Strategic Interaction Among AI Agents: Evidence from Payment Systems*

Furkan Danisman[1] Ajit Desai[2] Janet Jiang[2] Francisco Rivadeneyra[2] Talha Azad[2]

[1]University of Toronto & Vector Institute for AI

[2]Bank of Canada

Bank of Finland Seminar on Financial Market Infrastructures, August 2026

*The views expressed are solely those of the authors and do not necessarily reflect the views of the Bank of Canada.

Outline

1. Background & introduction
2. LLM agents experimentation setup
3. Results
   * Strategic capabilities of LLM agents
   * Role of strategy prompting
   * Environment with multiple equilibria
4. Summary and next steps

Background & introduction

AI is accelerating at pace

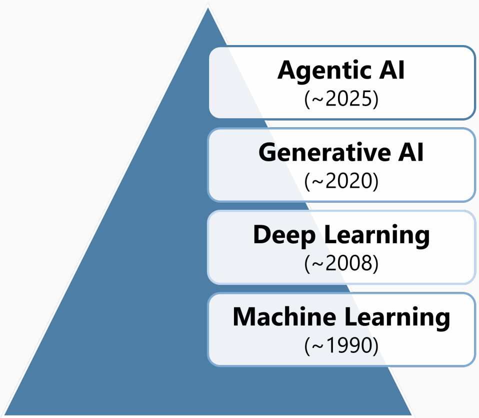

& learn to **act**

& learn to **create**

& learn to **perceive**

Models learn to **predict**

Agentic AI: a system that proactively plans, decides, and acts autonomously

It is an **AI system** with following core characteristics:

- **Interpretation**: understands user intent via prompt
- **Reasoning**: interpret context and formulate strategies
- **Planning**: decomposes the task into small steps
- **Tool use**: call APIs, databases, internet, etc.
- **Action and feedback**: executes, observes, and adapts

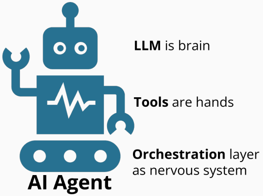

Source: https://docs.cloud.google.com/architecture/agentic-ai-overview

Agentic AI in financial systems: an emerging policy agenda

- FSB (2024): Financial stability implications of artificial intelligence
- FSB (2025): Monitoring adoption of artificial intelligence and related vulnerabilities in the financial sector
- FSB (2026): Sound practices for responsible adoption of artificial intelligence

- BIS Project Logos (2026): observing the behaviour of LLM-based agents in a simulated financial market environment

AI agents in interactive financial systems

- Many financial systems involve multiple financial institutions interacting: one FI decision affects the state, opportunities, or incentives faced by others.
  - Electronic trading
  - Repo and secured funding markets
  - Securities settlement and collateral management
  - Large value payment
- What happens as FIs deploy increasingly autonomous AI in settings where decisions interact with those of other FIs? How can its behavior be directed, traced, monitored, and explained?

What we do

Proof of concept of general purpose LLMs as cash managers in a stylized large value payment system:

- **Capability:** Can agents make strategically sophisticated decisions in an interactive system?
- **Instruction design (prompt):** How do instructions shape behavior and system outcomes?
- **Traceability:** Can recorded rationales help monitor and explain decisions?

Large value payment system

Core national financial infrastructure for settling large-value obligations between FIs

- Most systems are based on real-time gross settlement (**RTGS**)
- Lynx in Canada, T2 in Europe, and Fedwire in US
- Lynx: daily volume 55k; daily value $350 billion

LLM agents experimentation setup

Liquidity management in RTGS

Cash managers face a liquidity–delay trade-off in payment systems, which they manage through **strategic decisions** on:

- **Liquidity allocation**:
  amount of collateralized liquidity to secure
- **Payment choices**:
  pace at which payments are processed

