ACM HCOMP | CogSec Workshop | September 2026

Regulating Attention in Digital Platforms via Dynamic Taxation and Strategic Learning

Presented by **Elaine Ly**

Abhishek K. Umrawal, Eliot W. Robson, Alex Zheng, **Elaine Ly** , and Lav Varshney

1. Introduction
2. Problem Formulation
3. Preliminary Experiments
4. Future Work
5. Significance for CogSec Catalog
6. Questions

How do we moderate social media platforms from the perspective of *information overload*?

The problem is two-pronged.

Users are more distracted than ever.

Algorithms and AI-enabled platforms compete for user attention at unprecedented speeds.

When the volume of incoming content exceeds users' processing capacity, it reduces user welfare, productivity, and efficiency, and forces users to prioritize content that either is trending or reinforces prior beliefs already.

Only the top players get noticed.

The lack of equity in social media platforms means that echo chambers and viral engagement gets prioritized.

Algorithms look for watch time and engagement. Because top senders already have ample reach, their content rapidly accumulates views, signals the algorithm that it is "good," and gets pushed to even more users (the "rich-get-richer" effect).

This content may be "PRIME" information: prestigious, in-group, moral and emotional information.

We develop a reinforcement learning approach that dynamically taxes senders to create fairer moderation.

We draw on recent advances in AI economic design such as the “AI Economist” paper by Zheng et al. (2021), which uses reinforcement learning to explore tax policies in simulated economies.

We explore the effect of different tax policies on social media platforms.
Our setting differs in that it centers on attention as the scarce resource and on congestion that arises from messaging.

We use a Stackelberg game as our framework.

A sequential game in economics and game theory.

One player (the leader) makes a decision first.

The other players (the followers or senders) make the best possible responses afterward while observing the leader’s choice.

Our paper makes the following contributions.

1. Dynamic modeling of the attention economy

We model the tradeoff between message volume and the distribution of users' limited attention across senders.

2. Stackelberg game formulation

We don't try to find the equilibrium right away. In reality, a social media system is complex, and static penalties may not be optimal.

Instead, senders use trial-and-error to figure out the optimal number of messages.

3. Partially Observable Markov Decision Process (POMDP)

We show that the planner's problem can be cast as a POMDP, which allows the use of reinforcement learning algorithms for policy optimization.

4. Computational experiments

We develop a simulation framework that captures key features of real-world digital ecosystems.

There are budget constraints to send messages, limited attention capacity from users, and incentives that may vary over time depending on the actions of other individuals in the framework.

Environment and agents.

There are time steps $t = 1,\ldots,T$.
At each time step, planner sets a tax policy to which senders respond.

N senders indexed by $i = 1, \ldots, n$.
Total number of messages sent at time t is denoted by

$$l_t = \sum_{i=1}^{n} m_t^{(i)},$$

where $m_t^{(i)}$ is the number of messages sent by sender i at time t.

Sender Behavior.

The expected utility of sender *i* at time *t* is given by

$$r_t^{(i)} = \min\left\{1, \frac{C_t}{l_t}\right\} u_t^{(i)} - a_t\, m_t^{(i)},$$

where $u_t^{(i)}$ represents a stochastic process that assigns value to the sender's message, and $\min\{1, C_t/l_t\}$ shows the probability that a message is processed.

Each sender seeks to maximize its long-run cumulative reward as

$$R^{(i)} = \sum_{t=1}^{T} r_t^{(i)}.$$

Planner Objective.

We quantify equity using the negative Gini coefficient,

$$r_1(\mathbf{m}_t) = -\frac{\sum_{i=1}^{n}\sum_{j=1}^{n} |m_t^{(i)} - m_t^{(j)}|}{2n\sum_{i=1}^{n} m_t^{(i)}},$$

which rewards senders for more equitable distributions.

We quantify participation through a log-sum reward for individual sender message counts.

$$r_2(\mathbf{m}_t) = \sum_{i=1}^{n} \log\left(1 + m_t^{(i)}\right).$$

Sender Behavior.

