Deep Learning for Over-the-Counter Market Models*

Jonathan Payne† Adam Rebei‡ Yucheng Yang§

June 16, 2026

This is a companion note to Payne, Rebei, and Yang (2024), "Deep Learning for Search and Matching Models."[1] We use the DeepSAM method developed in that paper to solve an over-the-counter (OTC) bond market model with search and matching frictions, and we use the solution to study financial crises in the bond market. The model can be thought of as an extension to Duffie, Gârleanu, and Pedersen (2005) and Weill (2008) that expands the investor and asset heterogeneity and incorporates aggregate shocks.[2] We model bond duration explicitly and discuss how investor composition influences the emergent yield curve and bond market responses to financial crises. From a technical point of view, relative to the labor market models in Payne et al. (2024), this note introduces idiosyncratic type switching and asset trade.

1 Environment

Setting: Time is continuous with an infinite horizon. The economy has a collection of assets, indexed by $k \in \{1, \ldots, K\}$, which we interpret as bonds. Each asset k has positive net supply s_k, pays a flow dividend $\delta > 0$ each period and one good at

*We thank Nikolai Kurlovich for outstanding research assistance. We thank the Swiss NSF (#10003091) for generous financial support.

†Princeton University, Department of Economics. Email: jepayne@princeton.edu

‡Stanford University, Institute for Computational & Mathematical Engineering. Email: rebei@stanford.edu

§University of Zurich, Department of Finance and SFI. Email: yucheng.yang@uzh.ch

[1] This note was Section 5 in the working paper version of Payne, Rebei, and Yang (2024). It is now circulated as a separate companion note.

[2] To ensure tractability, Duffie, Gârleanu, and Pedersen (2005) restricts the model to two types of investors and one type of asset. Weill (2008) expands the set of investors and assets but simplifies trading between agents who have different assets and does not include aggregate shocks.

maturity, and matures at the rate $1/\tau_k$ (implying an average maturity of τ_k).

Investors and preferences: The economy is populated by a unit-mass continuum of infinitely-lived and risk-neutral investors indexed by $j \in \{1, \ldots, J\}$. Investors are heterogeneous in their discount rates and asset holding costs, which can be different types of asset market participants such as pension funds, hedge funds, or dealers. Investors of type j discount the future at rate $\rho_j > 0$. An investor gets a marginal utility of 1 from a non-storable numeraire good. In order to make payments, investors are endowed with a technology that instantly produces numeraire goods, at a unit marginal cost. An investor can hold either zero or one share of at most one type of asset. When investor j holds asset k, they get flow utility $\delta - \psi(j,k)$, where $\psi(j,k)$ is a "holding cost" that reflects institutional constraints. Investors are subject to independent idiosyncratic shocks of switching between types that follow a continuous time Markov chain. Let $\lambda_{i,j}$ denote the rate of switching from type i to j and let Λ denote the matrix of switching rates.

Financial crisis risk: The aggregate state in the economy is $z \in \{z_1, \ldots, z_n\}$, which follows a continuous time Markov process where $\zeta_{z,z'}$ denotes the rate at which the process switches from z to z'. We allow the aggregate state to affect agent switching rates and haircuts. Formally, at state z, the switching rate from agent type i to agent type j is given by $\lambda_{ij}(z)$. Changes to $\lambda_{ij}(z)$ impact the fraction of agents with low or high holding costs so, as in Duffie, Gârleanu, and Pedersen (2005), we interpret these shocks as changes to the "liquidity" or "institutional" constraints in the investor population. In addition, at state z, asset k pays a fraction $\phi(k,z)$ of the coupon and the principal. We interpret $1 - \phi(k,z) > 1$ as a "haircut" on the bond.

Primary market: When bonds mature, they are replaced by new bonds in the economy. We assume there is an exogenous primary market that allocates new bonds with maturity τ_k to type j investors who are not holding assets at the rate $\xi_{j,k}$.

