

---

# PI-LNO: Physics-Informed Latent Neural Operator for Coupled HJB-KFE Equilibria

---

Peter Li<sup>1,2</sup>

## Abstract

We introduce PI-LNO, a physics-informed neural operator for jointly learning structured probabilistic and functional outputs—a stationary density  $\mu(a, z)$ , value function  $V(a, z)$ , and endogenous price  $r^*$ —that must simultaneously satisfy adjoint PDE constraints (HJB and KFE) and a market-clearing closure. Constrained output heads enforce measure validity and value-policy consistency by construction; a three-stage curriculum gates physics residuals progressively to prevent destabilizing gradient feedback during density learning. On 7,500 solver-computed equilibria across a seven-dimensional parameter space, PI-LNO achieves 0.45% value-function MAE and 0.20% Gini MAE (Stage 3, five-seed mean) with  $7,590\times$  batched speedup enabling amortized simulation-based inference; FiLM conditioning contributes  $\sim 2.4\times$  on value error, and the surrogate’s speed permits identification diagnostics that expose weak global identifiability along  $(\sigma, \bar{z})$ .

## 1. Introduction

Learning structured outputs over function spaces poses challenges absent from standard regression: outputs must lie in constrained spaces (probability measures, convex dual pairs), satisfy coupling constraints across output heads, and handle boundary behavior that is singular rather than smooth. These challenges compound when the constraints arise from a coupled PDE system, where multiple output fields must simultaneously satisfy adjoint equations and an endogenous clo-

sure condition.

We study this problem through the stationary equilibria of heterogeneous-agent (HA) models, which require jointly learning a stationary probability density  $\mu(a, z)$ , a value function  $V(a, z)$ , and an endogenous price  $r^*$ . The density must be a valid probability measure;  $V$  and  $\mu$  are governed by adjoint operators (HJB and KFE); and  $r^*$  is determined by integrating  $\mu$  against the capital demand curve. This structure—constrained functional outputs, adjoint PDE coupling, and endogenous closure—is representative of mean-field game and McKean–Vlasov systems and arises naturally in quantitative social science. The computational cost of exact solvers (0.1–0.5s per evaluation) creates a bottleneck for simulation-based inference, which requires  $O(10^3\text{--}10^4)$  forward evaluations.

Our approach. Three structural properties of the joint output space must be preserved:

1. Measure consistency: the predicted  $\hat{\mu}$  must be a non-negative probability measure; un-normalized outputs make distributional moment evaluation meaningless.
2. Adjoint coupling: HJB and KFE are formal adjoint operators; violating this makes the predicted density non-stationary under the predicted policy.
3. Equilibrium closure: the endogenous price is determined by integrating the predicted density; all outputs are globally coupled.

PI-LNO enforces each via constrained architectural heads and adjoint-consistent physics residuals, with a staged curriculum that gates physics losses until supervised fit stabilizes.

Contributions. (1) We identify structural challenges in learning jointly constrained functional outputs (measure-valued density, value function, endogenous scalar) and design PI-LNO to address each. (2) We introduce adjoint-consistent HJB-KFE residuals and a staged physics curriculum respecting the forward-

---

<sup>1</sup>MIT Critical Data, Massachusetts Institute of Technology <sup>2</sup>Substrate Labs. Correspondence to: Peter Li <peterli@substrate-labs.org>.

backward coupling. (3) We evaluate on 7,500 equilibria over a seven-dimensional parameter space, demonstrating gains over purely supervised training and an MLP baseline (architecture ablation; mismatched training and parameter count, see Section 5). (4) We demonstrate surrogate-accelerated simulation-based inference and identification diagnostics; an exact-solver control establishes  $\theta^*$  as a well-separated basin (5 starts, 7D) with recovery bias attributed to weak global identification.

## 2. Problem: HJB-KFE Equilibrium Operator

The Aiyagari (1994) economy in continuous time (Achdou et al., 2022) consists of a unit mass of households facing idiosyncratic income shocks, choosing consumption  $c$  and wealth  $a$  subject to a borrowing constraint  $a \geq a_{\min}$ . Equilibrium requires solving three coupled components.

**HJB equation.**  $\rho V(a, z) = \max_c [u(c) + \partial_a V(a, z) \cdot s(a, z, r) + \sum_{z'} \Lambda_{zz'} V(a, z')]$ , where  $s = wz + ra - c$  is savings,  $u(c) = c^{1-\sigma}/(1-\sigma)$  (log  $c$  in the limit  $\sigma \rightarrow 1$ ) is CRRA utility, and  $\Lambda$  is the income transition generator. The first-order condition gives  $c^*(a, z) = (\partial_a V)^{-1/\sigma}$ .

**Kolmogorov Forward Equation.**  $0 = -\partial_a [s^*(a, z)\mu(a, z)] + \sum_{z'} \Lambda_{z'z}\mu(a, z')$ , where  $\mu(a, z) \geq 0$  integrates to one. The KFE operator is the formal adjoint of the HJB drift operator; this duality must be preserved at the discrete level.

**Market clearing.**  $\sum_z \int a \mu(a, z) da = K^d(r)$ , where  $K^d(r) = L \cdot \left(\frac{A\alpha}{r+\delta}\right)^{1/(1-\alpha)}$  is firm capital demand ( $L$  = labor supply,  $A$  = TFP). The interest rate  $r^*$  clears capital markets endogenously.

The solution operator is  $\mathcal{G} : \theta \mapsto (V^*, \mu^*, r^*)$ , where  $\theta \in \mathbb{R}^7$  collects  $(\rho, \sigma, a_{\min}, \lambda, \bar{z}, \alpha, \delta)$ . PI-LNO approximates  $\mathcal{G}$  with a single forward pass.

