Memory-Bounded Reciprocity in Decentralized Exchange Networks

Effect of memory on market fairness

Lars Paulussen | Supervisors: Georgios Iosifidis, Livia Chatzieleftheriou
TU Delft – EEMCS – CSE3000 Research Project

1. Motivation

- Peer-to-peer systems share resources through reciprocal exchange
- Reciprocity requires remembering past interactions
- Existing models assume unlimited history
- IoT devices and AI agents have limited memory

Take-away: Real systems cannot store infinite interaction histories

2. Research Gap

Previous work studied:

- Full-history memory
- Exponentially decaying memory

Missing:

- Hard memory constraints
- Fixed-size sliding windows

Take-away: We study what happens when agents remember only the last m_i interactions

3. Research Question

How does bounded memory affect fairness in decentralized resource exchange networks?

Sub-questions:

- **SQ1**: Effect of memory size?
- **SQ2**: Most robust strategy?
- **SQ3**: Effect of heterogeneous memory?
- **SQ4**: Effect of a single outlier?
- **SQ5**: Influence of topology?

4. Model

- Network: $G = (\mathcal{N}, \mathcal{E})$
- Agents exchange random surpluses: $D_i(t) \sim U(\alpha, \beta)$
- Sliding-window memory:

$$\hat{x}_{ji}^{m_i}(t) = \frac{1}{m_i} \sum_{\tau=t-m_i}^{t-1} x_{ji}(\tau)$$

- $x_{ji}(\tau)$: amount sent from j to i in round τ
- m_i: number of past rounds agent i retains

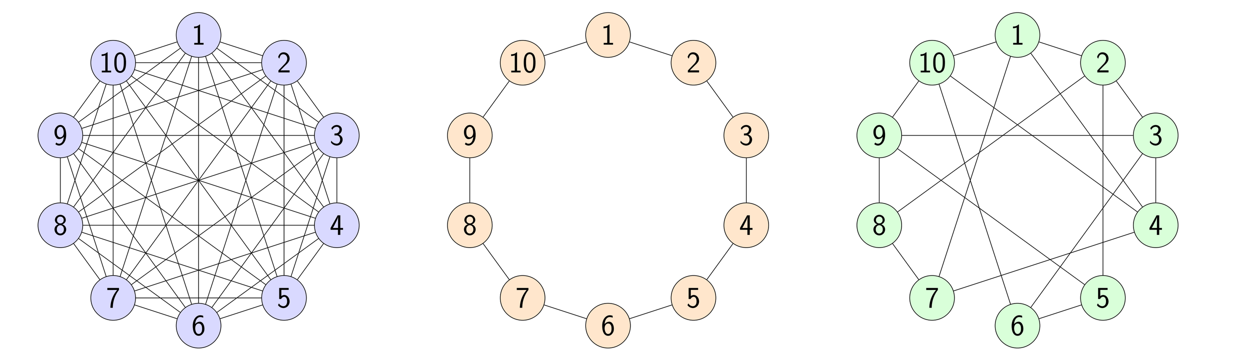

Complete (K_{10}) Ring (C_{10}) Erdős–Rényi ($G(10, 0.4)$)

Take-away: Memory cost is bounded, only takes the last m_i rounds into consideration

5. Reciprocity Strategies

1. Egalitarian (ϵ): split equally (memoryless)

$$x_{ij}^{\epsilon,m_i}(t) = \frac{D_i(t)}{|\mathcal{N}_i|}$$

2. Proportional (ϕ): split proportional to the amount received

$$x_{ij}^{\phi,m_i}(t) = D_i(t) \cdot \frac{\hat{x}_{ji}^{m_i}(t)}{\sum_{k \in \mathcal{N}_i} \hat{x}_{ki}^{m_i}(t)}$$