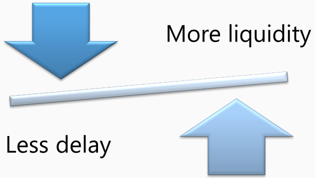

Liquidity and payment decisions: setup

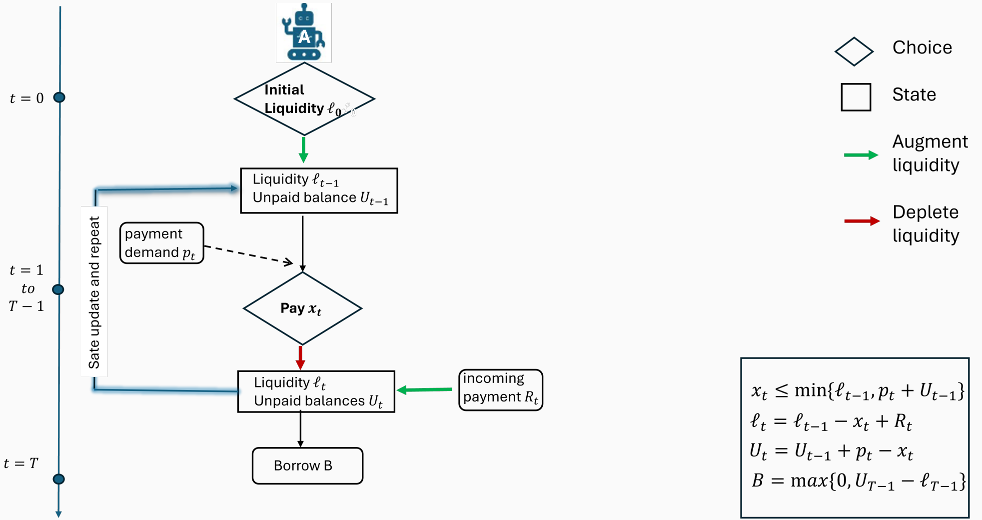

$$x_t \leq \min\{\ell_{t-1}, p_t + U_{t-1}\}$$
$$\ell_t = \ell_{t-1} - x_t + R_t$$
$$U_t = U_{t-1} + p_t - x_t$$
$$B = max\{0, U_{T-1} - \ell_{T-1}\}$$