## 3. PI-LNO Architecture

PI-LNO has three components: a parameter encoder, a FiLM-conditioned FNO backbone, and a constrained multi-head decoder. Figure 1 illustrates the architecture.

**Parameter encoder.** A 3-layer MLP  $E : \mathbb{R}^7 \rightarrow \mathbb{R}^{128}$  encodes  $\theta$  into a latent conditioning vector  $\mathbf{h}$ . Two linear projections produce broadcast features ( $d_p = d_\ell = 24$ ) replicated across the spatial grid and concatenated with the normalized asset coordinate.

FiLM-conditioned FNO backbone. A  $1 \times 1$  convolution lifts to hidden dimension  $d_h = 96$ . Six residual blocks each combine a spectral path (FNO,  $k_{\max} = 16$  modes) with a local pointwise convolution:

$$\begin{aligned} \mathbf{H}_b &= \mathbf{H}_{b-1} + \text{FiLM}_b^{(2)}(\tilde{\mathbf{H}}_b, \mathbf{h}), \\ \tilde{\mathbf{H}}_b &= \text{LN}\left(\text{MLP}_b\left(\sigma\left(\text{FiLM}_b^{(1)}(\text{LN}(\mathbf{S}_b), \mathbf{h})\right)\right)\right), \end{aligned} \quad (1)$$

where  $\mathbf{S}_b = \mathcal{F}^{-1}[R_b \cdot \mathcal{F}[\mathbf{H}_{b-1}]_{\leq k_{\max}}] + W_b^{\text{loc}} \mathbf{H}_{b-1}$ . Each FiLM layer (Pérez et al., 2018) computes  $\mathbf{x} \odot (1 + \tanh(\gamma)) + \beta$  with  $[\gamma, \beta] = W_{\text{film}} \mathbf{h}$ ; the tanh gate prevents early-training instability. FiLM is critical here because the output distribution  $\mu$  is highly parameter-dependent; the latent code must modulate the spectral path so that  $\partial_a V$  adapts to each  $\theta$ .

**Constrained decoder.** Value head: linear projection, no constraint ( $V$  is an unrestricted real-valued function).

**Density head:** softplus logits followed by normalization  $\hat{\mu}_{ij} \propto \text{softplus}(\ell_{ij}) + \epsilon$  with trapezoid quadrature weights (non-uniform grid spacing carried through). This guarantees  $\hat{\mu} > 0$  and  $\sum_j \int \hat{\mu} da = 1$  for any network weights—a necessary but not sufficient condition for stationarity; the KFE residual provides the sufficiency pressure. Crucially, the normalization uses quadrature-correct integration weights, so the predicted measure is a valid probability measure regardless of the network weights.

**Boundary mass head:** an MLP on left-edge features, global pool, and  $\mathbf{h}$ , designed to capture the Dirac mass of constrained households at  $a = a_{\min}$ . Our solver currently absorbs this mass into the first grid cell rather than reporting it separately, so the head is not actively supervised in the experiments below.

**Rate head:** MLP on global pool  $\oplus \mathbf{h}$ , outputting  $\hat{r} \in \mathbb{R}$ .

**Policy:** consumption  $\hat{c}_{ij} = (\partial_a \hat{V})^{-1/\sigma}$  computed via upwind finite differences on  $\hat{V}$ , not predicted independently. This ensures value-policy structural consistency; the surrogate cannot produce a consumption profile inconsistent with its own value function.

## 4. Structure-Preserving Training

**Adjoint-consistent residuals.** We compute HJB and KFE residuals using the same upwind discretization scheme on the non-uniform grid: both use the local spacing  $\Delta a_i = a_{i+1} - a_i$  in their finite-difference denominators, so the discrete KFE operator is the matrix transpose of the discrete HJB drift operator (the discrete-adjoint convention used by Achdou et al.

