Paper Review:

The AI Economist: Improving Equality and Productivity with AI-Driven Tax Policies

Abdullah Mamun

March 8, 2021

About this paper

The AI Economist: Improving Equality and Productivity with AI-Driven Tax Policies

Stephan Zheng[*,1], Alexander Trott[*,1], Sunil Srinivasa[1], Nikhil Naik[1], Melvin Gruesbeck[1], David C. Parkes[1,2], and Richard Socher[1]

[1]Salesforce Research
[2]Harvard University

April 28, 2020

Abstract

Tackling real-world socio-economic challenges requires designing and testing economic policies. However, this is hard in practice, due to a lack of appropriate (micro-level) economic data and limited opportunity to experiment. In this work, we train social planners that discover tax policies in dynamic economies that can effectively trade-off economic equality and productivity. We propose a two-level deep reinforcement

About this paper

- arXiv preprint by a team of **Salesforce Research** and **Harvard University** in 2020
- Cited by **15** as of March 8, 2021
- Reinforcement Learning-based model to make tax policies

Goal

- Balance wealth distribution (equality) and productivity by introducing an intelligent tax policy planner.

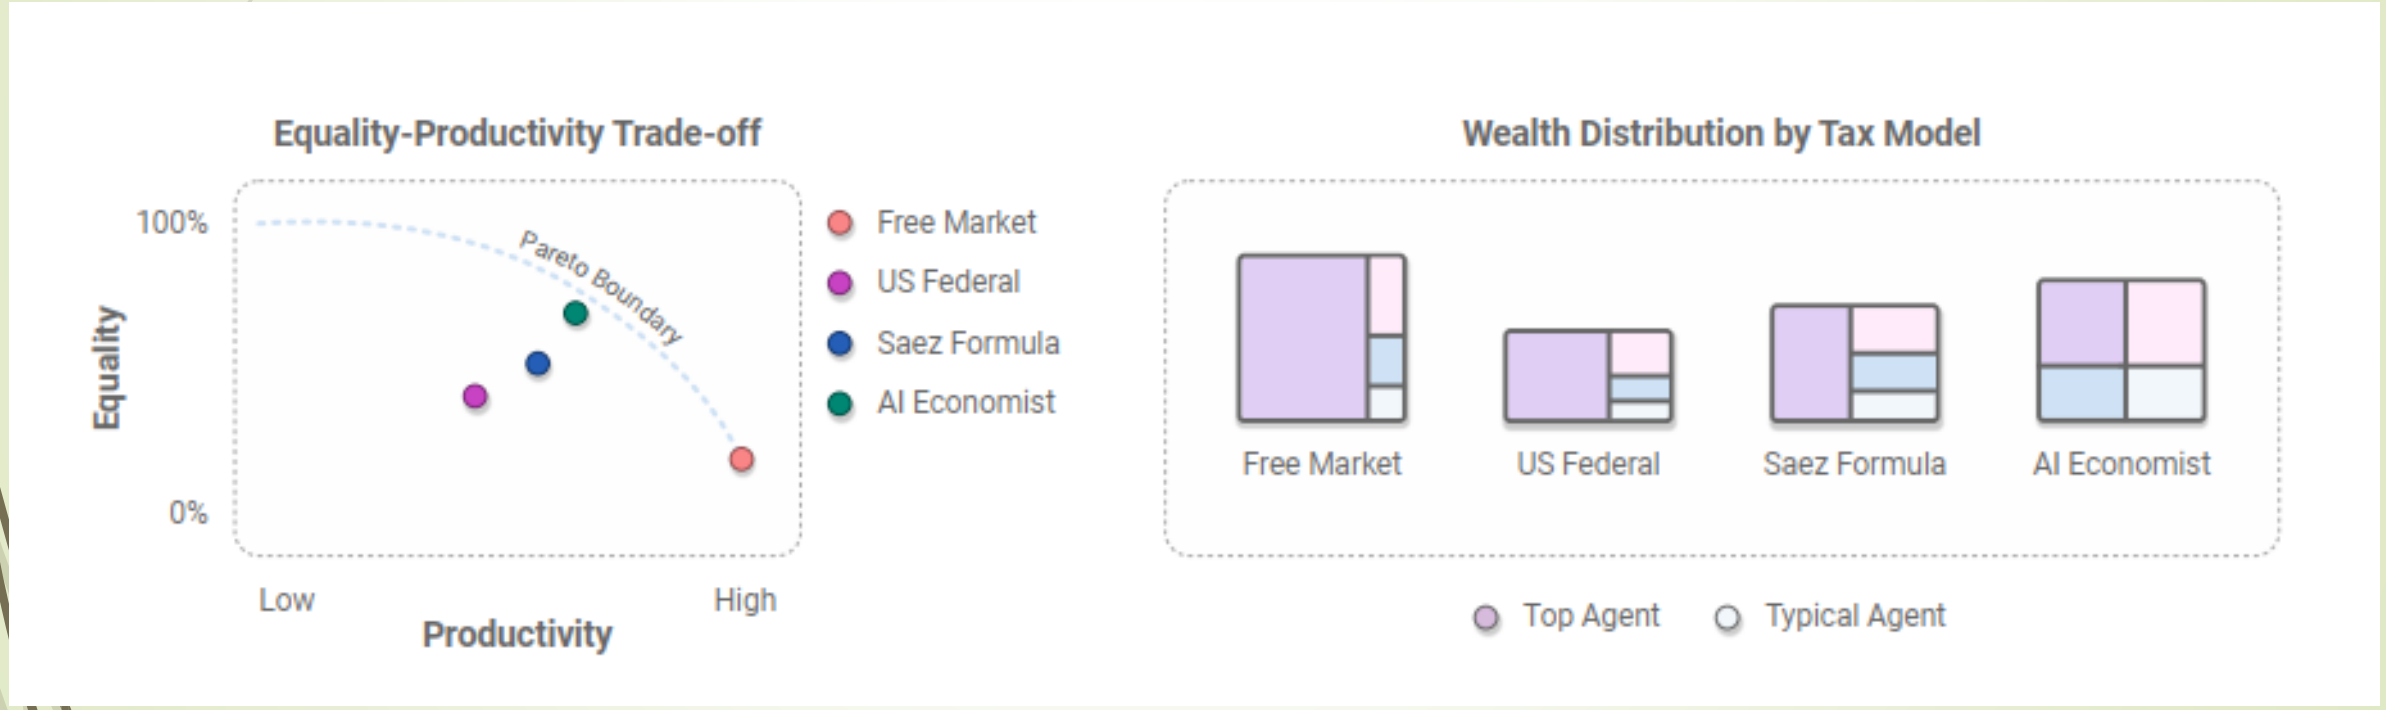

Experimental environment (Agents)

- AI agents of different levels of skills
- Agents can gather resources, trade, and build houses.
- Agents earn coins
- Policy planner imposes taxes to maximize equality and productivity
- Agents adapt their behaviors to maximize utility.

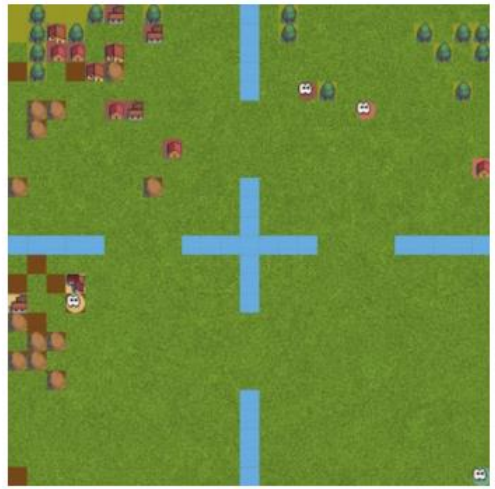

Figure 1: The Gather-and-Build game. Agents move around, collect resources (wood and stone) and build houses. Agents cannot move through each others' houses, or move through water. Agents can trade resources.

Experimental environment (human)

- Humans were rewarded in the amount of **USD bonus = Utility × 0.06**
- Moving around to gather resources were counted as labor cost
- On average $11.26 /30 mins was the reward for participants which is higher than the minimum wage.