3. Punish & Reciprocate (ψ): proportional, but only consider agents who sent over average
   - This is the set $\mathcal{R}_i = \{j : \hat{x}_{ji}^{m_i} \geq \frac{1}{|\mathcal{N}_i|} \Sigma_{k \in \mathcal{N}_i} \hat{x}_{ki}(t)\}$

$$x_{ij}^{\psi,m_i}(t) = \begin{cases} D_i(t) \cdot \frac{\hat{x}_{ji}^{m_i}}{\sum_{k \in \mathcal{R}_i} \hat{x}_{ki}^{m_i}} & j \in \mathcal{R}_i \\ 0 & \text{otherwise} \end{cases}$$

4. Greedy (π): send to the agent who received the least compared to what they contributed
   - Every agent announces $\hat{\rho}_i^{m_i}(t) = \frac{1}{m_i} \Sigma_{\tau=t-m_i}^{t-1} \Sigma_{j \in \mathcal{N}_i} x_{ji}(\tau) / \bar{D}_i$ each round

$$x_{ij}^{\pi,m_i}(t) = \begin{cases} D_i(t) & j = \arg\min_{k \in \mathcal{N}_i} \hat{\rho}_k^{m_k}(t) \\ 0 & \text{otherwise} \end{cases}$$

Performance metric (minimal sharing ratio):

$$\eta_{\min} = \min_i \rho_i = \min_i \left(\frac{\bar{R}_i(t)}{\bar{D}_i} \right)$$

Measures fairness of the worst-off agent.

Design goal: try to maximize the minimal sharing ratio

6. Experimental Setup

- 20 agents
- 5000 simulation rounds
- 30 random seeds
- Memory: $m \in \{1, 2, 5, 10, 20, 50, 100, 200, 500, 5000\}$
- Complete(, Ring and ER) graphs

Question: How much fairness is lost when memory shrinks?

7. Memory Sensitivity (SQ1)

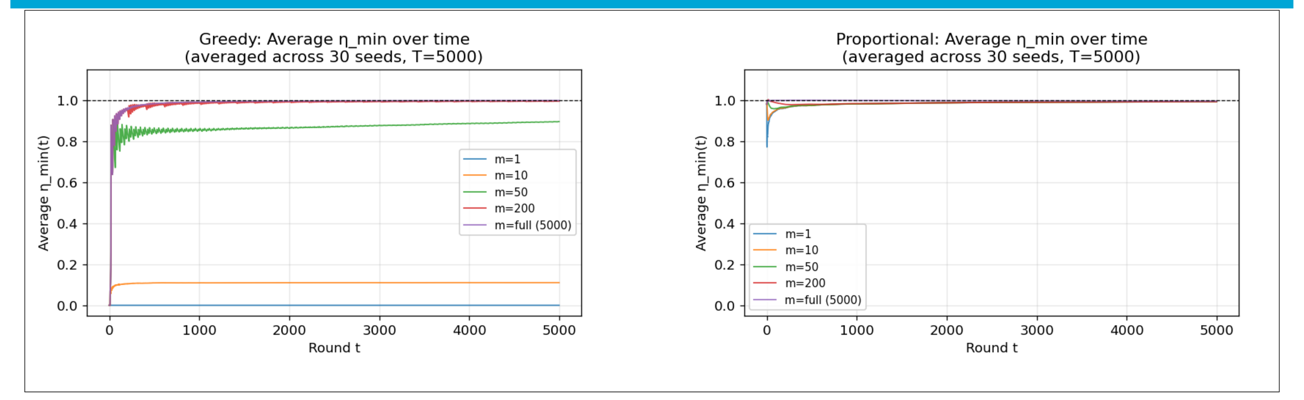

Key observations:

- Proportional reciprocity fairness is long term not dependent on memory size, while greedy is.

Take-away: Proportional converges easily, independent of memory size and greedy needs enough memory to converge quick to fair situation.