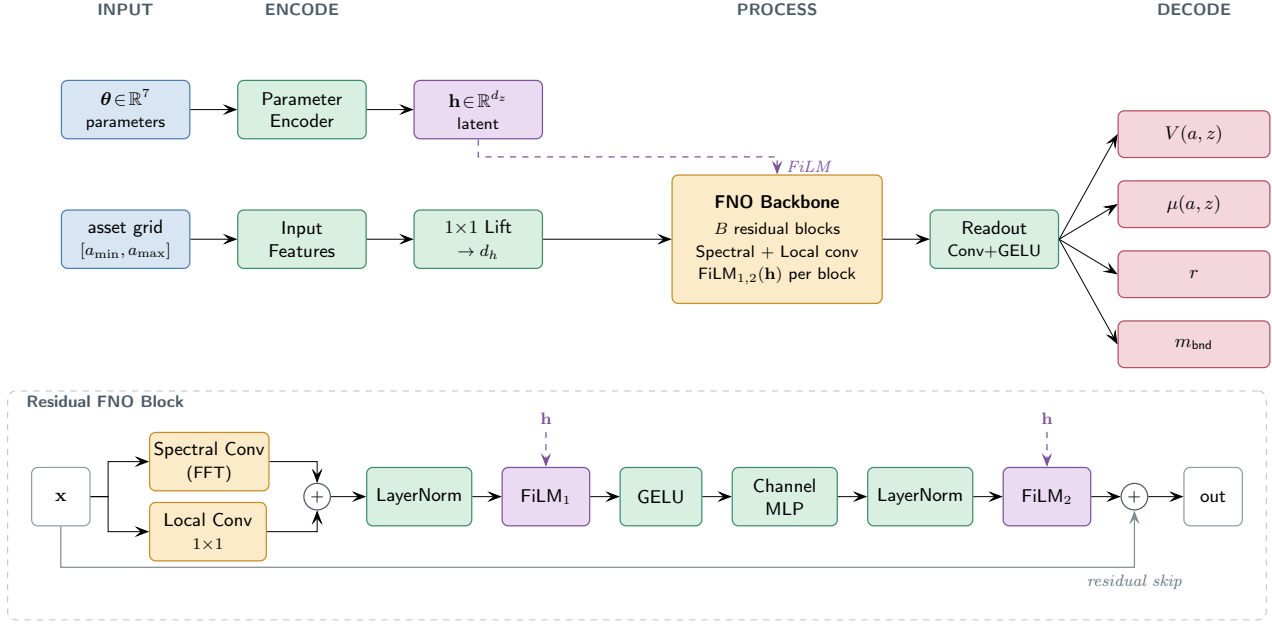


Figure 1. PI-LNO architecture. A parameter encoder maps  $\theta$  to latent  $\mathbf{h}$ ; FiLM-conditioned FNO blocks process the asset grid; multi-head decoder produces value, density, and interest rate.

(2022)). This is the consistency required for residual-based training; we do not impose a separate  $L_w^2$ -adjoint with quadrature weights. Using different schemes (e.g., central differences for KFE, upwind for HJB) would break this adjoint consistency and create irreconcilable conflicts in the loss.

Loss function. The total loss combines supervised fit, physics residuals, and boundary penalties:

$$\mathcal{L} = \mathcal{L}_{\text{data}} + w_{\text{HJB}} \mathcal{L}_{\text{HJB}} + w_{\text{KFE}} \mathcal{L}_{\text{KFE}} + w_{\text{clr}} \mathcal{L}_{\text{clear}} + w_{\text{bnd}} \mathcal{L}_{\text{bnd}} \quad (2)$$

$\mathcal{L}_{\text{data}}$  uses scale-normalized SmoothL1 for all output fields ( $V$ ,  $\mu$ ,  $r$ ,  $c$ ). SmoothL1 is chosen over MSE because its linear tail bounds the gradient contribution of large errors, reducing the pressure on the network to sharpen  $\partial_a \hat{V}$  at extreme calibrations, precisely the sharpening that causes FOC-derived consumption to diverge out of distribution (Section 5). Physics residuals use  $\log(1 + |\cdot|/\text{scale})$  to soften large early-training residuals.  $\mathcal{L}_{\text{HJB}}$  penalizes the HJB equation residual;  $\mathcal{L}_{\text{KFE}}$  penalizes the adjoint-consistent KFE residual;  $\mathcal{L}_{\text{clear}}$  penalizes capital market imbalance.

Three-stage curriculum. Stage 1 (60 epochs): supervised data losses only, all physics weights zero. Establishes a data-driven baseline and warms up the encoder and backbone.

Stage 2 (50 epochs): HJB residual only ( $w_{\text{HJB}} = 0.1$ , reduced lr =  $10^{-4}$ ). Refines value-function gradients

(which determine consumption via the FOC) without perturbing the density.

Stage 3 (80 epochs): full physics ( $w_{\text{HJB}} = 0.5$ ,  $w_{\text{KFE}} = w_{\text{clr}} = 10^{-4}$ ,  $w_{\text{bnd}} = 0.25$ ). Enforces distributional consistency and market clearing on top of the HJB-refined value function.

Physics losses are introduced via a warmup-then-ramp schedule (5-epoch warmup + 15-epoch ramp) at each stage to prevent destabilizing feedback before prior objectives have converged. All stages use AdamW with a 3-epoch LR warmup and gradient clipping ( $\|\cdot\|_2 \leq 0.5$ ).

Why the staged ramp is load-bearing. If KFE and clearing losses are active before the density head has learned meaningful shape, the physics gradient pushes the density toward a residual minimum that is far from the data-supervised target. As the ablation in Section 5 shows, skipping the ramp degrades density MAE by  $1.3\times$  and rate error by  $1.6\times$  at convergence, but the training trajectory is markedly unstable: consumption oscillates by three orders of magnitude across consecutive epochs before eventually settling, indicating that the ramp prevents destabilizing feedback rather than shifting the final basin.

## 5. Experiments

**Setup.** We generate 7,500 training, 1,500 validation, and 1,500 test equilibria using the implicit upwind finite-difference solver of Achdou et al. (2022) on a 64-point non-uniform wealth grid ( $J = 2$  income states). Parameters span a seven-dimensional box (Table 6 in Appendix). The solver accepts 7,500 of 10,820 attempted training draws ( $\sim 69\%$ ); rejected configurations fall into bracket-search failure or HJB-inadmissible regions of the parameter cube, so the surrogate is trained on the converged subset (selection note in Appendix). Training uses five independent seeds (42, 123, 456, 789, 1024); all seeds share the same training dataset (`DATASET_SEED = 0`), so variance across seeds reflects model initialization only. Unless noted, we report mean  $\pm$  s.d. across the five seeds.

### 5.1. In-Distribution Accuracy

Table 1 reports MAE on 1,500 held-out equilibria. Stage 3 (full physics) achieves 0.45% centered-value error, 0.20% wealth Gini MAE,  $9\times$  improvement in  $K/Y$  ratio MAE over Stage 1, and  $2.5\times$  improvement in market-clearing residual. The HJB residual falls from 0.0023 (Stage 1) to 0.00059 (Stage 3). Figure 2 shows representative predictions against solver ground truth.

Consumption is the one exception: Stage 2 achieves best consumption MAE (0.84%), while Stage 3 is 0.86% with substantially higher variance, because full physics imposes additional gradient constraints on the value function. Stage 3 is recommended for distributional/aggregate moments; Stage 2 for applications targeting consumption.

