at random from the intervals $(0.06, 0.16)$ and $(0.01, 0.04)$, respectively. The five guesses afterwards are chosen using Bayesian optimization.

5 Conclusion

In this paper, we proposed a new method for solving economics models with deep learning through physics-informed neural networks (PINNs). We demonstrated the advantages of this method: generality, simultaneous solution and estimation, leveraging the state-of-art machine-learning techniques, and handling large

state space. Our proposed method provides a general framework for solving PDEs, whereas in traditional methods, we usually need to design different algorithms for different problems. Furthermore, the inverse problem can be solved very easily and is no more than adding an extra loss term for additional information. We showed the effectiveness of the PINN method in solving two models, each featuring its own challenges.

There are a couple of directions for future research. First, currently we only allow for heterogeneous agent with stationary population distribution. Our method can be generalized to deal with mean-field games where the distribution is a state variable that changes over time. Second, we currently only solve exactly identified estimation problems where the number of moment conditions is equal to the number of parameters to be estimated. A more generalized method should be able to handle overidentified systems and provide error bounds on the estimated coefficients, similar to the classical GMM estimations. Finally, our method can be applied to a broader variety of economics and finance problems.

Acknowledgements

B.F., E.Q., and L.L thank MIT's PRIMES-USA program.

References

Bianchi, J. and Bigio, S. (2018). Banks, liquidity management and monetary policy. *NBER Working Paper No. 20490*.

Brunnermeier, M. K. and Sannikov, Y. (2014). A macroeconomic model with a financial sector. *American Economic Review*, 104(2):379–421.

Byrd, R. H., Lu, P., Nocedal, J., and Zhu, C. (1995). A limited memory algorithm for bound constrained optimization. *SIAM Journal on Scientific Computing*, 16(5):1190–1208.

Chen, H., Didisheim, A., and Scheidegger, S. (2021). Deep structural estimation: With an application to option pricing. *arXiv preprint arXiv:2102.09209*.

Chen, Y., Lu, L., Karniadakis, G. E., and Dal Negro, L. (2020). Physics-informed neural networks for inverse problems in nano-optics and metamaterials. *Optics Express*, 28(8):11618–11633.

Di Tella, S. (2019). Optimal regulation of financial intermediaries. *American Economic Review*, 109(1):271–313.

Drechsler, I., Savov, A., and Schnabl, P. (2014). A model of monetary policy and risk premia. *Journal of Finance*, pages n/a–n/a.

Duarte, V. (2018). Machine learning for continuous-time economics. *Available at SSRN 3012602*.

Fernandez-Villaverde, J., Nuno, G., Sorg-Langhans, G., and Vogler, M. (2020). Solving high-dimensional dynamic programming problems using deep learning. *Unpublished working paper*.

Gabaix, X., Lasry, J.-M., Lions, P.-L., and Moll, B. (2016). The dynamics of inequality. *Econometrica*, 84(6):2071–2111.

Gertler, M. and Kiyotaki, N. (2015). Banking, liquidity, and bank runs in an infinite horizon economy. *The American Economic Review*, 105(7):2011–2043.

Gomes, J., Kogan, L., and Zhang, L. (2003). Equilibrium cross section of returns. *Journal of Political Economy*, 111(4):693–732.

He, Z. and Krishnamurthy, A. (2013). Intermediary asset pricing. *The American Economic Review*, 103(2):732–770.

Hennessy, C. A. and Whited, T. M. (2007). How costly is external financing? evidence from a structural estimation. *The Journal of Finance*, 62(4):1705–1745.

Huang, J. (2018). Banking and shadow banking. *Journal of Economic Theory*, 178:124–152.

Karatzas, I. and Shreve, S. (1998). *Brownian motion and stochastic calculus*, volume 113, *Graduate Texts in Mathematics*. Springer.

Kingma, D. P. and Ba, J. (2014). Adam: A method for stochastic optimization.