Distribution: An investor's idiosyncratic state is made up of her type $j \in \{1, \ldots, J\}$ and her ownership status for each asset $k \in \{1, \ldots, K\}$. Hence the set of investor idiosyncratic states is:

$$A = \{1n, 2n, \ldots, Jn, 1o1, \ldots, 1oK, 2o1, \ldots 2oK, Jo1, \ldots, JoK\}$$

in which jn denotes an investor with type j who does not hold any asset, and jok denotes an type-j investor holding an asset k. For each $a \in A$, let $g(a)$ denote the fraction of total investors that have state a and $g = (g(a))_{a \in A}$ denote the mass function. Relative to the labor model, the support of the distribution has been expanded to account for type switching.

Meeting and Bargaining: The contact rate between investors with idiosyncratic states a and b is: $\mathcal{M}_{a,b} = \kappa_{a,b} g_a g_b$. When agents with states a and b meet, they engage in generalized Nash bargaining with bargaining power $\beta_{a,b}$ for the agent in state a.

2 Equilibrium

The aggregate states are (z, g). We denote the law of motion for the cross-sectional idiosyncratic state distribution, g, by the form $dg_t(a) = \mu^g(a, z, g)dt$ and agents' belief about the law of motion by $\breve{\mu}^g(a, z, g)$.

Trade and surplus division: Upon meeting, agents negotiate asset trades according to a generalized Nash Bargaining protocol. Let $V^n(in, z, g)$ denote the value function for an investor of type i without an asset, and let $V^o(iok, z, g)$ denote the value function for an investor of type i with an asset. If a type-i agent holding asset k meets a type-j agent without an asset, the total surplus from asset trade is:

$$S(iok, jn, z, g) = V^n(in, z, g) - V^o(iok, z, g) + V^o(jok, z, g) - V^n(jn, z, g)$$

The generalized Nash bargaining protocol implies that, if agents trade, the price paid for the asset, $p_k(in, jok, z, g)$, solves:

$$\begin{aligned} \Delta V_{i[ok \to n]} + p_k(in, jok, z, g) &= \beta_{iok,jn} S(iok, jn, z, g) \\ \Delta V_{j[n \to ok]} - p_k(in, jok, z, g) &= (1 - \beta_{iok,jn}) S(iok, jn, z, g) \end{aligned}$$

where $\Delta V_{i[ok \to n]}(z, g) := V^n(in, z, g) - V^o(iok, z, g)$ and $\Delta V_{j[n \to ok]}(z, g) := V^o(jok, z, g) -$

$V^n(jn, z, g)$. This implies that agents choose to trade if the surplus is positive.

Similarly, if a type-i agent holding asset k meets a type-j agent holding asset l, then the total surplus if the investors exchange assets is:

$$S(iok, jol, z, g) = V^o(iol, z, g) - V^o(iok, z, g) + V^o(jok, z, g) - V^o(jol, z, g)$$

Once again, the agents trade if the surplus is positive and the generalized Nash bargaining protocol implies that the net transfer between the agents $\Delta p_{k,l}$ satisfies:

$$\Delta V_{i[ok\to ol]} + \Delta p_{k,l}(iol, jok, z, g) = \beta_{iok,jol} S(iok, jol, z, g)$$
$$\Delta V_{j[ol\to ok]} - \Delta p_{k,l}(iol, jok, z, g) = (1 - \beta_{iok,jol}) S(iok, jol, z, g).$$

Hamilton-Jacobi-Bellman Equations: Given their belief $\tilde{\mu}$, the value function for a non-owner with type i satisfies the following HJB equation:

$$\begin{aligned} \rho_i V^n(in, g, z) = &\sum_a \kappa_{in,a} \alpha(in, a, g, z) \beta_{in,a} S(in, a, z, g) \\ &+ \sum_k \xi_{i,k} (V^o(iok, g, z) - V^n(in, g, z)) + \sum_{j\neq i} \lambda_{i,j}(z)(V^n(jn, g, z) - V^n(in, g, z)) \\ &+ \sum_{z'} \zeta_{z,z'} (V^n(in, g, z') - V^n(in, g, z)) + \sum_{a\in A} \partial_{g_a} V^n(in, g, z) \breve{\mu}^g(a, z) \end{aligned} \tag{2.1}$$

where $\alpha(in, jok, g, z)$ is an indicator function for whether the trade is accepted upon matching, which in equilibrium occurs if the surplus from the trade is positive $S(in, jok, g, z) > 0$. Likewise, the value function for an investor of type i holding asset k, $V^o(iok, g, z)$, is given by the following HJB equation:

$$\begin{aligned} \rho_i V^o(iok, g, z) = &\delta\phi(k, z) - \psi(i, k) + \frac{1}{\tau_k}(V^n(in, g, z) + \phi(k, z) - V^o(iok, g, z)) \\ &+ \sum_a \kappa_{iok,a} \alpha(iok, a, g, z) g_a \beta_{iok,a} S(iok, a, g, z) + \sum_{a\in A} \partial_{g_a} V^o(iok, g, z) \breve{\mu}^g(a, z) \\ &+ \sum_{j\neq i} \lambda_{i,j}(z)(V^o(jok, g, z) - V^o(iok, g, z)) + \sum_{z'} \zeta_{z,z'} (V^o(iok, g, z') - V^o(iok, g, z)). \end{aligned} \tag{2.2}$$

Kolmogorov Forward Equation: The law of motion for the distribution of non-owner states and owner states are given respectively by:

$$\frac{dg_{in}}{dt} = \mu^g(in, z, g) = \sum_{j \neq i} \lambda_{j,i}(z) g_{jn} + \sum_{j \neq i} \sum_k \kappa_{jn,iok} g_{jn} g_{iok} \alpha(jn, iok, g, z) \tag{2.3}$$
$$- \sum_{j \neq i} \lambda_{i,j}(z) g_{in} - \sum_{j \neq i} \sum_k \kappa_{in,jok} g_{in} g_{jok} \alpha(in, jok, g, z) + \sum_k \frac{1}{\tau_k} g_{iok} - \sum_k \xi_{i,k} g_{in}$$

$$\frac{dg_{iok}}{dt} = \mu^g(iok, z, g) = \sum_{j \neq i} \lambda_{j,i}(z) g_{jok} - \sum_{j \neq i} \kappa_{jn,iok} g_{jn} g_{iok} \alpha(jn, iok, g, z)$$
$$+ \sum_{j \neq i} \kappa_{in,jok} g_{in} g_{jok} \alpha(in, jok, g, z) - \sum_{j \neq i} \sum_{l \neq k} \kappa_{iok,jol} g_{iok} g_{jol} \alpha(iok, jol, g, z)$$
$$+ \sum_{j \neq i} \sum_{l \neq k} \kappa_{iol,jok} g_{iol} g_{jok} \alpha(iol, jok, g, z) - \sum_{j \neq i} \lambda_{i,j}(z) g_{iok} - \frac{1}{\tau_k} g_{iok} + \xi_{i,k} g_{in} \tag{2.4}$$

In equilibrium, the flows from assets maturing are equal to the flows from new assets being created so: $\sum_i \xi_{i,k} g_{in} = \frac{1}{\tau_k} \sum_i g_{iok} =: \frac{1}{\tau_k} s_k$.

Master equation: We cannot characterize equilibrium in the OTC market using a differential equation for surplus. Instead, we solve the two HJB equations (2.1) and (2.2) combined with belief consistency $\tilde{\mu}^g = \mu^g$ and the KF equations (2.3) and (2.4). To solve the problem with the DeepSAM method, we use neural networks to parameterize $V^n(in, g, z)$ and $V^o(iok, g, z)$, and solve the problem to minimize the weighted loss of equations (2.1) and (2.2) on the sampling data.

3 Endogenous Yield Curve and Financial Crises

We calibrate the model with four types of investors: $\{D, C, U, P\}$, where type D are dealers in the primary bond market, type C are liquidity constrained hedge funds, type U are unconstrained hedge funds, and type P are pension/insurance funds with a long investment horizon. This is reflected in their holding costs. On the asset side, we have four types of bonds with maturities $\tau = 0.25, 1, 5, 10$ years. Unconstrained hedge funds have no holding cost while liquidity-constrained hedge funds have a holding cost of 0.02 across all assets. Pension/insurance funds face holding costs of 0.02 for short maturity bonds ($\tau_1 = 0.25, 1.0$), 0.01 for bonds with $\tau = 5.0$, and no holding cost for long term bonds $\tau = 10.0$. We interpret this as reflecting regulatory constraints or financial frictions that encourage pension/insurance funds to hold long-term bonds.