**MLP baseline.** To isolate the contribution of the FNO backbone, we train a variant in which each FNO spectral block is replaced by a pointwise MLP block (same hidden dimension and depth, but no spectral mixing; 594K parameters vs. 1,685K for the FNO version due to the spectral weight tensors). The MLP baseline uses a single-stage training run with the same physics-loss weights and ramp schedule as Stage 3; it is not pre-initialized from Stages 1–2. This is not a fully matched comparison: the FNO benefits from the staged curriculum’s supervised pre-initialization while the MLP does not, so the gap conflates architecture and training procedure to some degree. What the current result establishes is that at matched hidden dimension and depth, the spectral architecture outperforms the pointwise baseline on every distributional metric. Across 5 independent seeds the MLP reaches usable but consistently weaker accuracy:  $1.4\times$  higher value-

function error,  $1.5\times$  higher Gini error,  $1.5\times$  higher rate error, and  $1.6\times$  higher  $\text{Std}(\log c)$  relative to Stage 3 PI-LNO. The parameter-count difference also favors the FNO (1,685K vs. 594K parameters); a wider or curriculum-trained MLP could narrow the gap further.

### 5.2. Out-of-Distribution Generalization

Table 2 evaluates at widened parameter bounds ( $\pm 10\%$  mild,  $\pm 30\%$  extreme). Physics training provides strong OOD regularization for value functions and distributional moments: Stage 3 achieves  $3.2\times$  better value-function error at mild OOD and  $8.5\times$  better Gini than Stage 1.

However, the ordering inverts for FOC-derived consumption. Per-seed counts of finite consumption MAE at mild OOD are  $3/5$  (Stage 1),  $1/5$  (Stage 2),  $1/5$  (Stage 3); all five seeds diverge at extreme OOD across every stage. The failure is not Stage 3-specific: both physics-trained checkpoints produce overflow on  $\geq 4$  of 5 seeds, while supervised-only Stage 1 retains stability on a majority. The mechanism is consistent with the structure-preservation framing: HJB training sharpens  $\partial_a \hat{V}$  to satisfy the gradient-based residual, and the FOC inversion  $\hat{c} = (\partial_a \hat{V})^{-1/\sigma}$  amplifies any sharpening that extrapolates outside the training box. Stage 1 leaves  $\partial_a \hat{V}$  smoother because no gradient-based loss pressures it, so the FOC inversion stays bounded more often. Because the exponent  $-1/\sigma$  ranges from  $-1$  to  $-0.25$  over the parameter box ( $\sigma \in [1, 4]$ ), the FOC inversion is most sensitive at low  $\sigma$ ; per- $\sigma$  stratification of the OOD failures (which the current aggregate-only eval does not provide) is left to future work. The choice of smooth $L_1$  over MSE for  $\mathcal{L}_{\text{data}}$  partially mitigates this effect: MSE’s quadratic penalty would impose even stronger gradient pressure on outlier calibrations, further sharpening  $\partial_a \hat{V}$  and worsening OOD consumption stability. For OOD consumption queries, solver verification is required regardless of stage; Stage 1 is the most stable surrogate option.

### 5.3. Ablations

Table 3 isolates three design choices. MLP baseline replaces the FNO spectral blocks with pointwise MLP blocks (same hidden dimension and depth, single-stage training with Stage 3 physics weights). The MLP is competitive on consumption and clearing but  $1.8\times$  worse on density and  $2.2\times$  worse on HJB residual, consistent with the FNO’s global receptive field improving gradient-based physics losses. Removing FiLM conditioning (3-stage curriculum with FiLM  $\equiv$  identity, 5-seed) degrades value error by  $2.4\times$ , rate by  $1.8\times$ , and Gini by  $2.1\times$ ; its role is to let the latent code mod-

Table 1. In-distribution accuracy (mean  $\pm$  s.d., 5 seeds, 1,500 test equilibria). Bold: best per row.

Metric	Stage 1	Stage 2	Stage 3	MLP <sup>†</sup>
Value (norm.)	0.018 $\pm$ 0.003	0.011 $\pm$ 0.002	<b>0.0045 <math>\pm</math> 0.0010</b>	0.0063 $\pm$ 0.0003
Value, centered	0.011 $\pm$ 0.002	0.0088 $\pm$ 0.0013	<b>0.0045 <math>\pm</math> 0.0010</b>	0.0078 $\pm$ 0.0003
Density MAE	0.00152 $\pm$ 0.00006	0.00116 $\pm$ 0.00003	<b>0.00106 <math>\pm</math> 0.00005</b>	0.00192 $\pm$ 0.00011
Interest rate (norm.)	0.0141 $\pm$ 0.0007	0.0106 $\pm$ 0.0008	<b>0.0089 <math>\pm</math> 0.0012</b>	0.0135 $\pm$ 0.0009
Consumption (norm.)	0.0104 $\pm$ 0.0011	<b>0.0084 <math>\pm</math> 0.0011</b>	0.0086 $\pm$ 0.0017	0.0092 $\pm$ 0.0003
Wealth Gini	0.031 $\pm$ 0.010	0.027 $\pm$ 0.005	<b>0.0020 <math>\pm</math> 0.0002</b>	0.0029 $\pm$ 0.0004
K/Y ratio MAE	0.112 $\pm$ 0.045	0.096 $\pm$ 0.020	<b>0.012 <math>\pm</math> 0.001</b>	0.0164 $\pm$ 0.0013
Std(log c) MAE	0.016 $\pm$ 0.003	0.013 $\pm$ 0.002	<b>0.0057 <math>\pm</math> 0.0003</b>	0.0090 $\pm$ 0.0007
Clearing residual ( $ K^s - K^d /K^d$ )	0.207 $\pm$ 0.055	0.194 $\pm$ 0.027	<b>0.082 <math>\pm</math> 0.003</b>	0.083 $\pm$ 0.005
HJB residual (norm.)	0.0023 $\pm$ 0.0001	0.0017 $\pm$ 0.0001	<b>0.00059 <math>\pm</math> 0.00010</b>	0.00130 $\pm$ 0.00003

<sup>†</sup>MLP baseline: FNO blocks replaced by pointwise MLP blocks (5 seeds, single-stage training with Stage 3 physics weights and warmup/ramp schedule; 594K params vs. 1,685K for FNO).