The expected utility of sender *i* at time *t* is given by

$$r_t^{(i)} = \min\left\{1, \frac{C_t}{l_t}\right\} u_t^{(i)} - a_t\, m_t^{(i)},$$

where $u_t^{(i)}$ represents a stochastic process that assigns value to the sender's message, and $\min\{1, C_t/l_t\}$ shows the probability that a message is processed.

Each sender seeks to maximize its long-run cumulative reward as

$$R^{(i)} = \sum_{t=1}^{T} r_t^{(i)}.$$

Planner Objective.

The overall reward for the planner at time *t* is given by

$$r_t = r_1(\mathbf{m}_t) + r_2(\mathbf{m}_t).$$

The planner's objective is to maximize cumulative reward over the time horizon,

$$\max_{a_1,\ldots,a_T} \sum_{t=1}^{T} r_t.$$

Reward Shaping

- Stackelberg games sometimes run into an issue where followers are stuck between two equally-good responses.
  - They can arbitrarily choose the suboptimal equilibrium.

- We explore using reward shaping to biases exploration towards the more socially optimal equilibrium.

$$r^{(j)}_{\text{shaped}} \coloneqq r_{\text{true}} + \lambda^{(j)} \cdot L,$$

Preliminary Experiments.

Each setting adds one complication.

Multi-sender means there are followers competing against one another, while multi-step adds more steps per iteration under a budget constraint.

One-step single-sender

One round. The platform sets a tax rate, one sender chooses how many messages to send (0–10), and the game ends.

One-step multi-sender

One round with several senders choosing at the same time.

Multi-step single-sender

Fifty rounds with one sender who has a fixed budget of 100 messages in total. The platform sets a tax rate each round, and the sender has to decide how to spread its budget over time.

Multi-step multi-sender

Fifty rounds with several senders choosing at the same time with fixed budgets.

Preliminary Experiments.

Each setting adds one complication.

Multi-sender means there are followers competing against one another, while multi-step adds more steps per iteration under a budget constraint.

One-step single-sender

Converges, but graphs show that equations may need more refining.

One-step multi-sender

Converges, but graphs show that equations may need more refining.