We consider three aggregate states: good, normal, and bad, where the bad state is interpreted as a financial crisis. We impose that the dealers and pension/insurance funds have constant types over time. By contrast, hedge funds switch from type U to C (i.e. become liquidity constrained) at the rate 0.3 in the good aggregate state, 0.5 in the normal aggregate state, and 0.7 in the crisis aggregate state. In all aggregate states, they switch from type C to U (i.e. become unconstrained) at the rate 0.1. We calibrate other parameters so that the ergodic yield curve matches the average high-grade corporate yield curve over the past 50 years documented by Payne and Szőke (2024) and the haircut rates in the crisis state to match Chen, Cui, He, and Milbradt (2017). We explain the calibration in more detail in Appendix A.

Figure 1a shows the ergodic mean bond yields as a function of maturity at the ergodic steady state of our economy. Evidently, longer maturity bonds have higher yields indicating an upward sloping yield curve. This shape reflects relative investor willingness to hold short and long maturity bonds in the economy. Hedge funds prefer to hold short-maturity bonds because they are worried that they will end up stuck with long-maturity bonds if they become liquidity-constrained. By contrast, the pension/insurance fund prefers to hold long-maturity bonds. Under our calibration, it is the first effect that dominates and so the yield curve is upward-sloping.

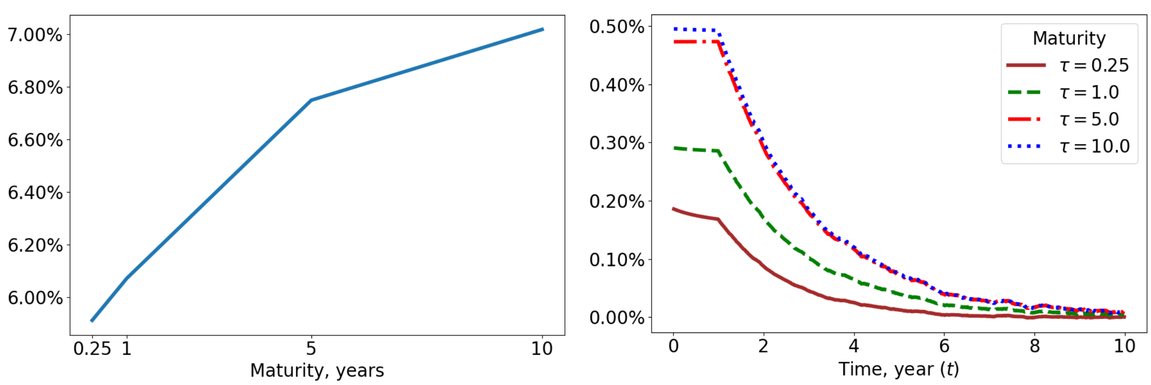

(a) Ergodic Yield Curve (b) IRF of Yield to Maturity by Bond

Figure 1: Yield curve at ergodic steady state and impulse responses. Plot (b) shows the proportional bond yield change compared to the ergodic yield at each maturity following a one-year recession. To calculate the figures, we simulate 3000 paths and calculate the mean.

We use our model to examine the impact of a financial crisis in our OTC bond market. Specifically, the economy starts at the ergodic mean, then moves to the bad "financial crisis" state z_B for one year, and then follows the stochastic z_t process.

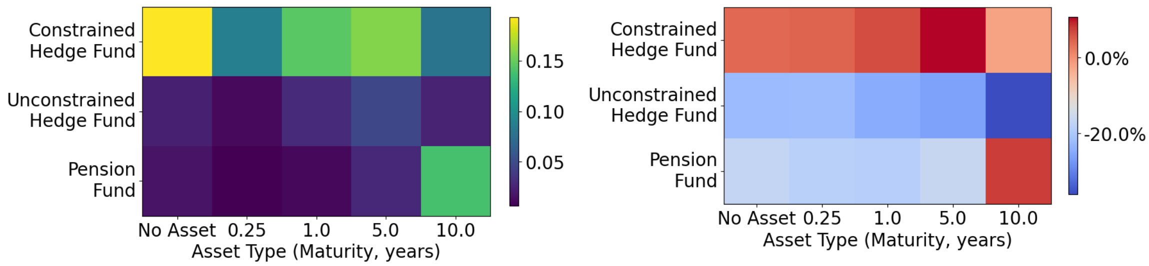

(a) Ergodic Distribution g^{ergodic}

(b) Change after Crisis $\frac{g(z_B)-g^{\text{ergodic}}}{g^{\text{ergodic}}}$

Figure 2: Distribution response to a crisis shock.