Kissas, G., Yang, Y., Hwuang, E., Witschey, W. R., Detre, J. A., and Perdikaris, P. (2020). Machine learning in cardiovascular flows modeling: Predicting arterial blood pressure from non-invasive 4d flow mri data using physics-informed neural networks. *Computer Methods in Applied Mechanics and Engineering*, 358:112623.

Krishnamurthy, A. and Li, W. (2020). Dissecting mechanisms of financial crises: Intermediation and sentiment. Technical report, National Bureau of Economic Research.

Lu, L., Meng, X., Mao, Z., and Karniadakis, G. E. (2021a). DeepXDE: A deep learning library for solving differential equations. *SIAM Review*, 63(1):208–228.

Lu, L., Pestourie, R., Yao, W., Wang, Z., Verdugo, F., and Johnson, S. G. (2021b). Physics-informed neural networks with hard constraints for inverse design. *SIAM Journal on Scientific Computing*, 43(6):B1105–B1132.

Matvos, G. and Seru, A. (2014). Resource allocation within firms and financial market dislocation: Evidence from diversified conglomerates. *The Review of Financial Studies*, 27(4):1143–1189.

Maxted, P. (2020). A macro-finance model with sentiment.

Pavliotis, G. A. (2014). *Stochastic processes and applications: diffusion processes, the Fokker-Planck and Langevin equations*, volume 60. Springer.

Raissi, M., Perdikaris, P., and Karniadakis, G. (2019). Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378:686–707.

Raissi, M., Yazdani, A., and Karniadakis, G. E. (2020). Hidden fluid mechanics: Learning velocity and pressure fields from flow visualizations. *Science*, 367(6481):1026–1030.

Ramachandran, P., Zoph, B., and Le, Q. V. (2017). Searching for activation functions.

Risken, H. (1996). *The Fokker-Planck equation: Methods of solution and applications*, volume 18. Springer.

Sahli Costabal, F., Yang, Y., Perdikaris, P., Hurtado, D. E., and Kuhl, E. (2020). Physics-informed neural networks for cardiac activation mapping. *Frontiers in Physics*, 8.

Snoek, J., Larochelle, H., and Adams, R. P. (2012). Practical bayesian optimization of machine learning algorithms. page 2951–2959.

Tartakovsky, A. M., Marrero, C. O., Perdikaris, P., Tartakovsky, G. D., and Barajas-Solano, D. (2020). Physics-informed deep neural networks for learning parameters and constitutive relationships in subsurface flow problems. *Water Resources Research*, 56(5):e2019WR026731.

Whited, T. M. and Wu, G. (2006). Financial constraints risk. *The review of financial studies*, 19(2):531–559.

Yazdani, A., Lu, L., Raissi, M., and Karniadakis, G. E. (2020). Systems biology informed deep learning for inferring parameters and hidden dynamics. *PLoS Computational Biology*, 16(11):e1007575.

Yu, J., Lu, L., Meng, X., and Karniadakis, G. E. (2021). Gradient-enhanced physics-informed neural networks for forward and inverse PDE problems. *arXiv preprint arXiv:2111.02801*.

Appendices

A Setup and Derivations for the Firm Industrial Dynamic Model

The appendix will contain detailed model setup and derivations for each problem.

Proposition 1. *In bank industrial dynamic models, stationary distribution exists.*