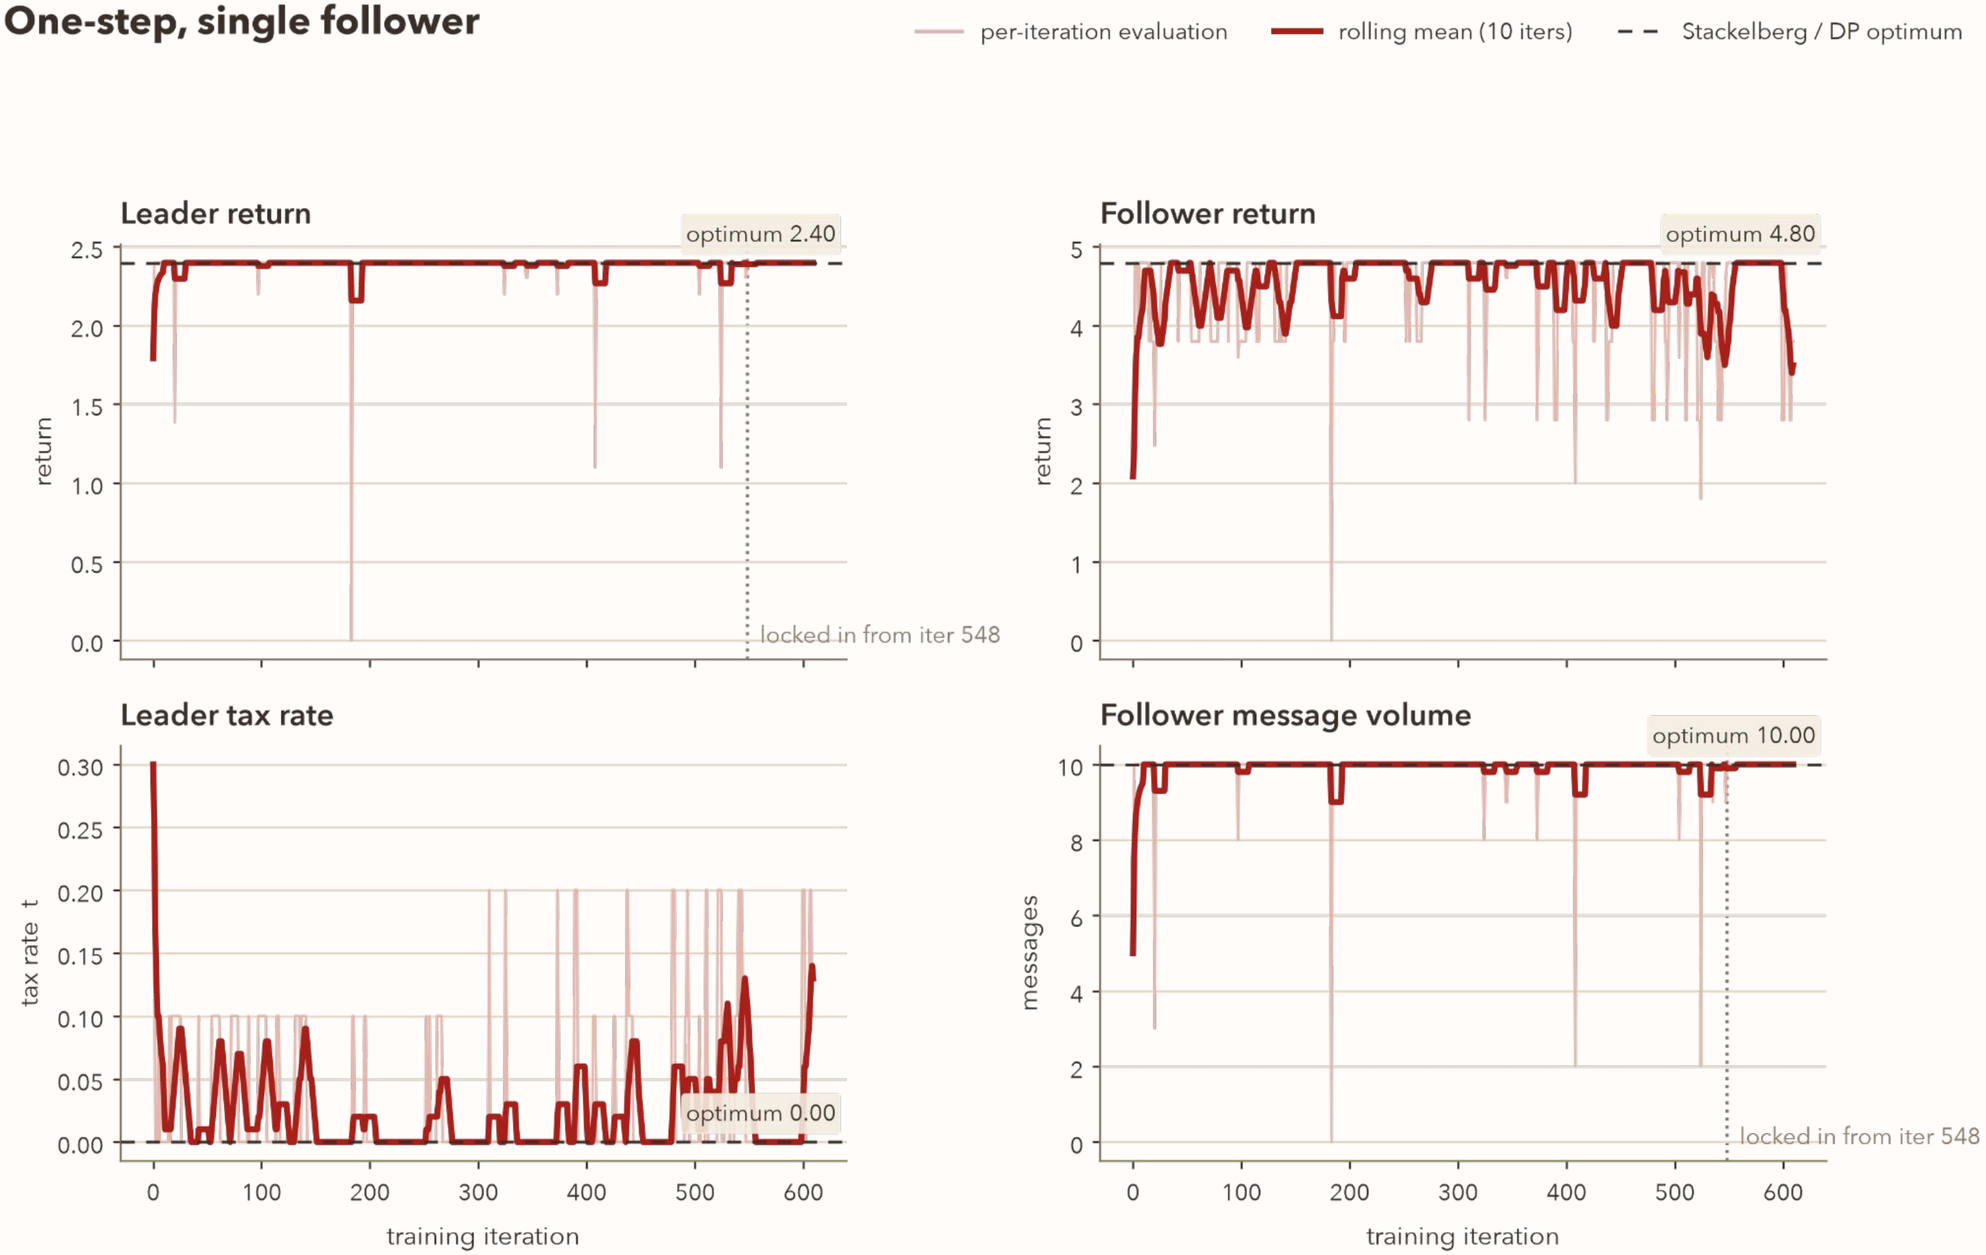
Multi-step single-sender

Converges, but graphs show that equations may need more refining.