Figure 2. Representative predictions (Stage 3, seed 42) vs. solver ground truth for four held-out equilibria spanning the parameter space. Left: value functions  $V(a, z)$ . Right: stationary densities  $\mu(a, z)$ . The surrogate captures both the smooth interior and the boundary-constraint region.

ulate the spectral path so  $\partial_a V$  is parameter-adapted. Removing the staged ramp produces training instability (consumption oscillates by  $10^3 \times$  across consecutive epochs) that eventually recovers but at the cost of distributional accuracy.

#### 5.4. Computational Speedup

Single-query PI-LNO takes 9.1 ms on an NVIDIA L4 GPU vs. 234 ms for the FD solver on a CPU:  $26 \times$  speedup. Batched inference exhibits superlinear scaling; a batch of 1,000 configurations completes in 30.8 ms (0.031 ms/sample), a  $7,590 \times$  speedup. The FD solver cannot parallelize across parameter configurations (each requires independent nested iteration); the neural operator exploits GPU parallelism fully. The comparison reflects the intended deployment scenario (a GPU surrogate replacing a CPU solver in an estimation loop), and a GPU-accelerated FD solver would narrow the gap.

## 6. Application: Surrogate-Accelerated Inference

The  $7,590 \times$  batched speedup makes simulation-based inference (McFadden, 1989) practical:  $10^4$  surrogate evaluations cost seconds rather than hours, enabling amortized inference over the full 7D parameter space.

Synthetic recovery. The standard surrogate validation generates moment targets from the exact solver

at a known  $\theta^*$ , recovers  $\hat{\theta}$  via surrogate-based SMM, and reports the bias  $\hat{\theta} - \theta^*$ . This tests surrogate reliability conditional on model correctness, distinct from fitting empirical targets, which tests model specification.

We evaluate three interior calibrations spanning the parameter space: point A is a central calibration, point B has low risk aversion and loose borrowing, point C has high risk aversion ( $\sigma = 3.0$ ) and a tighter borrowing constraint. All seven structural parameters are recovered jointly via differential evolution over the training box.

The key finding is in Tables 4–5: the surrogate matches its own target moments to near-machine precision (SMM objectives  $2.8 \times 10^{-7}$ ,  $2.3 \times 10^{-8}$ ,  $5.0 \times 10^{-6}$ ), yet the estimated parameters are far from  $\theta^*$  on  $\sigma$  and  $\bar{z}$ . A control test using the FD solver as the moment function (Section 6, “Control: FD-SMM”) shows  $\theta^*$  is a well-separated basin (5 starts, 7D space), but identification is weak along the  $(\sigma, \bar{z})$  direction (elasticity-Jacobian condition number  $\approx 30$  at Point A; Table 5). A small ( $\sim 3\%$ ) surrogate moment-surface error then amplifies into a large  $(\sigma, \bar{z})$  parameter bias along that ridge.

Identification diagnostics. To diagnose this, we compute the elasticity Jacobian  $E_{ij} = (\partial m_i / \partial \theta_j)(\theta_j / m_i)$  at  $\theta^*$  via central differences ( $2 \times 7 = 14$  surrogate evaluations,  $< 1$  second total) and take its SVD (Newey & McFadden, 1994). Table 5 reports the results.

Table 2. OOD generalization. Mean  $\pm$  s.d., 5 seeds, 200 equilibria per tier. Bold: best per column. Consumption shown separately: Stage 3 is unstable OOD due to FOC gradient amplification.

Stage	Value (norm.)		Wealth Gini	
	Mild ( $\pm 10\%$ )	Extreme ( $\pm 30\%$ )	Mild ( $\pm 10\%$ )	Extreme ( $\pm 30\%$ )
1 (Supervised)	0.028 $\pm$ 0.003	0.275 $\pm$ 0.017	0.027 $\pm$ 0.008	0.034 $\pm$ 0.008
2 (HJB)	0.019 $\pm$ 0.003	0.250 $\pm$ 0.017	0.024 $\pm$ 0.004	0.030 $\pm$ 0.006
3 (Full physics)	<b>0.0088 <math>\pm</math> 0.0016</b>	<b>0.229 <math>\pm</math> 0.013</b>	<b>0.0032 <math>\pm</math> 0.0006</b>	<b>0.016 <math>\pm</math> 0.002</b>

Consumption (norm.) finite-seed counts at mild OOD: S1 = 3/5, S2 = 1/5, S3 = 1/5; all stages diverge at extreme OOD. FOC instability is driven by HJB training, not by Stage 3 alone.

Table 3. Ablations on the matched 3-stage pipeline (smooth- $L_1$ , LR warmup, gradient clipping throughout). All rows: 5-seed mean  $\pm$  s.d. on 1,500 test equilibria.  $^\dagger$ same 3-stage curriculum as Full, with FiLM  $\rightarrow$  identity.  $^\ddagger$ single-stage with Stage-3 physics weights.  $^\S$ MLP baseline: single-stage with Stage-3 physics weights and warmup/ramp; FNO blocks replaced by pointwise MLPs (594K params vs. 1,685K for FNO).