Proof. As the volatility $\sigma(z)$ is constant, and Ornstein–Uhlenbeck process is recurrent, we only need to check if the linear growth condition for (μ_e, μ_z) is satisfied:

$$\begin{aligned}
||\vec{\mu}(e,z)|| \leq ||\mu_e(e,z)|| + ||\mu_z(e,z)|| &= ||\theta_0(z-z_m)|| + ||\pi^*(e,z) - \mathbf{1}_{(e_t,z_t)\in\mathscr{C}^c}\kappa(\phi e - f(z,l^*))^+|| \\
&\leq ||\theta_0(z-z_m)|| + ||\pi^*(e,z)|| + ||\kappa(\phi e - f(z,l^*))^+|| \\
&\leq ||\theta_0(z-z_m)|| + ||\pi^*(e,z)|| + |\kappa|(||(\phi e)|| + ||f(z,l^*)||) \\
&\leq ||\theta_0(z-z_m)|| + ||\pi^*(e,z)|| + |\kappa|(|\phi||x| + ||f(z,l^*)||)
\end{aligned}$$

We can conclude our proof by noting that both instantaneous profit $\pi^*(e,z)$, and deposit production $f(z,l^*)$ are scaled by productivity z. □

B Setup and Derivations for the Macroeconomic Model with a Financial Sector

The model in Brunnermeier and Sannikov (2014) considers an economy populated with one unit of experts (indexed by $i, i \in \mathbb{I} = [0,1]$) and one unit of households (indexed by $j, j \in \mathbb{J} = [1,2]$), in an infinite horizon setup. Time is continuous here. The physical capital k_t held by experts produces output at rate: $y_t = ak_t$, while households produces at rate: $\underline{y}_t = \underline{a}k_t$. Denoting the investment rate per unit of capital as: $\iota_t, \underline{\iota}_t$ for experts and households respectively (i.e., $\iota_t k_t$ is the total investment rate of experts), the corresponding capital evolves as:

$$\begin{cases} dk_t = (\Phi(\iota_t) - \delta)k_t dt + \sigma k_t dZ_t, \\ d\underline{k}_t = (\Phi(\underline{\iota}_t) - \underline{\delta})\underline{k}_t dt + \sigma \underline{k}_t dZ_t, \end{cases}$$

where Z_t here is exogenous aggregate Brownian shocks, $\Phi(\cdot)$ is standard investment technology with convex adjustment costs, i.e., $\Phi(0) = 0, \Phi'(0) = 1, \Phi'(\cdot) > 0$ and $\Phi''(\cdot) < 0$, $\delta, \underline{\delta}$ are depreciation rate (assume $\underline{\delta} > \delta$), and σ is the volatility. Experts and households are risk neutral, households have discount rate r and they may have positive or negative consumption, which means households provide fully elastic lending at risk-free rate r. The equilibrium market price of capital is postulated as a GBM:

$$dq_t = \mu_t^q q_t dt + \sigma_t^q q_t dZ_t,$$

where μ_t can be viewed as the time dependent price drift and σ_t^q is the price volatility. Both μ_t^q and σ_t^q are determined by market equilibrium. We focus on the case that equilibrium price $q \in [\underline{q}, \overline{q}]$, where $\underline{q} = \max_\iota \frac{\underline{a}-\iota}{r-(\Phi(\iota)-\underline{\delta})}$ and $q = \max_\iota \frac{a-\iota}{r-(\Phi(\iota)-\delta)}$. The return of the capital managed by experts and households can then be expressed by:

$$\begin{aligned}
dr_t^k &= \frac{dD_t + d(k_t q_t)}{k_t q_t} = \frac{(a-\iota_t)\ \not{k}_t}{\not{k}_t q_t} + (\Phi(\iota_t) - \delta + \mu_t^q + \sigma\sigma_t^q)\,dt + (\sigma + \sigma_t^q)dZ_t, \\
d\underline{r}_t^k &= \frac{d\underline{D}_t + d(k_t q_t)}{k_t q_t} = \frac{(\underline{a}-\underline{\iota}_t)\ \not{k}_t}{\not{k}_t q_t} + (\Phi(\underline{\iota}_t) - \underline{\delta} + \mu_t^q + \sigma\sigma_t^q)\,dt + (\sigma + \sigma_t^q)dZ_t.
\end{aligned}$$