Figure 1b shows the impulse responses for bond yields following the shock. For the short-maturity bonds, the yields move very little whereas the yields of long-maturity bonds increase significantly. Figure 2b shows the change in the distribution of investors (relative to the ergodic mean in Figure 2a) when the economy stays in the crisis state for a long time. Evidently, the crisis increases the likelihood that a hedge fund becomes constrained and so increases the proportion of liquidity-constrained hedge funds. This heightens hedge fund concern that they will end up stuck with long-maturity bonds while liquidity-constrained. Consequently, the relative demand for long-term bonds falls during the crisis leading to the relatively large yield increase for long-term bonds in Figure 1b.

Although our OTC model is highly stylized, it illustrates DeepSAM's potential for connecting different asset pricing literatures. By incorporating aggregate crisis risk, institutional heterogeneity, and bond maturities into Duffie et al. (2005), our model introduces default risk premia and term premia into the original OTC model with liquidity premia. This connects OTC models to both traditional bond default models and yield curve models with institutional investors (e.g. Vayanos and Vila (2021)) where, similar to our model, it is the variation in institutional willingness to hold the asset that generates the yield curve. We make these connections while maintaining endogenous trading patterns in our OTC market. That is, agent preferences lead to bonds with particular maturities endogenously emerging as being more or less actively traded and having more of less volatile yields. Ultimately, we believe DeepSAM opens up the possibility to study how a wide range of risks are priced in OTC environments.

4 Conclusion

In this paper, we developed a new method for globally solving and estimating search and matching models with aggregate shocks and heterogeneous agents. This allows us to study dynamics in models where agent decisions depend upon the distribution so the model is not "block-recursive". We believe our methodology is a major breakthrough in the literature of search and matching models and will open up many potential applications in the labor, finance, and spatial literature.

References

Chen, H., R. Cui, Z. He, and K. Milbradt (2017): "Quantifying Liquidity and Default Risks of Corporate Bonds over the Business Cycle," *The Review of Financial Studies*.

Duffie, D., N. Gârleanu, and L. H. Pedersen (2005): "Over-the-counter markets," *Econometrica*, 73, 1815–1847.

Payne, J., A. Rebei, and Y. Yang (2024): "Deep Learning for Search and Matching Models," *Swiss Finance Institute Research Paper*.

Payne, J. and B. Szőke (2024): "Convenience Yields and Financial Repression," .

Vayanos, D. and J.-L. Vila (2021): "A preferred-habitat model of the term structure of interest rates," *Econometrica*, 89, 77–112.

Weill, P.-O. (2008): "Liquidity premia in dynamic bargaining markets," *Journal of Economic Theory*, 140, 66–96.

Appendix

A Additional Details For the OTC Model

A.1 Numerical Illustration

We now consider a calibration of the model that draws on Weill (2008), Chen et al. (2017), Payne and Szőke (2024), and incorporates our agent and asset specification.

Economic parameters: We consider an environment with four types of agents: $\{D, C, U, P\}$, where type D are interpreted as dealers in the primary bond market, type C are interpreted as liquidity-constrained hedge funds, type U are unconstrained hedge funds, and type P are pension/insurance funds with a long investment horizon. Formally, the matrices for holding costs, $\psi(i, \tau)$, switching rates, $\lambda_{ij}(z)$, and primary market participation, $\xi(i, \tau)$, are given in Tables 1, 2, and 3 respectively. The dealer agents (type D) are the only agents who are assigned assets in the primary market. They do not get a net benefit from holding the asset but instead only from trading the asset. The hedge funds randomly switch between getting net benefit from holding any asset (type U) and getting net loss from holding all assets (type C). In this sense, they face the risk of becoming "liquidity constrained" and highly incentivized to sell assets. The pension/insurance funds face a penalty for holding short-maturity assets, interpreted as a regulatory constraint on short-asset exposure.

