A Heterogeneous Learning Agent Approach to Financial Stability

Steve Y. Yang[1], Victor Luo[1]

1. Motivation, Focus and Significance

Recent advances in macroeconomics and macro-finance research call for the incorporation of financial frictions and heterogeneity into structural models, as well as credible deviations from the conventional rational expectation paradigm (Christiano, Eichenbaum and Trabandt, 2018[2]). We model the U.S. banking sector network with a heterogeneous agent-based framework where agents (i.e. banks) trade off between exploration and exploitation actions according to reinforcement learning principles to maximize expected long-term rewards. This model features endogenous network formation, establishing a mapping between the empirical distributions of all banks' balance sheets in the financial system to each individual bank's action. We estimate structural parameters of the model using Simulated Method of Moments (SMM) techniques, based on quarterly observed balance sheets data collected by FDIC and SEC covering around 10, 000 banks from March, 2001 to December, 2018. The flexibility of our model enables us to test alternative theoretical explanations on the origin and propagation mechanisms of the systemic risks associated with the 2007-09 financial crisis and European Sovereign Debt Crisis of 2010.

2. Proposed Method and Validation

Existing banking network models delivers valuable economic insights using stylized assumptions (Stiglitz, 2010; Bisias, 2012; Glasserman and Young, 2015b; Roukny et al., 2018). Nevertheless, by abstracting from the obvious heterogeneity across different banks' balance sheets, these models remain silent on how heterogeneity may affect the formation and evolution of banking networks. Our goal is to build models that speak closer to the empirically observed organizational decision patterns that include the following elements: a) performance driven decisions based on financial accounting ratios, such as liquidity ratio, leverage ratio, lending/borrowing ratio, etc.; b) behaviors governed by a reinforcement learning scheme with two key characteristics, i.e. reputation and relationship, to determine to whom to lend and from whom to borrow; c) decision biases/preferences; d) idiosyncratic behaviors based on statistics of individual institutions.

When a bank decides to send a borrowing request to other banks, the borrower considers all the available state variables, Φ(s), to maximize its chance to obtain additional funding. Therefore, value function V(s) along with all the state variables should predict the probability of whether the lenders will extend credits to the borrower. State variables include asset sizes, volume share (Lender Preference Index), excess reserve, amount of loans and funds borrowed, etc. However, the choice of the state variables is one major research task. Upon evaluating the current state, the agent selects an action that maximizes their expected cumulative reward in the long run. The process of selecting an action based on a given state is typically specified in agent's policy. The environment will return a reward signal that agent can use to learn and improve their future evaluation of states. In the proposed model, agents adopt the temporal difference approach, TD(λ), to improve their policy in the current state based on the environment. For example, banks utilize an S-shape function to assess the chance that lending bank i would settle new debts to borrowing bank j, and pick partners. $S_{i,j}^{\{total\}} = \omega \cdot S_{i,j}^{\{relation\}} + (1-\omega) S_{i,j}^{\{reputation\}}$ reputation where $S_{i,j}$ is the score that borrower i assigns to

[1] School of Business at Stevens Institute of Technology. Postal address: 1 Castle Point on Hudson, Hoboken, NJ 03070; Emails: syang14}@stevens.edu and victor.luo@stevens.edu.

[2] Christiano, Lawrence J., Martin S. Eichenbaum, and Mathias Trabandt. "On DSGE models." Journal of Economic Perspectives 32, no. 3 (2018): 113-40.

lender *j*. It is the weighted average of relationship score and reputation score, e.g. rating or capitalization, of bank j. The weights ω determine banks' behaviors.

In our model environment, banking networks endogenously emerge and evolve as a result of decisions made and actions taken by all banks in the financial system. Using the learning agent model and Monte Carlo simulation, we plan to employ inverse reinforcement learning method to establish the connection between the individual bank decisions and the network structures. Flexibility is a significant advantage of our modeling framework. First, it is possible in our framework to characterize banks' policy functions and resulting banking networks even when rational expectation assumption has been relaxed. Second, we can solve for banking network structures under alternative levels of information transparency of the interbank market. Third, we can flexibly disturb the existing banking network structure by placing empirically plausible shocks to different nodes and study the propagation of such shocks in a more realistic setting. In fact, the identification of model parameter depends critically on empirical banking network patterns during the 2007-09 financial crisis and European Sovereign Debt Crisis of 2010.

3. Computational Resource and Data Challenges

However, there exists a significant computational hurdle when we try to calibrate and simulate the model with the real U.S. banking system. The study relies on four major data sources: the FDIC/FFIEC Bank Call Report data, the SEC corporate financial reports, and the Thompson Reuters Financial data, and Bloomberg Market data. In a preliminary work (Liu et al. 2017[3]), we implemented a statistical Bayesian agent model where 6,600 bank agents were implemented and consequently Monte Carlo simulation was conducted to examine the bank's decision implications. The computation time has reached 4 hours with 8 Core CPU for 1 single simulation run. As we are adding adaptive learning agents, the normal workstations are not feasible for model parameter estimation because all the economic agents need to observe large amount of information and make reinforcement inferences for future decisions. Even with the scale-down model, we start to face degraded performance and lose fidelity to the real banking system. We therefore propose to develop a Python package and run the model on Amazon Web Services (AWS) Cloud. We plan to apply AWS Cloud Credits for Research and make the data and basic models available through an open source license agreement. The goal of this open source Python package is to allow other academic researchers and regulators to conduct new research and perform system stress testing.

4. Outcome and Broad Impact

By providing an explicit modeling of the banking sector, we hope to contribute to policy discussion on the merits or drawbacks of imposing regulations on banks' balance sheet or banking network structure. For example, the U.S. Federal Reserve has recently proposed "single-counterparty credit limits" policy that specifically target and/or impose network structure. Through the lenses of our model, we can investigate the effect of imposing balance sheet holding/trading regulation on different asset classes, i.e., MBS, ABS, CDS, IRS. Each of these asset classes will have certain effects on banks' balance sheets. The likelihood of systemic risk will be calculated, and mitigation measures will be proposed and assessed using this large-scale computational model. Different mechanisms that will trigger detrimental systemic events will be investigated and a tool for bank regulators to stress-test the interbank system through examining the severity of outcomes will be developed.

---

[3] Liu, Anqi, Mark Paddrik, Steve Y. Yang, and Xingjia Zhang. "Interbank contagion: An agent-based model approach to endogenously formed networks." Journal of Banking & Finance (2017).