Liquidity and payment decisions: costs

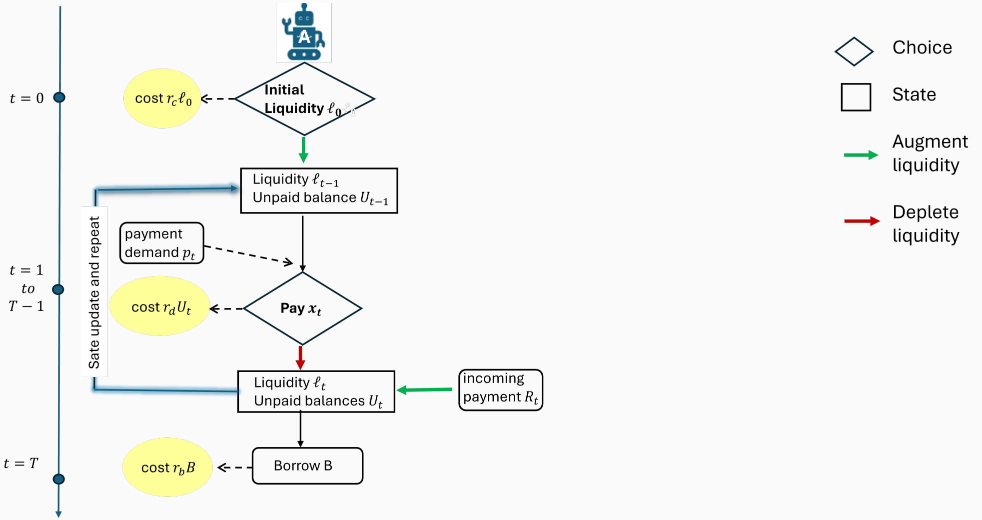

Liquidity and payment decisions: game

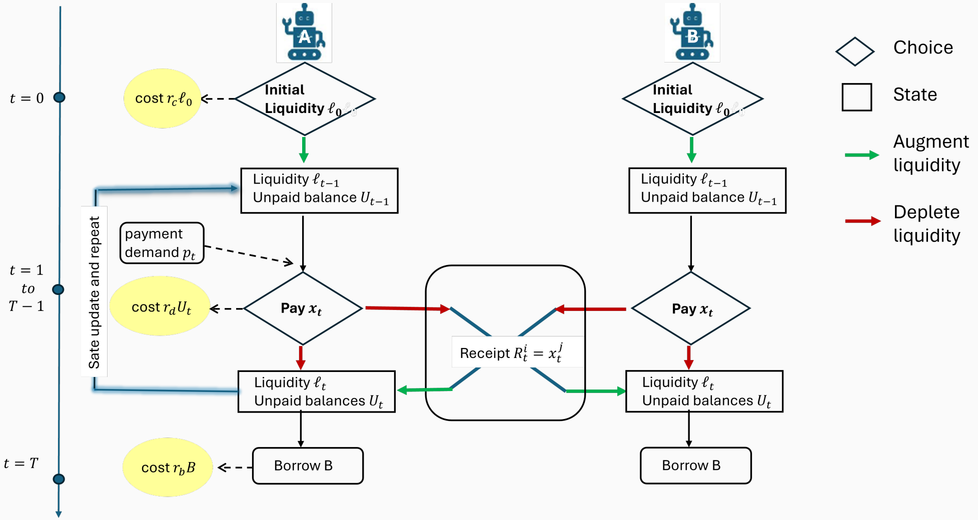

LLM implementation framework: ChatGPT 5.2 reasoning model

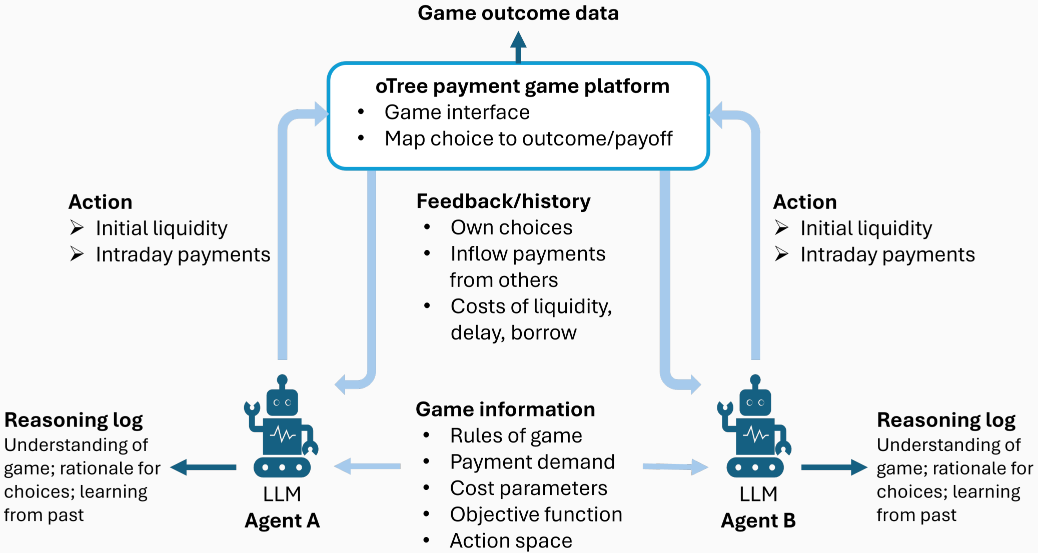

Payment game

- Two agents A,B; Periods: 0,1,2,3
- Payment profile $[p_1, p_2]$ for both agents
- Period 2 payment is lumpy for agent B: pay either full or nothing
- Subgame perfect pure-strategy Nash equilibrium (SPE)
- Payment strategy given liquidity
  - A: pay as much as possible given available liquidity in both periods
  - B period 2: if enough liquidity to pay indivisible payment, pay it and use rest to pay carry-over divisible payment; otherwise, max out carry-over divisible payment
  - B period 1: strategically chooses how much to pay/retain

Treatments

Treatment	Payment profile (p_1, p_2)	Cost parameters (r_c, r_d, r_b)	No. of eqa
1: Strategy-informed	(1.6, 2)	(0.8, 0.2, 0.5)	1
2: Rules-only	(1.6, 2)	(0.8, 0.2, 0.5)	1
3: Multiple equilibria	(2, 2)	(0.8, 0.2, 0.38)	4

- Treatments 1 and 2: unique eqm featuring strategic delay ($\ell_0^B > x_1^B$)
- Treatment 3: multiple eqm; two with strategic delay
- Each session connects two AI agents interacting with each other and repeats the game for 50 rounds

LLM: System prompt sample (all treatments)

"You are a cash management agent in a large-value payment system to manage liquidity and payments in a series of four-period games. You are playing the game with another cash management agent. Each four-period game is called a round. Your goal is to minimize the total cost in each round by optimizing liquidity and payment decisions while learning from previous rounds. Always respond in valid JSON with two keys:

- **Reasoning**: a concise justification of your choice, focusing on strategic choices. Explain the steps when making algebraic calculations.