8. Which Strategy Is Best? (SQ2)

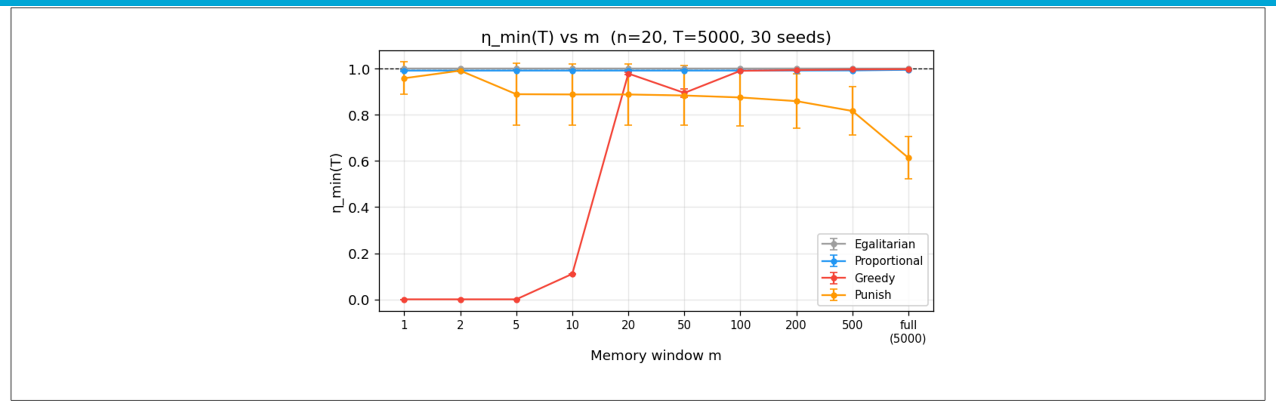

Key observations:

- Proportional remains near 0.99 fairness
- Greedy collapses for small memory, but has a sharp threshold around $m \approx 10$
- Punish-and-reciprocate performs worst with full memory