The cumulative consumption of a household and an expert are denoted as $\underline{c}_t$ and c_t, respectively. Then utilities are given by:

$$\mathbb{E}\left[\int_0^\infty e^{-rt} d\underline{c}_t\right] \text{ (households)} \quad \text{and} \quad \mathbb{E}\left[\int_0^\infty e^{-\rho t} dc_t\right] \text{ (experts)}.$$

The net worth of an expert n_t evolves as:

$$\frac{dn_t}{n_t} = x_t dr_t^k + (1-x_t)rdt - \frac{dc_t}{n_t},\ x_t \text{ is the fraction of capital,}$$

the first part is return on risky assets, the second part is return on safe assets, and the third part is consumption. Similarly, for households, we have:

$$\frac{d\underline{n}_t}{\underline{n}_t} = \underline{x}_t d\underline{r}_t^k + (1-\underline{x}_t)rdt - \frac{d\underline{c}_t}{\underline{n}_t},\ \underline{x}_t \text{ is the fraction of capital.}$$

Both households and experts maximize their utility. Households can have negative consumption while experts cannot, i.e., $dc_t \geq 0$. Households and experts' problems can be written as:

$$\max_{\underline{x}_t \geq 0, d\underline{c}_t, \underline{\iota}_t} \mathbb{E}\left[\int_0^\infty e^{-rt} d\underline{c}_t\right] \text{ (households)},\ \max_{x_t \geq 0, dc_t \geq 0, \iota_t} \mathbb{E}\left[\int_0^\infty e^{-\rho t} dc_t\right] \text{ (experts)},$$

subject to net worth's equation of motion.

Definition for Equilibrium. *Given initial wealth distribution, k_0^i, k_0^j, an equilibrium is described by the stochastic process: $\{q_t, n_t^i, \underline{n}_t^i \geq 0, n_t^i, \underline{k}_t^i \geq 0, \iota_t^i, \underline{\iota}_t^i, dc_t^i \geq 0, d\underline{c}_t^i\}$ such that: (1) initial net worth is $n_0^i = k_0^i q_0, \underline{n}_0^j = k_0^j q_0$; (2) each expert and agent solves their problems, given capital price q_t; (3) markets for consumption goods and capital are clear, i.e.,*

$$\int_{\mathbb{I}} dc_t^i di + \int_{\mathbb{J}} d\underline{c}_t^j dj = \int_{\mathbb{I}} (a - \iota_t^i) k_t^i di + \int_{\mathbb{J}} (\underline{a} - \underline{\iota}_t^j) \underline{k}_t^j dj,$$

$$\int_{\mathbb{I}} k_t^i di + \int_{\mathbb{J}} \underline{k}_t^j dj = K_t, dK_t \text{ is: } \left(\int_{\mathbb{I}} (\Phi(\iota_t^i)) - \delta) k_t^i di + \int_{\mathbb{J}} (\Phi(\underline{\iota}_t^i)) - \underline{\delta}) \underline{k}_t^j dj\right) dt + \sigma K_t dZ_t.$$

Solution. First, households and experts' investment choices $\iota_t, \underline{\iota}_t$ are solved by maximizing $dr_t^k, d\underline{r}_t^k$, which means

$$\iota_t, \underline{\iota}_t \in \arg\max_{\iota} \Phi(\iota) - \iota/q_t \Rightarrow \iota_t = \underline{\iota}_t = \Phi'^{-1}(1/q_t).$$

Second, denoting ψ_t as the fraction of capital held by experts, we are led to equilibrium condition $\mathbb{E}_t[d\underline{r}_t^k]/dt \leq r$, with equality if $1 - \psi_t > 0$. This condition means that when expected return is less than risk-free asset, households will not hold any capital. Also, risk-neutral households will hold a fraction of capital when the return equals the risk-free rate.