		Maturity (τ)			
		$\tau_1 = 0.25$	$\tau_2 = 1.0$	$\tau_3 = 5$	$\tau_4 = 10$
Agent Type (i)	D	$\delta\phi(1, z)$	$\delta\phi(2, z)$	$\delta\phi(3, z)$	$\delta\phi(4, z)$
	C	0.02	0.02	0.02	0.02
	U	0.0	0.0	0.0	0.0
	P	0.02	0.02	0.01	0.00

Table 1: Holding costs: $\psi(i, \tau)$.

	D	C	U	P
D				
C			0.1	
U		0.7		
P				

(a) $\lambda(i,j)$ for $z = z_L$.

	D	C	U	P
D				
C			0.1	
U		0.5		
P				

(b) $\lambda(i,j)$ for $z = z_M$.

	D	C	U	P
D				
C			0.1	
U		0.3		
P				

(c) $\lambda(i,j)$ for $z = z_H$.

Agent Type (j) — columns; rows: agent type (i).

Table 2: Switching rates $\lambda(i,j)$ across different aggregate states.

Agent Type (i)	Maturity (τ): $\tau_1 = 0.25$	$\tau_2 = 1.0$	$\tau_3 = 5$	$\tau_4 = 10$
D	ξ_1	ξ_2	ξ_3	ξ_4
C	–	–	–	–
U	–	–	–	–
P	–	–	–	–

Table 3: Primary market participation: $\xi(i,\tau)$.

We consider the following matching rates, which specify that agents can trade more quickly with the dealers than with each other (following Chen et al. (2017)):

$$\kappa_{a,b} = \begin{cases} 50, & \text{if } (a,b) = (in, jok) \text{ and } i,j \neq A, \\ 50, & \text{if } (a,b) = (iok, jok) \text{ and } i,j \neq A, \\ 75, & \text{if } (a,b) = (in, Aok) \text{ and } i \neq A, \\ 0, & \text{if } (a,b) = (iok, Aol) \text{ and } \forall i, \\ 0, & \text{if } (a,b) = (in, jn) \text{ and } \forall i,j, \end{cases}$$

We impose that agents have equal bargaining power unless they match with a dealer, in which case they have bargaining power 0.05 (following Weill (2008) and Chen et al.

(2017)):

$$\beta_{a,b} = \begin{cases} 0.5, & \text{if } (a,b) = (in, jok) \text{ and } i,j \neq A, \\ 0.5, & \text{if } (a,b) = (iok, jol) \text{ and } i,j \neq A, \\ 0.05, & \text{if } (a,b) = (in, Aok) \text{ and } i,j \neq A, \end{cases}$$

The other economic parameters are listed in Table 4. We calibrate the model at the annual frequency. Where possible, we take standard parameters from the literature.

Parameter	Interpretation	Value	Target/Source
ρ	Discount rate	0.05	Average short rate
δ	Bond Coupon Rate	0.01	
Aggregate State: $z \in \{z_B, z_N, z_G\}$			
$\phi(z)$	Coupon haircut	$(0.986, 0.991, 0.997)$	Chen et al. (2017)
$\pi(z)$	Principal haircut	$(0.986, 0.991, 0.997)$	Chen et al. (2017)
$\zeta_{N,B}$, $\zeta_{N,G}$	Transition rate: normal to bad/good	0.1	Crisis every 10 years
$\zeta_{B,N}$, $\zeta_{G,N}$	Transition rate: bad/good to normal	0.5	Average crisis duration 2 years

Table 4: Economic Parameters.

Neural network parameters: We describe the details of the neural network approximation and sampling in Table 5. We use a fully connected feed-forward network with 8 layers, 100 neurons per layer, and a combination of GELU$(\cdot)$ activation functions. The training loss is shown in Figure 3.

Parameter	Value
Number of layers	8
Neurons per layer	100
Activation function	GELU$(\cdot)$
Initial learning rate	10^{-4}
Final learning rate	10^{-6}
Initial sample size per epoch	256
Sample size per epoch	1024
Convergence threshold for target calibration	10^{-7}

Table 5: Neural network parameters

Figure 3: Loss function along training epochs