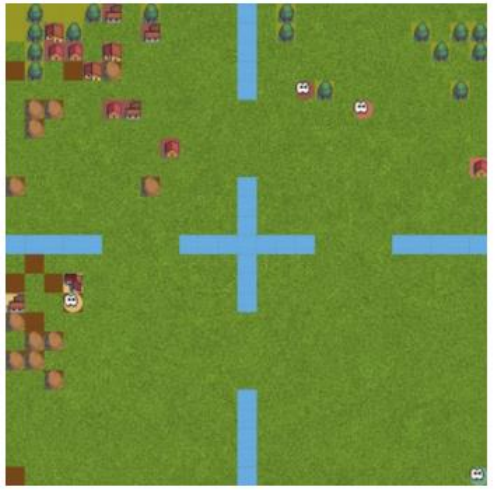

Figure 1: The Gather-and-Build game. Agents move around, collect resources (wood and stone) and build houses. Agents cannot move through each others' houses, or move through water. Agents can trade resources.

Experimental environment (human)

- Two-level reinforcement learning

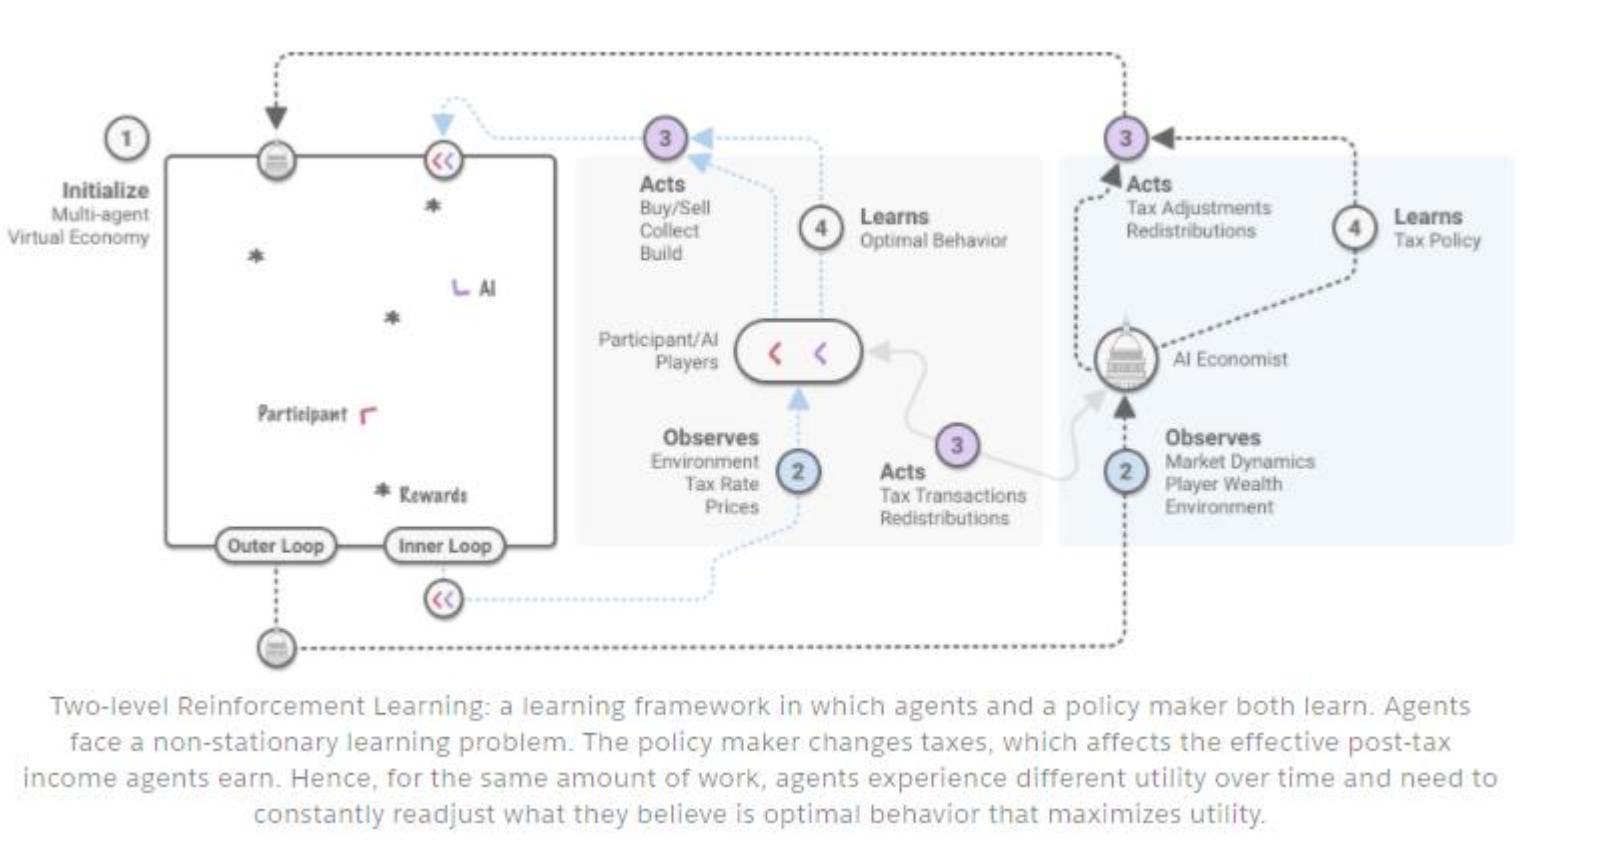

Two-level Reinforcement Learning: a learning framework in which agents and a policy maker both learn. Agents face a non-stationary learning problem. The policy maker changes taxes, which affects the effective post-tax income agents earn. Hence, for the same amount of work, agents experience different utility over time and need to constantly readjust what they believe is optimal behavior that maximizes utility.