- **Answer**: your final decision (e.g., "Yes" to confirm understanding)

LLM: Task prompt sample (all treatments)

In each experiment, acting as an cash manager, you will play a repeated multi-agent cash-management game with a specified payment profile.

At the beginning of each round, you choose an initial liquidity level, which is costly but can be used to settle payment obligations in subsequent periods.

...

Your objective is to minimize total round cost, defined as the sum of initial liquidity cost, delay cost, and final borrowing cost.

LLM: Strategy-informed prompt (treatments 1 and 2)

Interaction with the Other Agent

- Your inflows R_1 and R_2 depend on the other agent's liquidity and payment decisions. Conversely, your payments affect the other agent's inflows and available liquidity.
- This creates reciprocity across periods. In particular, your period-1 payment, x_1, can affect your own inflow in period 2, R_2, through the following mechanism:
  1. Period 1: the payment you send, x_1, becomes the other agent's inflow R_1, increasing their liquidity carried into period 2 (denoted ℓ_1).
  2. Period 2: if that additional liquidity relaxes the other agent's liquidity constraint, they may be able to pay you more, so part or all of x_1 can come back to you as higher R_2. If the other agent is not liquidity-constrained, sending extra in period 1 will not change your R_2.

LLM: Strategy-informed prompt (treatments 1 and 2)

"If you consistently end up with *idle* liquidity by the end of the round, or $\ell_3 > 0$, this is a signal that you may be holding more initial liquidity than necessary. In that case, consider reducing your initial liquidity choice ℓ_0 in future rounds, as lowering ℓ_0 reduces the upfront liquidity cost without increasing delay or borrowing costs (as long as you can still meet payment and settlement needs)."