Third, to solve the experts' problem, we introduce multiplier θ_t for experts' future utility, i.e., $\theta_t n_t \equiv \mathbb{E}_t\left[\int_0^\infty e^{-\rho(s-t)} dc_s\right]$. Introducing unit worth's consumption $d\zeta_t$, the experts optimal trading strategy is:

$$\rho\theta_t n_t = \max_{\hat{x}_t, d\zeta_t} n_t d\zeta_t + \mathbb{E}d[\theta_t n_t]$$

Considering a finite process $d\theta_t/\theta_t = \mu_t^\theta dt + \sigma_t^\theta dZ_t$, for optimal strategy, the solution of θ_t features: (1) it is always true that $\theta_t \geq 1$, $d\zeta_t > 0$ only when $\theta_t = 1$; (2) $\mu_t = \rho - r$; (3) either $x_t > 0$ when $\mathbb{E}_t[dr_t^k]/dt - r = -\sigma_t^\theta(\sigma + \sigma_t^q)$ (the risk premium), or $x_t = 0$ when $\mathbb{E}_t[dr_t^k]/dt - r < -\sigma_t^\theta(\sigma + \sigma_t^q)$.

Denote the experts' wealth share as $\eta_t \equiv \frac{N_t}{q_t K_t} \in [0,1]$, where all functions can be functions of η_t. By applying Itô's lemma, we get ($\langle \cdot, \cdot \rangle$ is the quadratic variation):

$$d\eta_t = \frac{dN_t}{q_t K_t} + N_t d\left(\frac{1}{q_t K_t}\right) + \left\langle dN_t, d\left(\frac{1}{q_t K_t}\right)\right\rangle \equiv \eta_t(\mu_t^\eta dt + \sigma_t^\eta dZ_t)$$

$$\Rightarrow \mu_t^\eta = -\sigma_t^\eta(\sigma + \sigma_t^q + \sigma_t^\theta) + \frac{a - \iota(q_t)}{q_t} + (1 - \psi_t)(\underline{\delta} - \delta),\ \sigma_t^\eta = \frac{\psi_t - \eta_t}{\eta_t}(\sigma + \sigma_t^q).$$

Equilibrium conditions. Optimal strategies of households and experts imply

$$\frac{\mathbb{E}_t[dr_t^k - d\underline{r}_t^k]}{dt} - \sigma_t^\theta(\sigma + \sigma_t^q) = 0 \Rightarrow \psi(\eta),$$

which means for experts, the opportunity cost of holding capital is the risk premium, if the equilibrium has an interior solution. When the implied solution $\psi > 1$, the above equation does not necessarily hold, as it is always profit profitable to hold capital in this case. Next, by Itô's formula, we can solve the equilibrium price $q(\eta)$ and multiplier θ_t from

$$\begin{cases} \mu_t^q q(\eta) = q'(\eta)\mu_t^\eta \eta + \frac{1}{2}(\sigma_t^\eta)^2\eta^2 q''(\eta), \\ \mu_t^\theta \theta(\eta) = \theta'(\eta)\mu_t^\eta \eta + \frac{1}{2}(\sigma_t^\eta)^2\eta^2\theta''(\eta). \end{cases}$$

Given the stochastic process of $d\eta_t = \eta_t\mu_t^\eta dt + \eta_t\sigma_t^\eta$, the distribution $f(\eta, t)$ evolves as

$$\frac{\partial}{\partial t}f(\eta,t) = -\frac{\partial}{\partial \eta}\left(\mu^\eta f(\eta,t)\right) + \frac{1}{2}\frac{\partial^2}{\partial \eta^2}\left(\sigma_\eta^2(\eta) f(\eta,t)\right).$$

Prior to solving the invariant distribution, we need to show that the distribution function exists[2].

Proposition 2. *In Brunnermeier and Sannikov (2014), the stationary distribution exists, if* $2(\rho - r)\sigma^2 < \Lambda^2, \Lambda = \frac{\mathbb{E}_t[dr_t^k - dr_t^k]}{dt} = \frac{a-\underline{a}}{\underline{q}} - (\delta - \underline{\delta})$.