Configuration	Value (norm.)	Consumption (norm.)	Density MAE	Interest rate (norm.)
Full (Stage 3)	<b>0.0045 <math>\pm</math> 0.0010</b>	<b>0.0086 <math>\pm</math> 0.0017</b>	<b>0.00106 <math>\pm</math> 0.00005</b>	<b>0.0089 <math>\pm</math> 0.0012</b>
MLP baseline $^\S$	0.0063 $\pm$ 0.0003	0.0092 $\pm$ 0.0003	0.00192 $\pm$ 0.00011	0.0135 $\pm$ 0.0009
w/o FiLM $^\dagger$	0.0109 $\pm$ 0.0024	0.0111 $\pm$ 0.0015	0.00165 $\pm$ 0.00019	0.0164 $\pm$ 0.0037
w/o staged ramp $^\ddagger$	0.0057 $\pm$ 0.0014	0.0109 $\pm$ 0.0011	0.00136 $\pm$ 0.00004	0.0141 $\pm$ 0.0017

Condition numbers of 27–84 confirm the 5-moment system weakly identifies all 7 parameters across the parameter space. The smallest singular value ( $\approx 0.09$ – $0.14$ ) is consistent across all three points, indicating a structurally under-identified direction regardless of calibration.  $\sigma$  does not appear as the weakest local direction in any of the SVDs, yet has the largest recovery biases. The FD-SMM control resolves the apparent contradiction: the moment surface is locally curved enough that local sensitivity analysis sees  $\sigma$  as identifiable, but globally weakly identified along the  $(\sigma, \bar{z})$  direction, and surrogate moment-surface errors exploit precisely that weak global direction. This type of global non-identification, invisible to local sensitivity analysis, is diagnosable in practice only because the surrogate makes dense exploration of the moment surface cheap.

Multi-seed recovery at Point A. Re-running the recovery against each of the 5 trained checkpoints (independent training seeds, same Point A target) gives SMM objective  $< 1.7 \times 10^{-6}$  in all cases. Bias means  $\pm$  s.d. across the 5 seeds:  $\sigma = +0.74 \pm 0.59$ ,  $\bar{z} = +0.16 \pm 0.06$ ,  $\rho = +0.008 \pm 0.008$ ,  $\lambda = -0.022 \pm 0.016$ ,  $a_{\min} = +0.012 \pm 0.002$ ,  $\alpha = +0.003 \pm 0.018$ ,  $\delta = +0.001 \pm 0.005$ .

Control: FD-SMM at the same target. To distinguish surrogate-shape error from true non-identification, we re-ran the same SMM problem using the exact FD solver as the moment function (no surrogate), with multi-start L-BFGS-B from  $\theta^*$  and four perturbations of it (4,000 solver evaluations,  $\sim 4$  minutes). The FD-

SMM optimizer started at  $\theta^*$  converges to  $\theta^*$  exactly with loss 0; perturbed starts converge to local minima within  $|\Delta\sigma| \leq 0.4$  of  $\sigma^*$  with loss  $4 \times 10^{-4}$ – $5 \times 10^{-2}$  (all substantially worse than the global optimum at loss 0, confirming they are local traps rather than near-equivalent optima). Independently, the surrogate’s recovered  $\hat{\theta}$  from each seed produces FD-evaluated moments off by 2–5% relative. Interpretation: with five starts in seven dimensions these results cannot rule out other global optima, but every perturbed start lands in a substantially worse local trap, making  $\theta^*$  a well-separated basin; identification is weak along the  $(\sigma, \bar{z})$  direction (condition number  $\approx 30$ , Table 5), so a small ( $\sim 3\%$ ) surrogate moment-surface error amplifies into a large  $(\sigma, \bar{z})$  parameter bias. The seed-to-seed s.d. in  $\sigma$  reflects different surrogates’ shape errors interacting with the same weakly-identified ridge. Non-collinear parameters ( $\rho, a_{\min}, \lambda, \alpha, \delta$ ) are recovered consistently across seeds with bias s.d. under 0.02 because they are well-identified and unaffected by the same shape error.

Limitations. The Stage 3 surrogate is reliable for in-distribution value and distributional-moment queries. OOD consumption queries require solver verification (Table 2). Approximation bias propagates into parameter estimates; FD refinement at the surrogate’s SMM optimum is recommended before reporting empirical estimates. The 8% market-clearing residual at Stage 3 is small enough not to dominate the moment errors used in SMM but is the largest residual reported here; tightening it likely requires increasing the clearing-loss weight or sharper rate prediction.

Table 4. Synthetic recovery bias  $\hat{\theta} - \theta^*$  (seed 42 Stage 3, differential evolution, 300 iterations). SMM objectives at convergence: A =  $2.8 \times 10^{-7}$ , B =  $2.3 \times 10^{-8}$ , C =  $5.0 \times 10^{-6}$ ; max relative moment error 0.04% (A), 0.01% (B), 0.19% (C). Large  $\sigma/\bar{z}$  biases reflect weak global identification along the  $(\sigma, \bar{z})$  direction (Table 5) interacting with  $\sim 3\%$  surrogate moment-surface error; the FD-SMM control (§6) shows  $\theta^*$  is a well-separated basin (5 starts, 7D).