For the agent who faces a lumpy payment in treatment 2, we remind it "Do not ignore strategies to keep part of ℓ_0 in period 1 to ensure the payment of indivisible p_2 in period 2."

Results

Treatment 1

- Unique equilibrium featuring strategic delay
- Strategy prompt

$\rightarrow$ Strategic capabilities of LLM agents

Treatment 1: decisions

- LLM Session 2 choices (showing 20 rounds; flat after that)
- Initial liquidity "o"; period 1 payment "$\star$"

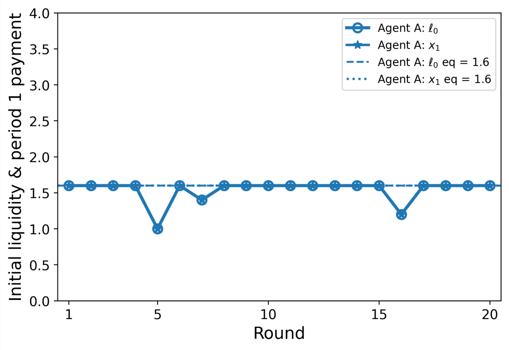

Agent A

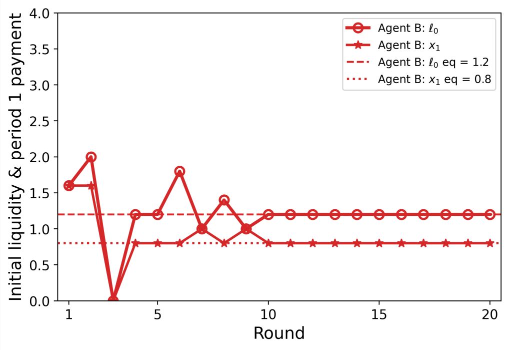

Agent B

Treatment 1: results by session

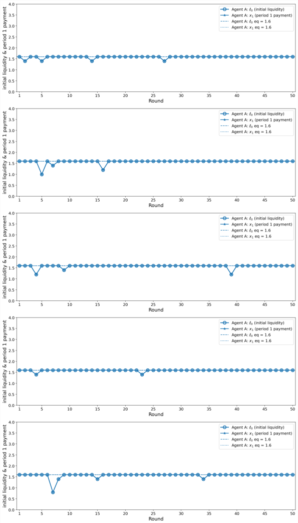

Agent A

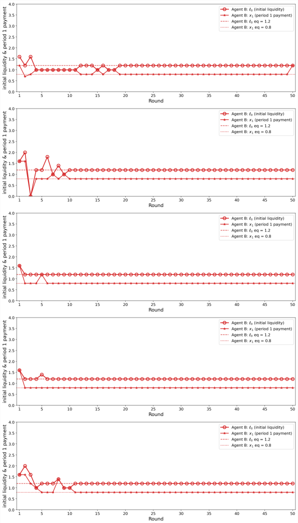

Agent B

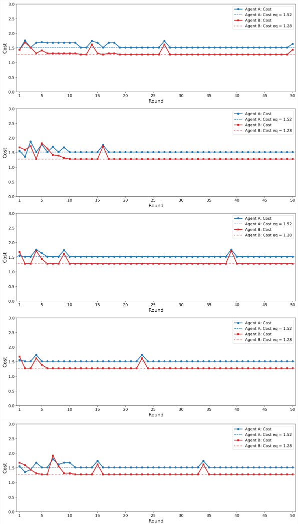

Costs

Treatment 1: results by session

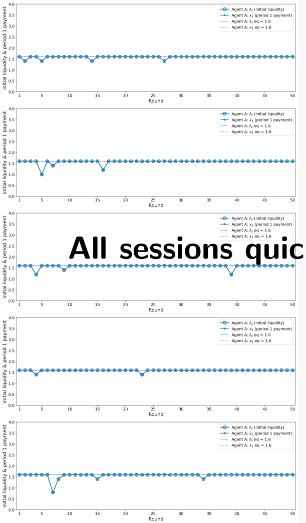

Agent A

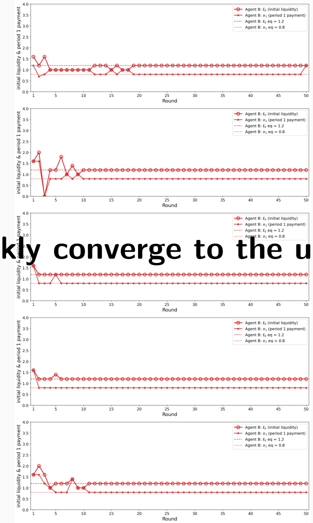

Agent B

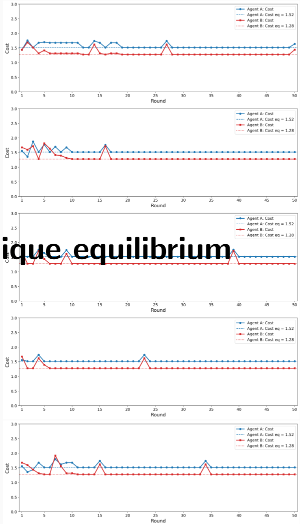

Costs

All sessions quickly converge to the unique equilibrium

Treatment 1: reasoning → considered factors

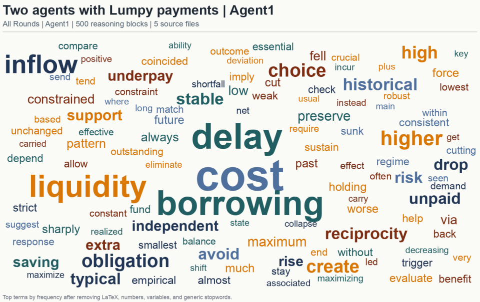

Agent A

Agent B

Treatment 1: reasoning $\rightarrow$ strategies

Traceability and explainability

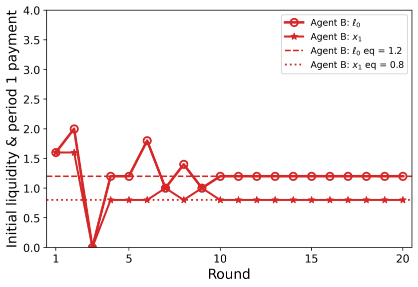

- **Round 1**: no history, assume $R_1 = R_2 = 0$, compare $\ell_0 = 0$ and 1.6, choose 1.6
- **Round 2**: see $R_1 = 1.6$, set $\ell_0 = 2$, $x_1 = 1.6$ to pay both p_1, p_2 in full
- **Round 3**: see $R_1 = R_2 = 1.6$ and assume they will continue; choose $\ell_0 = x_1 = 0$ to save on liquidity; got $R_2 = 0$ and an increase in delay and borrowing costs
- **Round 4**: recognize $x_1 = R_2$; assume R_1 continues at 1.6, choose $\ell_0 = 1.2, x_1 = 0.8$ to minimize total cost