Take-away: Proportional is the most robust strategy across all memory budgets

9. Memory Heterogeneity (SQ3)

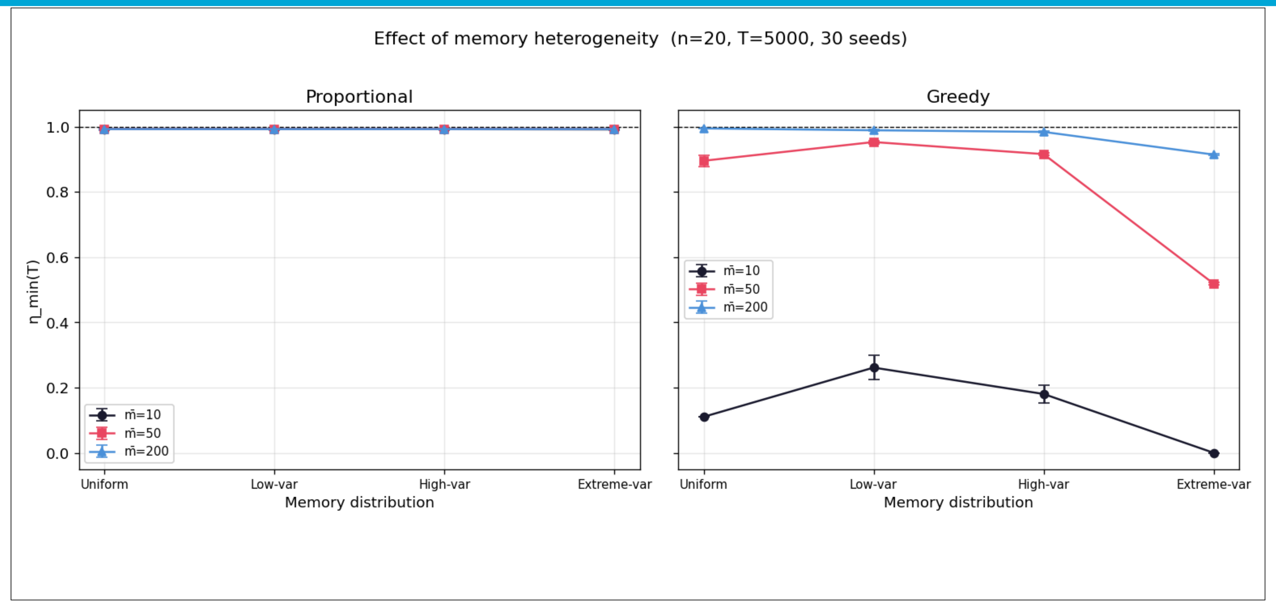

- **Low-var**: agents uniformly between $[\bar{m} - \bar{m}/4, \bar{m} + \bar{m}/4]$
- **High-var**: agents uniformly between $[\bar{m} - \bar{m}/2, \bar{m} + \bar{m}/2]$
- **Extreme-var**: agents uniformly between $[1, 2\bar{m} - 1]$

Key observations:

- Proportional fairness largely variance-independent, greedy is dependent
- A little variance can outperform uniform population for greedy compared to no variance
- Greedy is still very reliant on enough base memory

Take-away: This market performs best when only a little variance is present

10. Outlier Effects (SQ4)

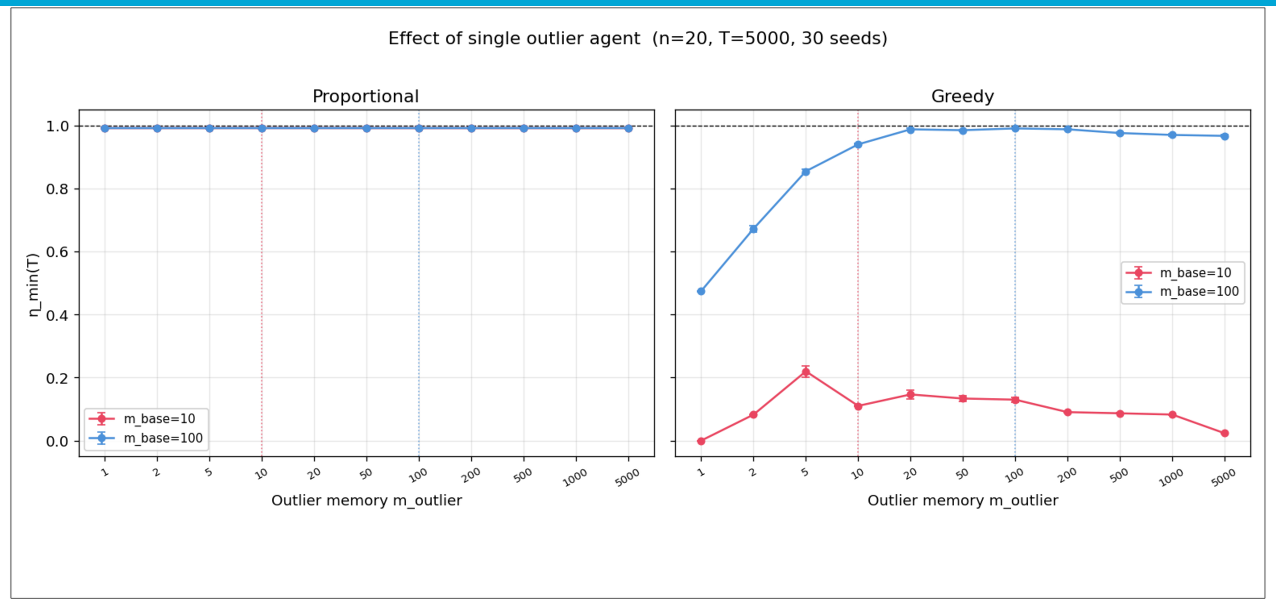

Key observations:

- Proportional fairness is outlier-independent
- A single greedy outlier at $m_{\text{out}} = 1$ reduces $\eta_{\min}(T)$ by 52%, but a very high-memory outlier (i.e. $m_{\text{out}} = 5000$) also mildly hurts performance

Take-away: For greedy, both very low- and very high-memory outliers tend to worsen fairness. Proportional is immune in both directions