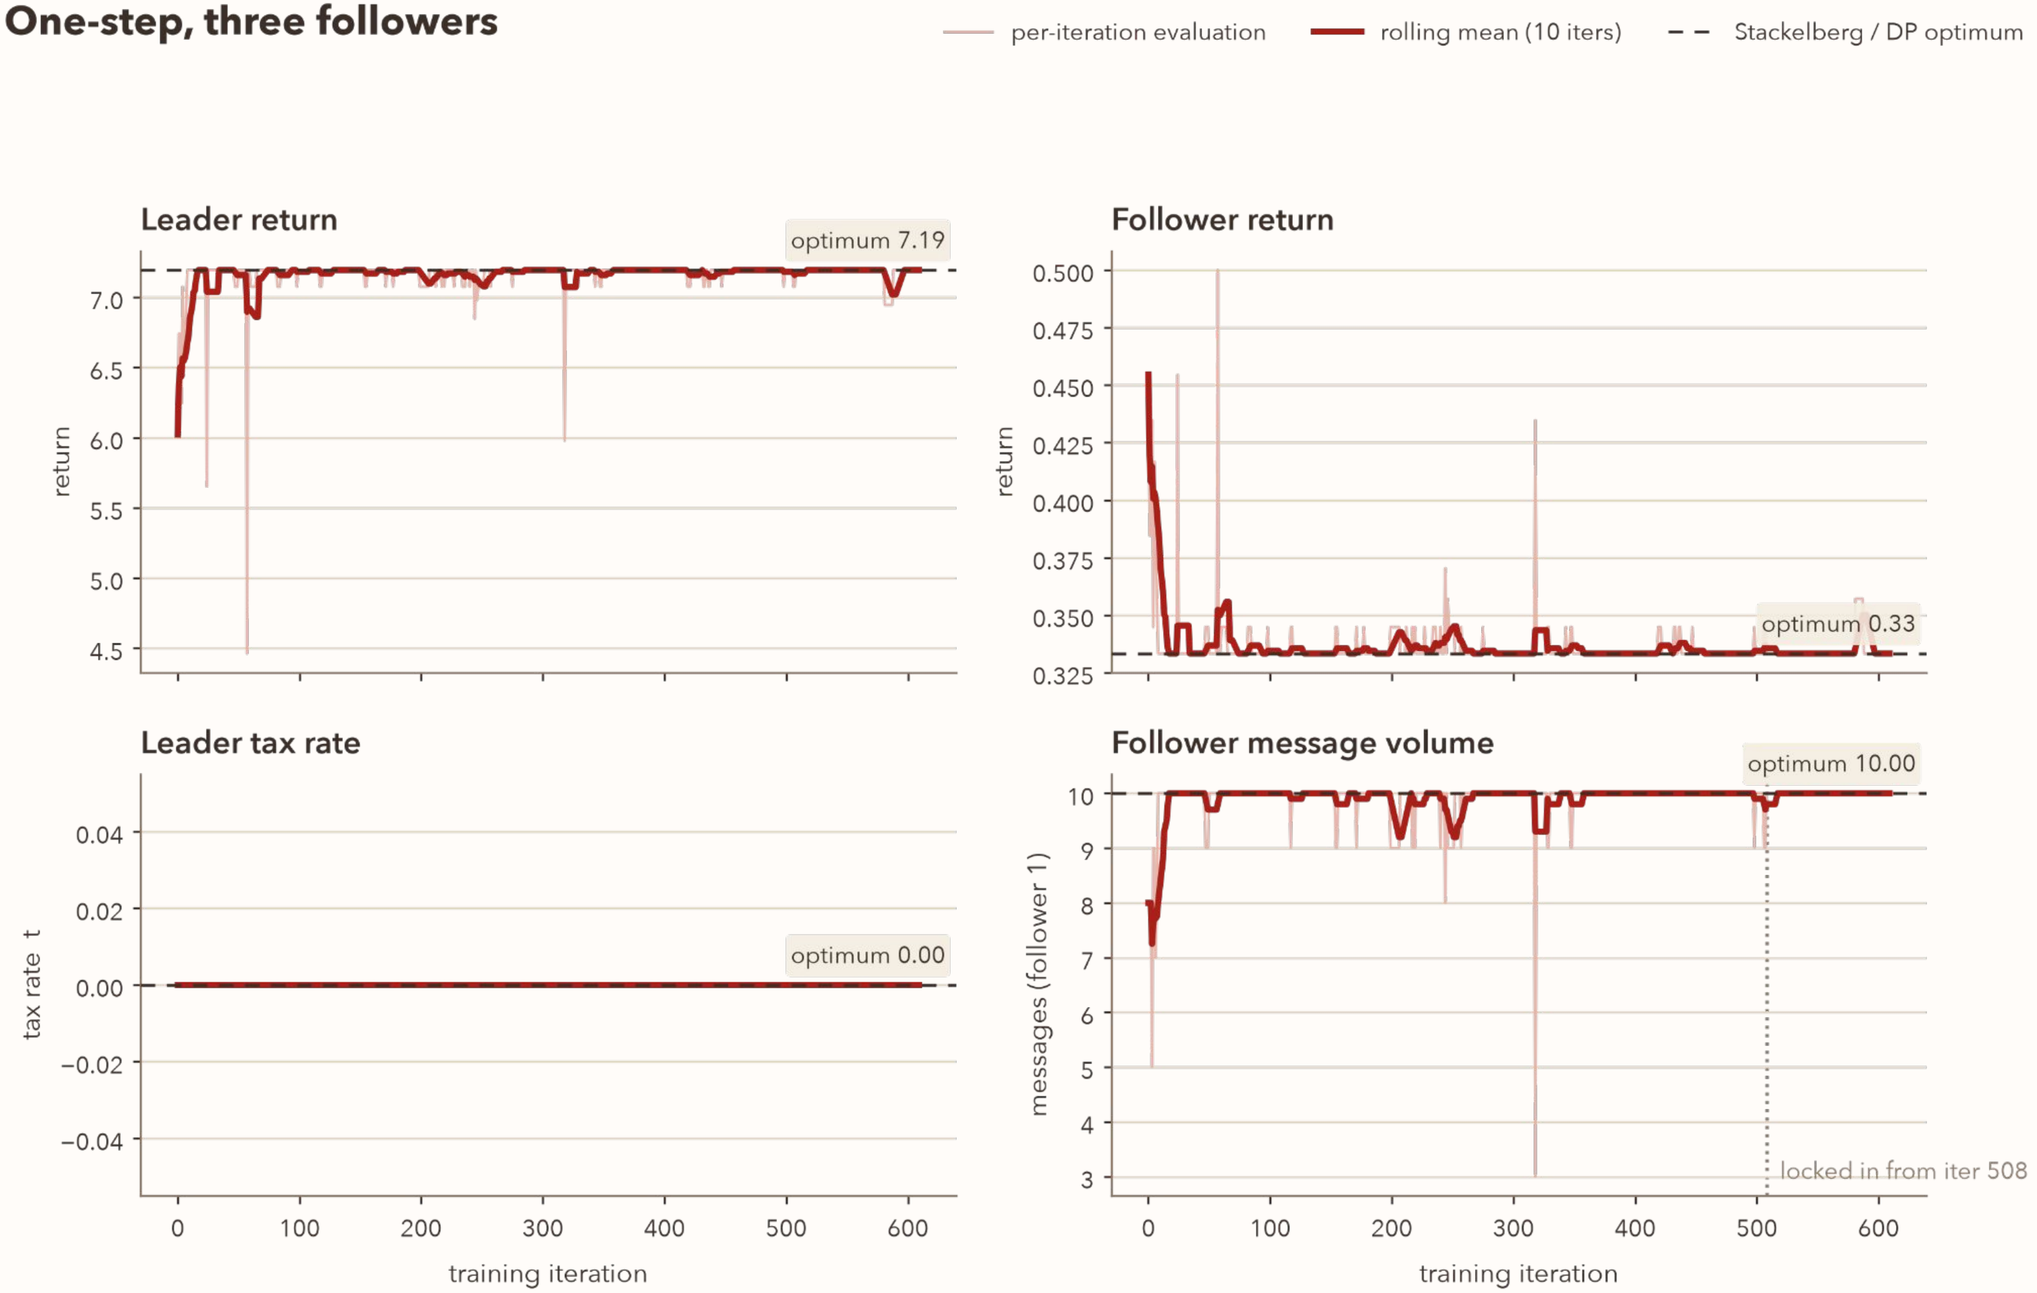
Multi-step multi-sender

Still in progress.