Proof. Stochastic process η_t's recurrence is equivalent to the inequality, $\mu^\eta > \frac{(\sigma^\eta)^2}{2}$, when $\eta = 0^+$ (see one dimensional case in Risken (1996)). Asymptotic ansatz when $\eta \to 0$ (in the online appendix of Brunnermeier and Sannikov (2014)): $\mu_t^\eta = \hat{\mu} + o(1), \sigma_t^\eta = \hat{\sigma} + o(1), \psi(\eta) = C_\psi \eta + o(\eta), q(\eta) = \underline{q} + C_q\eta^\alpha + o(\eta^\alpha), \theta(\eta) = C_\theta \eta^{-\beta} + o(\eta^{-\beta})$ $(\alpha, \beta > 0)$. By plugging into equilibrium condition, we have $\hat{\sigma} = \Lambda/\beta\sigma$. From the equations for $q(\eta)$ and $\theta(\eta)$, we find

$$2\frac{(\rho - r)\beta^2\sigma^2}{\Lambda^2} = -\beta\frac{2\hat{\mu}}{\hat{\sigma}^2} + \beta(\beta+1) \to \frac{2\hat{\mu}}{\hat{\sigma}^2} = \beta + 1 - \frac{2(\rho - r)}{\Lambda^2}\sigma^2\beta.$$

□

C Order Reduction in Forward Equation

This section is a technical note. We first discuss the structure of Kolmogorov Forward Equation (KFE), and then give a continuity equation's interpretation of it. Denote the density function as f; the KFE can be heuristically written as:

$$\frac{\partial}{\partial t}f = -\sum_i \frac{\partial}{\partial x_i}\left(\mu_i(x) f\right) + \sum_{i,j}\frac{\partial^2}{\partial x_i \partial x_j}\left(\left(\sigma^2\right)_{ij}(x) f\right) \equiv \widehat{L}^* f = -\nabla \cdot \vec{J},$$

where $\hat{L}^*$ is the Kolmogorov Forward Operator and the *density flux* $\vec{J}$ is defined as

$$J_i = \frac{\partial}{\partial x_i}\left(\mu_i(x) f\right) - \sum_j \frac{\partial^2}{\partial x_i \partial x_j}\left(\left(\sigma^2\right)_{ij}(x) f\right), \quad \text{for the } i\text{–th column of } \vec{J}.$$

To study the stationary distribution we are interested in, we assume the existence of stationary distribution, which means the following assumption holds.

Assumption 1. *Regularity assumptions in Karatzas and Shreve, 1998, Ch. 5.*

1. *$\sigma^2(x)$ is uniform elliptic, i.e.:*
$$\vec{y}^T[\sigma^2(x)]\vec{y} \geq \alpha|\vec{y}|^2, \ \forall \vec{y} \in \mathbb{R}^n.$$

[2]For example, consider a Geometric Brownian Motion: $dX_t = \mu X_t dt + \sigma X_t dZ_t$ with reflecting boundary at $0, D$, the stationary distribution solved from Appendix C is: $f(x) = \frac{\frac{2\mu}{\sigma^2}-1}{D^{\frac{2\mu}{\sigma^2}-1}} x^{\frac{2\mu}{\sigma^2}-2} \times \mathbf{1}_{x\in[0,D]}$, we can see that it cannot be an invariant distribution when $\frac{2\mu}{\sigma^2} - 1 < 0$, as density is negative.