Treatment 2: role of strategy prompting

- Same game as treatment 1: single equilibrium
- Remove strategy prompting

Treatment 2: results by session

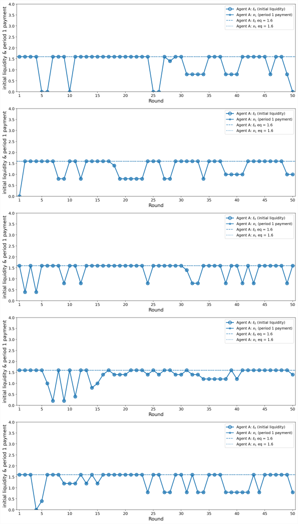

Agent A

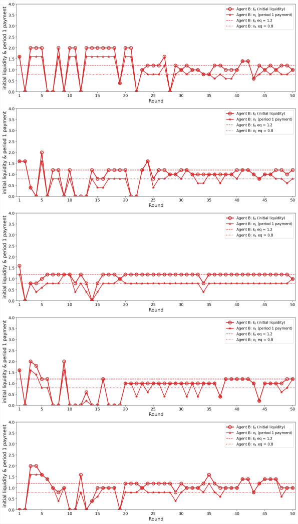

Agent B

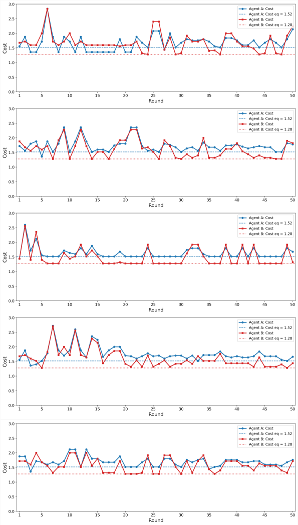

Costs

Treatment 2: results by session

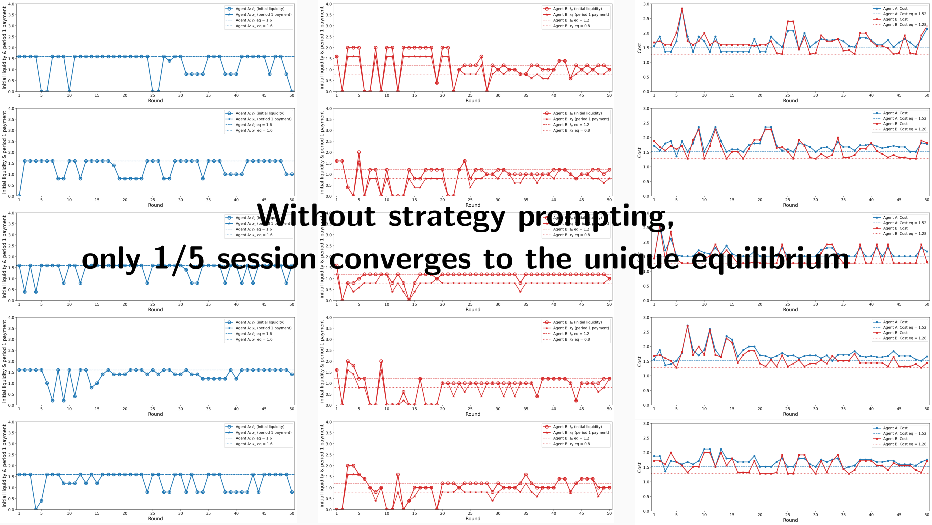

Without strategy prompting, only 1/5 session converges to the unique equilibrium

Agent A

Agent B

Costs

Treatment 2: role of strategy prompting

B's action or mention	treat 1	treat 2
mention reciprocity	21%	3%
mention $R_2 \approx x_1$	44%	33%
engage in strategic delay ($\ell_0 > x_1$)	90%	51%
mention keep liquidity for p_2	14%	3%

Treatment 3: multiple eqa

- Prompting same as treatment 1; parameterization leads to multiple eqa
- Results

Session	Long-run outcome	
1	$\ell_0^A = x_1^A = 0.2$, $\ell_0^B = x_1^B = 1.8$	Equilibrium 3
2	In the neighborhood of Equilibria 3 and 4	Near Equilibrium 3/4
3	$\ell_0^A = x_1^A = 1.8$, $\ell_0^B = 1.2, x_1^B = 1$	Equilibrium 1
4,5	$\ell_0^A = x_1^A = 0$, $\ell_0^B = x_1^B = 2$	Equilibrium 4

Treatment 3: results by session

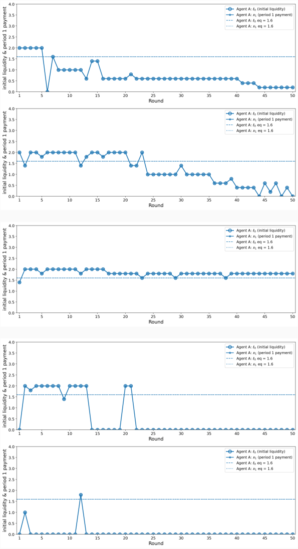

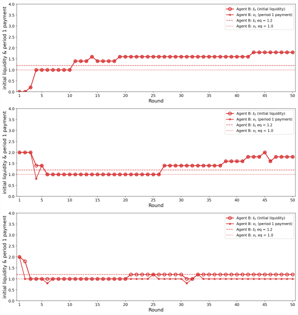

Scenario 3 – Agent B

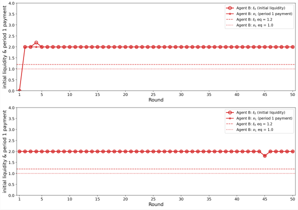

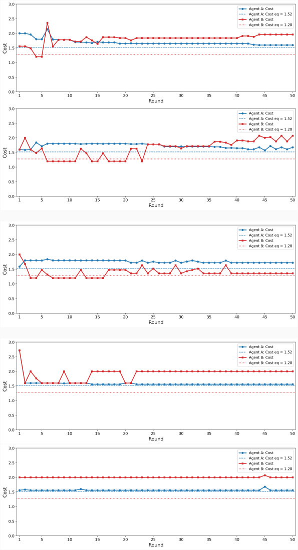

Agent A

Agent B

– Costs

Treatment 3: results by session

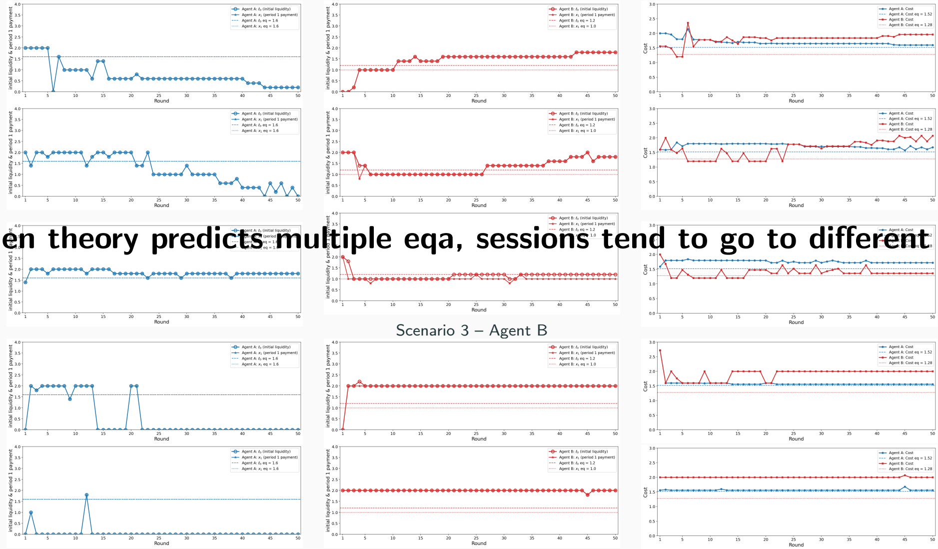

When theory predicts multiple eqa, sessions tend to go to different eqa

Scenario 3 – Agent B

Agent A

Agent B

– Costs

Supplementary experiments

Two agents with divisible payments

- Same parameterization as treatment 3
- Difference: all payments are divisible $\rightarrow$ no strategic delay