Parameter	Point A (central)		Point B (low $\sigma$ )		Point C (high $\sigma$ )	
	$\theta^*$	Bias	$\theta^*$	Bias	$\theta^*$	Bias
$\rho$ (discount)	0.050	+0.006	0.040	+0.013	0.060	+0.017
$\sigma$ (risk aversion)	2.0	<b>+0.966</b>	1.5	<b>+1.568</b>	3.0	<b>-0.844</b>
$a_{\min}$ (borrowing)	0.10	+0.016	0.05	-0.009	0.20	+0.007
$\lambda$ (separation)	0.15	-0.009	0.10	-0.002	0.20	+0.056
$\bar{z}$ (productivity)	1.00	<b>+0.175</b>	1.20	<b>+0.417</b>	1.00	+0.024
$\alpha$ (capital share)	0.33	-0.002	0.30	+0.060	0.35	+0.017
$\delta$ (depreciation)	0.060	+0.005	0.050	+0.015	0.070	-0.023

Table 5. SVD identification diagnostics at each true  $\theta^*$ .  $\sigma_1$  = largest singular value;  $\sigma_5$  = smallest;  $\kappa = \sigma_1/\sigma_5$ . Weakest-direction parameter = largest absolute loading in the last right singular vector.

Point	$\sigma_1$	$\sigma_5$	$\kappa$	Weakest direction
A (central)	4.25	0.141	30	$\rho$ (loading +0.87)
B (low $\sigma$ )	2.62	0.096	27	$\delta$ (loading -0.66)
C (high $\sigma$ )	7.74	0.093	84	$\lambda$ (loading -0.80)

Scale. The current setup uses  $J = 2$  income states and a 64-point non-uniform wealth grid. At this resolution our  $k_{\max} = 16$  Fourier modes cover 25% of the spectrum, which is high relative to typical FNO benchmarks; on a 256-point grid with the same  $k_{\max}$  this drops to 6%, where the spectral structure becomes a stronger architectural prior. Richer HA models (HANK, life-cycle, multi-asset) involve finer grids and more income states; changing  $J$  requires regenerating training data and rederiving the discrete operators.

## 7. Related Work

Structured inference and probabilistic surrogates. Learning structured outputs that must lie in constrained spaces (probability measures, dual function pairs) is a recurring challenge in probabilistic ML. PINNs (Raissi et al., 2019) embed PDE residuals in the loss but solve one configuration per training run; PI-LNO is an operator that amortizes solution across the full parameter family, enabling downstream simulation-based inference over thousands of evaluations. Curriculum strategies that progressively introduce constraints have been explored for single-PDE settings (Karniadakis et al., 2021), and concurrent work (Marcandelli et al., 2026) independently proposes multi-stage curricula for FNO-class architectures. Our three-stage design is specialized to the adjoint-coupled HJB-KFE system, where the activation order (HJB be-

fore KFE) is dictated by the forward-backward structure.

Neural operators. FNO (Li et al., 2021) and DeepONet (Lu et al., 2021) learn function-space mappings for parametric PDEs. Applications to fluid dynamics, weather forecasting, and materials science typically involve single PDEs with well-behaved regularity. Our setting introduces forward-backward coupling, endogenous prices, and constraint singularities absent from standard PDE surrogate benchmarks; the FNO’s global receptive field is particularly well-suited to the non-local propagation of boundary singularity errors.

Surrogates for coupled PDE systems in science. The most directly comparable work is Grzeskiewicz (2025), which applies PINN-style training to the same Aiyagari-class system but solves the time-dependent transition for one fixed parameterization; PI-LNO solves the stationary parameter-to-equilibrium operator for any  $\theta$  in a 7-dimensional training box, amortizing cost across the parameter family. Han et al. (2026) and Fernández-Villaverde et al. (2020) apply neural networks to heterogeneous-agent and McKean–Vlasov problems; Kase et al. (2025) target likelihood-based estimation of nonlinear HA models. Yang et al. (2025) studies HA equilibrium computation with aggregate shocks; Zhong (2024) applies operator learning to stationary HA equilibria across parameter space.

## 8. Conclusion

PI-LNO demonstrates that structure preservation is essential for reliable surrogates that jointly learn constrained probabilistic and functional outputs. Architecturally enforced measure normalization ensures the density head always produces a valid probability measure; adjoint-consistent physics losses enforce forward-backward consistency; FiLM conditioning ( $\sim 2.4\times$  im-

provement on value error, 5-seed matched-protocol ablation) lets the latent code modulate spectral representations so outputs adapt correctly across the parameter family; and the staged curriculum stabilizes training of the coupled output heads. The resulting surrogate achieves sub-1% error on equilibrium distributions and value functions, with batched speedup that enables simulation-based inference workflows previously constrained by computation. The surrogate’s speed permits dense identification diagnostics that are computationally prohibitive with an exact solver; an exact-solver control test confirms  $\theta^*$  is a well-separated basin, with recovery bias attributable to weak global identification amplifying small surrogate moment-surface error. The FNO backbone and FiLM conditioning are architecture-agnostic and transfer to other coupled systems where outputs must jointly satisfy adjoint PDE constraints; extension to  $J \geq 3$  states requires rederiving the adjoint-consistent discrete operators, a natural next step.