2. Coefficients are smooth and satisfy linear growth conditions:

$$\exists M \in \mathbb{R},\ s.t.\ ||\sigma^2(x)|| \leq M,\ ||a(x)|| \leq M(1+||x||),\ ||b(x)|| \leq M(1+||x||),$$

where $a(x), b(x)$ are defined as

$$a_j(x) = -\mu_j(x) + \sum_{i=1}^{n} \partial_{x_i}[\sigma^2(x)]_{ij},$$
$$b_j(x) = \frac{1}{2}\sum_{i,k} \partial^2_{x_i,x_j}[\sigma^2(x)]_{i,k} - \sum_i \partial_{x_i}\mu_i(x).$$

*3. The stochastic process X_t is **recurrent**.*

For the reflecting boundary[3] in our problem, we have $\vec{J}\cdot\hat{n} \equiv 0$ at $\partial\Omega$. **Solution Determination Problem** can be formally written as:

$$\begin{cases} \nabla\cdot\vec{J} = 0, \\ \vec{J}\cdot\hat{n}|_{\partial\Omega} = 0, \\ J_i = -\mu_i(x)f + \sum_j \frac{1}{2}\frac{\partial}{\partial x_j}((\sigma^2(x))_{ij}f), \end{cases}$$

with normalization condition $\int_\Omega f dV = 1$.

Proposition 3. *Under assumption 1, the solution determination problem is equivalent to the problem:*

$$\vec{J} = 0 \ \textit{and}\ \int_\Omega f dV = 1.$$

Proof. We first show that in the one dimensional case[4], the probability flux is always zero under the reflecting boundary condition. This is because:

$$J(x) \equiv -\mu(x)f(x) + \frac{1}{2}\frac{d}{dx}\left(\sigma^2(x)f(x)\right) = \int_{\underline{x}}^{x} \hat{L}^* f(x')dx' + J(\underline{x}) = 0,$$

where $\underline{x}$ is the lower boundary.

In higher dimensional cases, $\sigma^2(x)$ is the covariance matrix and $\vec{\mu}$ is the drift vector. According to Pavliotis (2014), define $Q(x)$ as: $\vec{Q}(x) = \left(\sigma^2(x)\right)^{-1}(2\vec{\mu}(x) - \nabla\sigma^2(x))$, where $\nabla\sigma^2(x)$ is defined as $\sum_{i,j}\frac{\partial}{\partial x_j}\sigma^2(x)_{ij}\vec{e}_i$. Then[5] $\vec{J}\equiv 0$ **if and only if** $\frac{\partial Q_j}{\partial x_i} = \frac{\partial Q_i}{\partial x_j}$, for all $x \in \Omega$. The density f can be solved as:

$$f(\vec{x}) = A\exp\left(-\frac{1}{2}\int_{x_0\in\partial\Omega}^{\vec{x}} \vec{Q}\cdot d\vec{x}\right) \quad \text{and} \quad \frac{1}{A} = \int_\Omega \exp\left(-\int_{x_0\in\partial\Omega}^{x} \vec{Q}(\vec{x}')\cdot d\vec{x}'\right) d^n\boldsymbol{x}.$$

The above formula implies that once we obtain the solution for the KFE with reflecting boundary, we can conclude the solution is **unique** if assumption 1 holds here. □

[3]Indifference condition for value function at the boundary implies the reflecting boundary for distribution. Usually three boundary conditions are considered. They are: (1) *refecting boundary*: $\vec{J}\cdot\hat{n}|_{\partial\Omega} = 0$; (2) *absorbing boundary*: $f|_{\partial\Omega} = 0$; (3) *periodic boundary*: $\vec{J}|_{x=a} = \vec{J}|_{x=b}$.

[4]Also see in Gabaix et al. (2016).

[5]This condition implies that the path integral exists, or $d\vec{Q}$ is integrable. Intuitively, for a constant $\sigma^2(x)$, this condition means that the drift term is *curl-free*. Accordingly, a counter example in 2D can be constructed as: $\vec{\mu} = \frac{A}{r}\hat{\tau}$, where $\hat{\tau} = \frac{-y}{\sqrt{x^2+y^2}}\vec{e}_x + \frac{x}{\sqrt{x^2+y^2}}\vec{e}_y$.