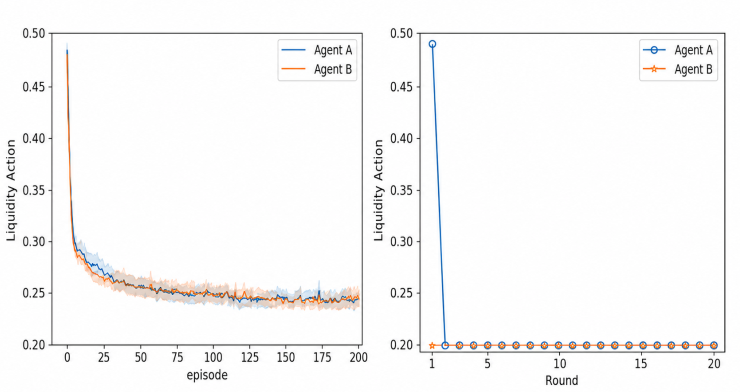

Fig: Initial liquidity choice of RL in Castro et al. (2025) (left) and LLM (right)

Three agents with divisible payments

Agents: A, B, **and** C:

- $p_{AB} = [1, 2, 0]$; $p_{AC} = [1, 0, 0]$;
- $p_{BA} = [0, 1, 2]$; $p_{BC} = [0, 0, 0]$;
- $p_{CA} = [0, 0, 2]$; $p_{CB} = [0, 0, 0]$;
- $r_c = 0.8$, $r_d = 0.2$, $r_b = 0.38$
- In eqm, agent A prioritizes payment to B in period 1 because the liquidity will be returned in period 2

	ℓ_0^A	ℓ_0^B	ℓ_0^C	x_1^{AB}	x_1^{AC}	C_A	C_B	C_C	$C_A + C_B + C_C$
eqm	1	0	0	1	0	2.36	0.76	0.76	3.88

Three agents with divisible payments: LLM Agent A choices (session 1)

- Initial liquidity: *O*
- Payment to B in period 1: -B-
- Payment to C in period 1: -C-

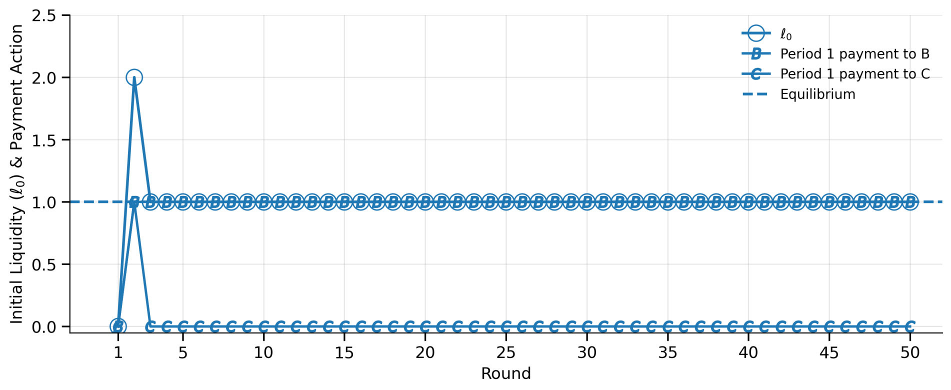

Summary and next steps

Summary

- AI agents can quickly learn constraints and incentives and adopt strategically coherent policies to respond to observed counterpart behaviors
  - Understand trade-off of the cost of liquidity, delay and borrowing
  - Understand liquidity recycling and interdependence between participants
  - Combine algebra and inferences from historical patterns
  - Use strategies to retain a liquidity buffer for future lumpy payments
  - Prioritize payments that will return in subsequent periods
- Strategy prompting is important
- Reasoning text important for traceability and explainability

Next steps: other aspects of RTGS

- Delay propagation
- Gridlock dynamics
- Measures to reduce gridlock
- Can LLMs be prompted to achieve more desirable outcomes?