## References

- Achdou, Y., Han, J., Lasry, J.-M., Lions, P.-L., and Moll, B. Income and wealth distribution in macroeconomics: A continuous-time approach. *The Review of Economic Studies*, 89(1):45–86, 2022. doi: 10.1093/restud/rdab002.
- Fernández-Villaverde, J., Nuño, G., Sorg-Langhans, G., and Vogler, M. Solving high-dimensional dynamic programming problems using deep learning. Manuscript, 2020.
- Grzeskiewicz, M. Solving heterogeneous agent models with physics-informed neural networks, 2025.
- Han, J., Yang, Y., and E, W. DeepHAM: A global solution method for heterogeneous agent models with aggregate shocks. *Quantitative Economics*, 17:297–341, 2026. URL <https://arxiv.org/abs/2112.14377>.
- Karniadakis, G. E., Kevrekidis, I. G., Lu, L., Perdikaris, P., Wang, S., and Yang, L. Physics-informed machine learning. *Nature Reviews Physics*, 3(6):422–440, 2021. doi: 10.1038/s42254-021-00314-5.
- Kase, H., Melosi, L., and Rottner, M. Estimating nonlinear heterogeneous agent models with neural networks. BIS Working Papers 1241, Bank for International Settlements, January 2025. URL <https://www.bis.org/publ/work1241.htm>.
- Li, Z., Kovachki, N., Azizzadenesheli, K., Liu, B., Bhattacharya, K., Stuart, A., and Anandkumar, A. Fourier neural operator for parametric partial differential equations. In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=c8P9NQVtmnO>.
- Lu, L., Jin, P., Pang, G., Zhang, Z., and Karniadakis, G. E. Learning nonlinear operators via DeepONet based on the universal approximation theorem of operators. *Nature Machine Intelligence*, 3(3):218–229, 2021. doi: 10.1038/s42256-021-00302-5.
- Marcandelli, P., Mathur, N., Markidis, S., Siena, M., and Mariani, S. Unsupervised physics-informed operator learning through multi-stage curriculum training, 2026.
- McFadden, D. A method of simulated moments for estimation of discrete response models without numerical integration. *Econometrica*, 57(5):995–1026, 1989. doi: 10.2307/1913621.
- Newey, W. K. and McFadden, D. Large sample estimation and hypothesis testing. In Engle, R. F. and McFadden, D. L. (eds.), *Handbook of Econometrics*, volume 4, pp. 2111–2245. Elsevier, 1994. doi: 10.1016/S1573-4412(05)80005-4.
- Pérez, E., Strub, F., de Vries, H., Dumoulin, V., and Courville, A. FiLM: Visual reasoning with a general conditioning layer. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018. doi: 10.1609/aaai.v32i1.11671. URL <https://ojs.aaai.org/index.php/AAAI/article/view/11671>.
- Raissi, M., Perdikaris, P., and Karniadakis, G. E. 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, 2019. doi: 10.1016/j.jcp.2018.10.045.
- Yang, Y., Wang, C., Schaab, A., and Moll, B. Structural reinforcement learning for heterogeneous agent macroeconomics, 2025. URL <https://ssrn.com/abstract=5988734>. SSRN working paper.
- Zhong, Y. Operator learning in macroeconomics, 2024. URL [https://warwick.ac.uk/fac/soc/economics/staff/yzhong/operator\\_learning\\_crepe\\_day.pdf](https://warwick.ac.uk/fac/soc/economics/staff/yzhong/operator_learning_crepe_day.pdf). Presentation slides, The CREPE Conference Day.

## A. Parameter Space and Training Details

Table 6. Parameter space. All ranges are uniform distributions.

Parameter	Description	Min	Max
$\rho$	Discount rate	0.03	0.08
$\sigma$	Risk aversion	1.0	4.0
$a_{\min}$	Borrowing limit	0.0	0.5
$\lambda$	Job separation	0.05	0.30
$\bar{z}$	Mean productivity	1.0	1.8
$\alpha$	Capital share	0.28	0.38
$\delta$	Depreciation	0.04	0.08

Full architecture hyperparameters. Hidden channels  $d_h = 96$ ; latent dim  $d_z = 128$ ; parameter/latent broadcast  $d_p = d_\ell = 24$ ; FNO blocks  $B = 6$ ; Fourier modes  $k_{\max} = 16$ ; MLP expansion  $2.0\times$ ; encoder hidden dim 256; rate/boundary MLP dim 128; density floor  $\epsilon = 10^{-8}$ . The income process uses a symmetric two-state Poisson process with spread  $\eta = 0.30$  around the mean (so  $z_1 = \bar{z}(1 - \eta)$ ,  $z_2 = \bar{z}(1 + \eta)$ ) and TFP  $A = 0.1$  in the Cobb-Douglas production function.

Training hyperparameters. Stage 1: 60 epochs, lr =  $10^{-3}$ , batch 32, cosine LR decay with 3-epoch warmup. Stage 2: 50 epochs, lr =  $10^{-4}$ , same schedule. Stage 3: 80 epochs, lr =  $10^{-4}$ . All stages: AdamW weight decay  $10^{-5}$ , gradient clipping  $\|\cdot\|_2 \leq 0.5$ . Physics warmup 5 epochs + ramp 15 epochs per stage where active. Total wall-clock training time (all three stages):  $\sim 14$  hours per seed on a TPU v6e-8; dataset generation (7,500 equilibria) takes  $\sim 40$  minutes on a single CPU core.

Dataset selection bias. The solver accepts 7,500 of 10,820 attempted training draws (69.3%); rejected configurations split between bracket-search failure (1,421/seed 42) and HJB-inadmissible regions (1,893/seed 42). The surrogate is therefore trained on a converged subset of the parameter cube, not a uniform draw. This is consistent across the five seeds and the val/test splits (acceptance 69.3%/69.3%/69.1%). Implications: (i) OOD evaluations sample from the same converged manifold, not the full  $\pm 10\%/ \pm 30\%$  box; (ii) parameter regions where the FD solver itself struggles are absent from training, which is one mechanism behind FOC consumption instability when the surrogate is queried outside the training manifold.

Code availability. Code and trained checkpoints will be released upon acceptance.