"ESTIMATING POLICY FUNCTIONS IN PAYMENTS SYSTEMS USING REINFORCEMENT LEARNING"

BY

P. S. CASTRO A. DESAI H. DU R. GARRATT F. RIVADENEYRA

Discussion by

MICHAEL GOFMAN

University of Rochester

October 21, 2021

ECONOMICS OF PAYMENTS X

This Paper

- New methodology to estimate banks' policies in a large value payment system.
- Reinforcement learning can successfully approximate analytical solutions for a simple two period model with two banks.
- Promising results for a more complex environment with two banks and multiple periods.
- An interesting case study for how AI can be used in large value payment systems. Showcases AI's ability to make financial decisions under uncertainty.

Illustrative Example

- 2 banks – Bank A and Bank B
- Bank A needs to send $10B to Bank B at 8:01AM
- Bank B needs to send to Bank A $10B at 8:02AM
- Common knowledge about the payment needs
- Q: How much collateral should Bank A pledge at 8:00AM?
- A: $10B
- Q: How much collateral should bank B pledge at 8:00AM?
- A: Northing! Because Bank B uses the collateral from Bank A.

  Bank B free-rides on the liquidity provided by Bank A.

Economic trade-offs in the model

- Collateral is costly so pledging less is better
- Delays in sending payments are even more costly so pledging more is better, unless there is a reasonable chance that there is an incoming payment that is backed by other bank's collateral ("free collateral")
- Banks want to reuse the free collateral from another bank rather than to pledge their own collateral
- The timing of the payments is not (fully) predictable, and the actions of other banks are not observable
- The paper shows that reinforcement learning allows AI to account for the above trade-offs

Comments

Large Value Transfer System (LVTS)

- Canadian LVTS has two tranches to send payments
  - Tranche 1 (RTGS): high stability, low efficiency (collateral overuse)
  - Tranche 2 (DNS): high efficiency, low stability (credit limits instead of collateral)
    - Banks mainly use Tranche 2. The paper models Tranche 1 only.
    - The true risk is that banks cut credit limits, collateral constraints become binding and critical payments cannot be made, causing a high-frequency systemic risk (Chapman, Gofman, Jafri 2019)
    - Stress measures: rejected payments, delayed payments, credit limit reductions, delays of credit limits extension.
- Can reinforcement learning be applied in a hybrid system like LTVS? Can it help to predict jumps in systemic risk? Can it help regulators to decide when to inject liquidity by accepting broader collateral?

Should AI learn from Humans?

- The key question for training an AI model is what is the objective function. Do we try to replicate/ approximate human behavior? Do we try to create a superhuman AI system? How do we measure success? Do we want to have an autonomous AI-based large value payments system? Is it better than a hybrid system where both humans and AI make decisions?

- Do human operators of LVTS play an optimal equilibrium strategy in a multi-bank game of liquidity hoarding? Should we train AI to approximate this strategy? What if it is suboptimal?

- There is a significant number of failed payments that should never happen in equilibrium. Does it suggest that humans are not perfect?

- If there is a good and a bad equilibrium of the liquidity injection game or a good and a bad equilibrium of the credit limits provision game, should be train AI to avoid bad equilibria? Can AI be better in coordination?

Cooperation or Competition?

- Do Banks play a non-cooperative game in LVTS?
- Evidence that large Canadian banks in LVTS injected extra liquidity at times of systemic stress (Chapman, Gofman, Jafri 2019).
- Inconsistent with the goal to hoard liquidity at the beginning of the day.
- Banks want clients to receive payments as much as they want clients to send payments.
- Banks do not want counterparties to fail to repay them their loans because they lack the collateral.
- Banks are not interested in the system's failure to process payments.
- Can AI help with collateral management at the system level?

Model Misspecification

- 2 banks vs. many banks (Fedwire or Target2 have thousands)
- Banks can pledge more collateral during the day
- Payments come continuously. Some are urgent. We observe only sent payments, do not observe delays.
- Interbank market for liquidity.
- Borrowing from the Bank of Canada at the end of the day requires collateral.
- Will AI be able to deal with the real-life setting? Can we learn from AI's solution to a simplified environment? What if we feed AI with less data than what human operators have?

Conclusion

- Reinforced Learning has been very successful in different domains (Chess, Go).
- It is great to see the technique is used for estimating policy functions for initial collateral pledging to LVTS. Very promising research agenda with a lot of potential.
- I hope to see more AI work on estimating size and timing of payment needs, predicting rejected payments, delays in payments and collateral shortages