11. Topology Effects (SQ5)

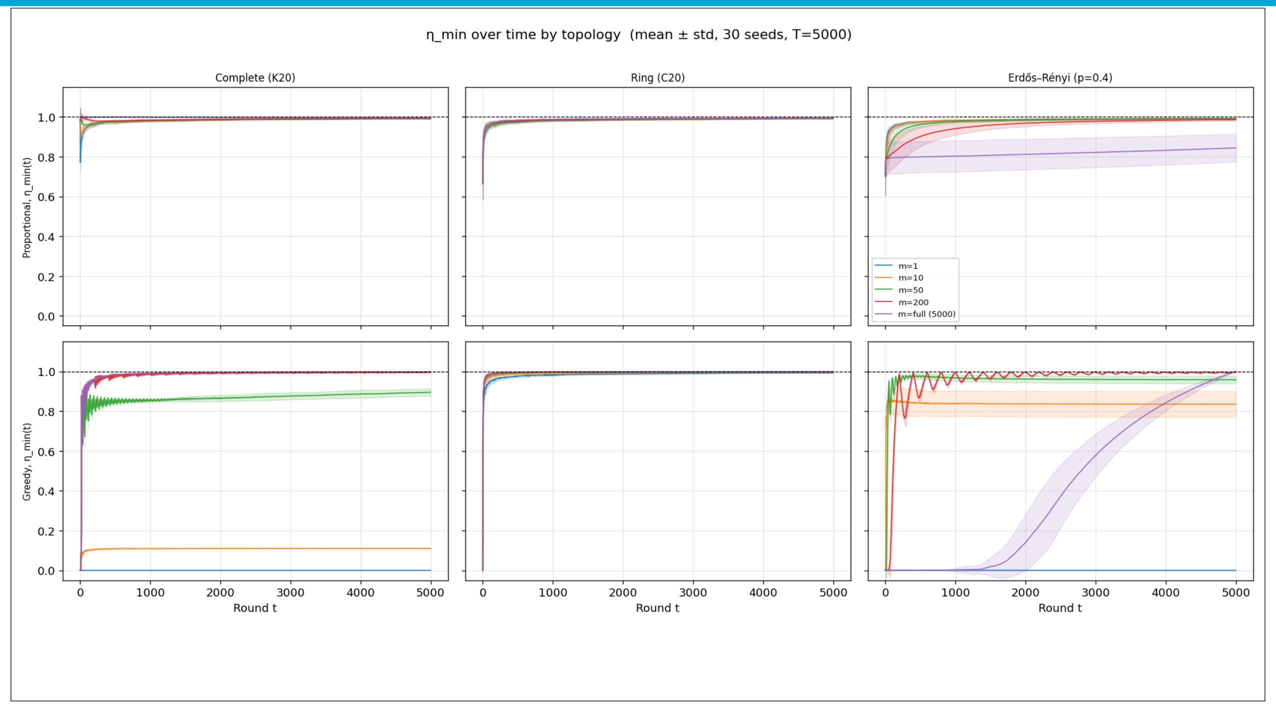

Key observations:

- Greedy's critical memory threshold scales with node degree: ring (degree 2) needs $m \geq 1$; ER (degree ≈ 8) needs $m \approx 5 - 10$; complete (degree 19) needs $m \approx 20$ — up to a 90% memory saving on sparse graphs
- Proportional stays topology-invariant for $m \leq 100$; at very large m it degrades on ER graphs ($\eta_{\min}(T) \approx 0.84$) due to structural heterogeneity that persists across random-graph draws

Take-away: Sparse networks let greedy get away with far less memory. Proportional only breaks down on irregular graphs once memory is large enough to make structural disadvantage permanent

12. Conclusions

1. Memory matters, but only for some strategies
2. Proportional remains fair across all memory sizes
3. Greedy can collapse under low memory
4. A single low-memory agent can destabilize greedy markets
5. Topology influences the convergence speed for both greedy and proportional (see thesis)

Code & Paper

GitHub: `github.com/larspaulussen/mbx`