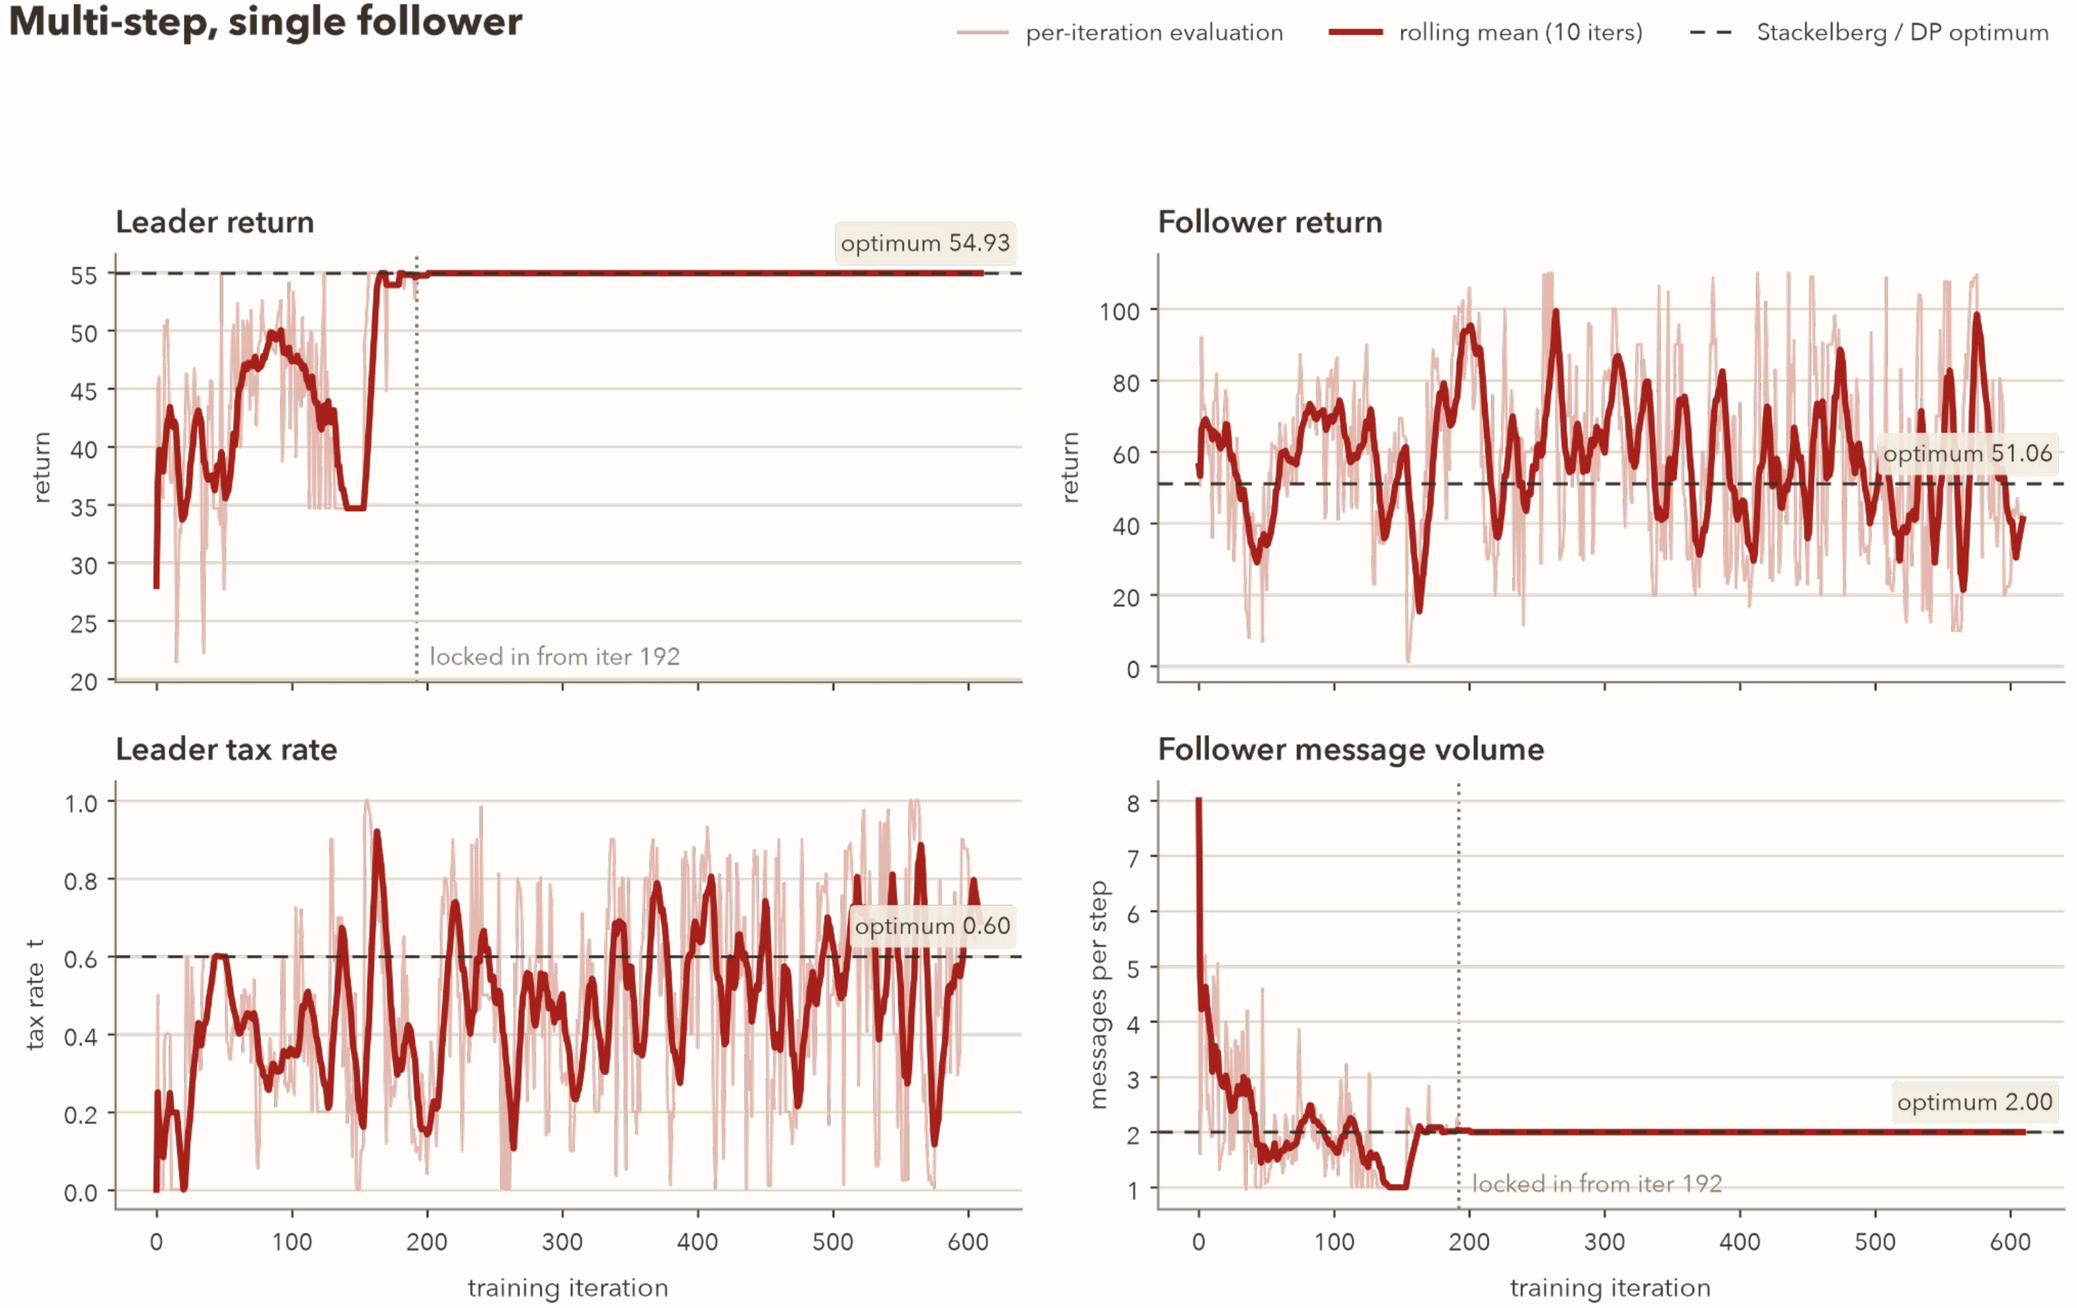
One-step, single follower



One-step, three followers



Multi-step, single follower



Plans for future work.

1. Continue experimentation.
2. Refine equations.
3. Look into real-data simulations.

Any other suggestions? (Wait for closing slide.)

We identify *information overload* as a cognitive vulnerability that inhibits people's decision-making and information-gathering skills, individually and collectively.

Thank you! Questions? Suggestions?

Please email questions or suggestions
about the paper to:
ly.el@northeastern.edu

We will work to answer your queries as
a team.