Next steps: Project Agenti (jointly with BISIH, BoC and SNB)

Aims at developing an **AI Payment Simulator (AIPS)** – a controlled, research-grade environment that enables realistic modelling of payment system participants:

	Present	Future
Operations	**RTGS operations** *Human-operated payment systems.* Real-world settlement where banks' cash managers make release, prioritisation, and liquidity decisions manually.	**Autonomous RTGS ecosystem** *AI-driven payment systems with human oversight.* Payment and liquidity decisions executed by agentic AI in production, with humans supervising outcomes and governance frameworks.
Simulations	**RTGS simulators** *Machine models of today's system.* Tools like the Bank of Finland simulator replicate current payment system dynamics using rule-based logic to analyse liquidity, gridlock, or operational scenarios.	**Project Agenti** *The future of RTGS participant simulation.* Central banks use AIPS to capture more dynamic participant behavior, evaluate how autonomous financial agents may affect systemic risk, and stress-test policies prior to implementation.

Objectives of Project Agenti

- **Understand future liquidity dynamics**: How can central banks prepare for a future in which agentic AI changes intraday liquidity behavior?
- **Identify emerging risk channels**: Automated behaviors could alter offsetting, gridlock formation, and shock transmission.
- **Strengthen stress testing**: Dynamic, adaptive agent behavior enables more realistic stress-scenario analysis.
- **Inform policy and design choices**: A safe sandbox lets central banks test liquidity tools and guardrails before the private sector scales AI adoption.

Thank you!

Literature on liquidity-delay trade-off

Different methods to **study and mitigate** the trade-off between liquidity and delay:

- Game theory model (Bech and Garratt, 2003)
- Agent-based model (Galbiati and Soramäki, 2011)
- Auction-based simulations (Ates et al., 2025)
- Reinforcement learning (Castro et al., 2025)
- Generative AI (Aldasoro and Desai, 2025)

$\Rightarrow$ This paper: explores potential of **LLM agents** for strategic cash management

Standard setup for cash management in wholesale systems

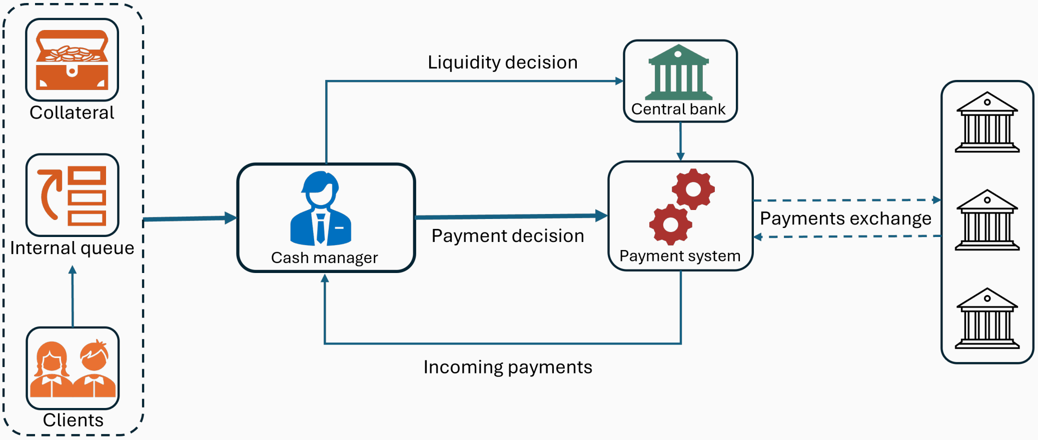

Deep Reinforcement Learning (DRL) VS LLM Agents

Agent type: DRL vs LLM:

- **State:** numerical, structured and sparse **vs** richer contextual inputs, including rules, constraints as numbers, equations and text
- **Objective:** explicit per-round numerical reward **vs** cost parameters and policy objectives expressed in natural language
- **Training:** network updates through repeated training **vs** contextual updating through prompts and histories

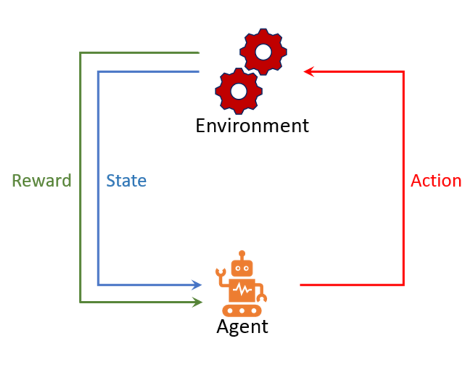

References

Aldasoro, I. and A. Desai (2025). AI agents for cash management in payment systems. Staff Working Papers 35, Bank of Canada.

Ates, C., M. L. Bech, R. Garratt, and M. N. Petrič (2025). Auction-based liquidity saving mechanisms.

Bech, M. L. and R. Garratt (2003, April). The intraday liquidity management game. *Journal of Economic Theory 109*(2), 198–219.

Castro, P., A. Desai, H. Du, R. Garratt, and F. Rivadeneyra (2025). Estimating policy functions in payment systems using reinforcement learning. *ACM Transactions on Economics and Computation 13*(1), 1–31.

Galbiati, M. and K. Soramäki (2011). An agent-based model of payment systems. *Journal of Economic Dynamics and Control 35*(6), 859–875.