Results

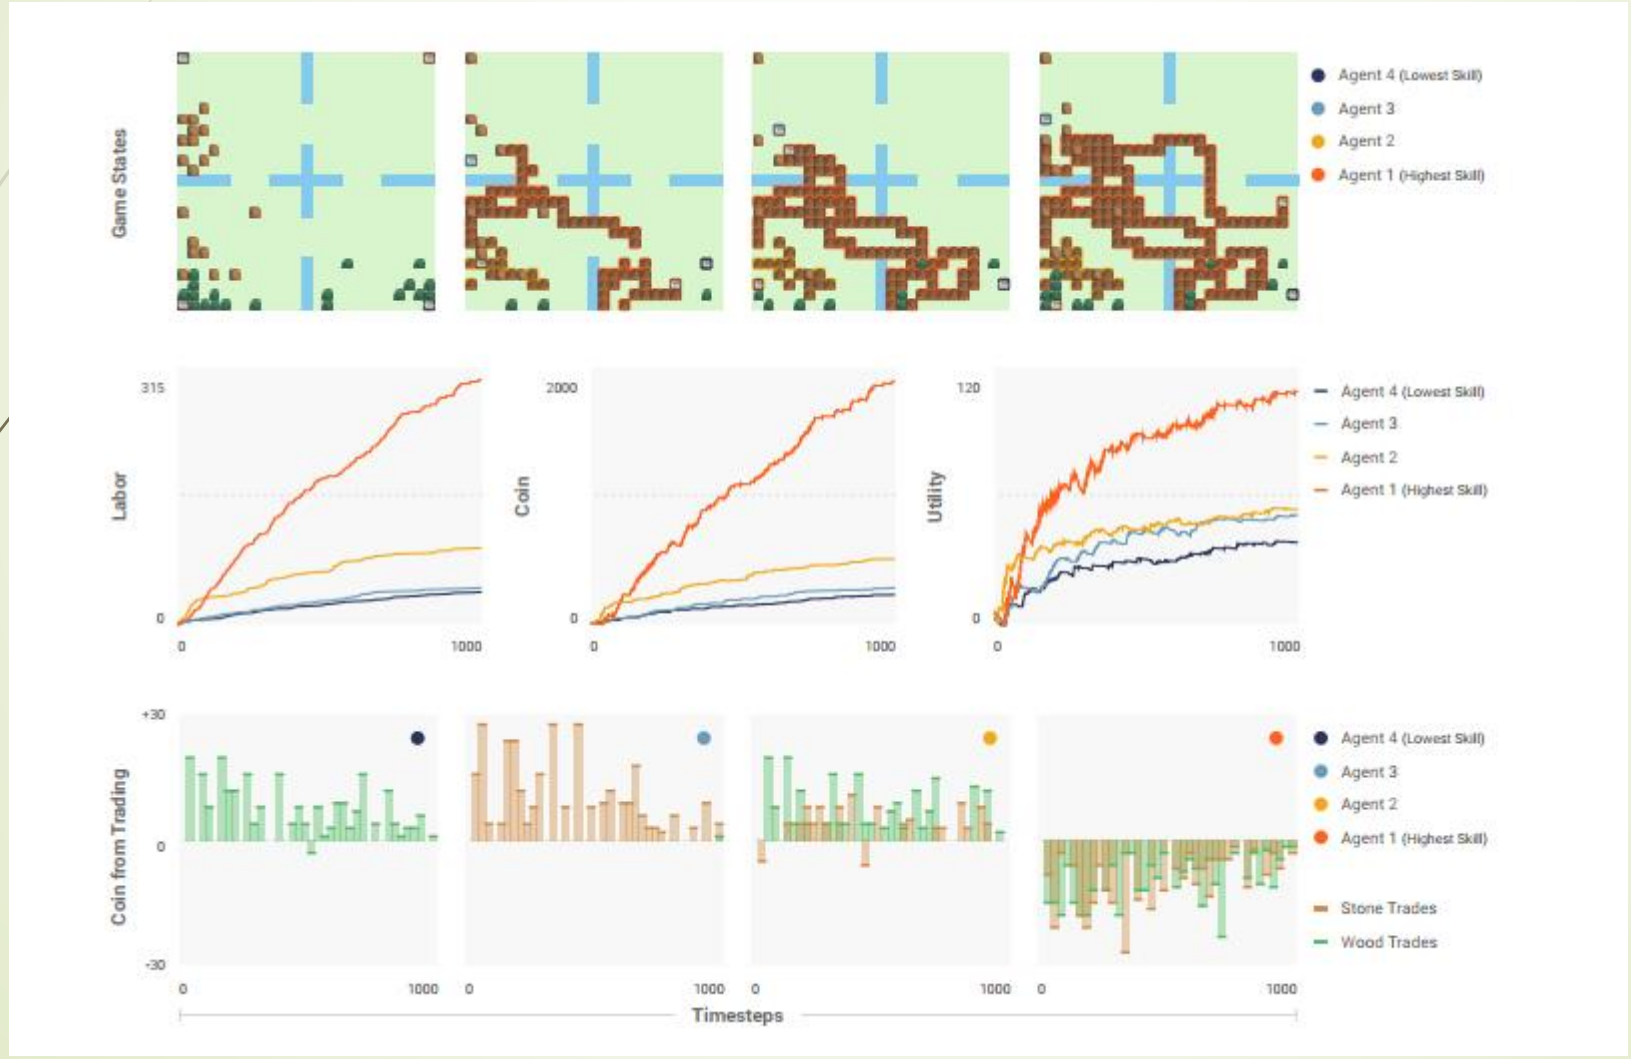
Game States
Agent 4 (Lowest Skill)
Agent 3
Agent 2
Agent 1 (Highest Skill)
Labor
315
0
0
1000
Coin
2000
0
0
1000
Utility
120
0
0
1000
Agent 4 (Lowest Skill)
Agent 3
Agent 2
Agent 1 (Highest Skill)
Coin from Trading
+30
0
-30
0
1000
Timesteps
Agent 4 (Lowest Skill)
Agent 3
Agent 2
Agent 1 (Highest Skill)
Stone Trades
Wood Trades

Results (simulation)

- Achieves better equality*productivity in the cost of imposing higher taxes on the middle-income people.

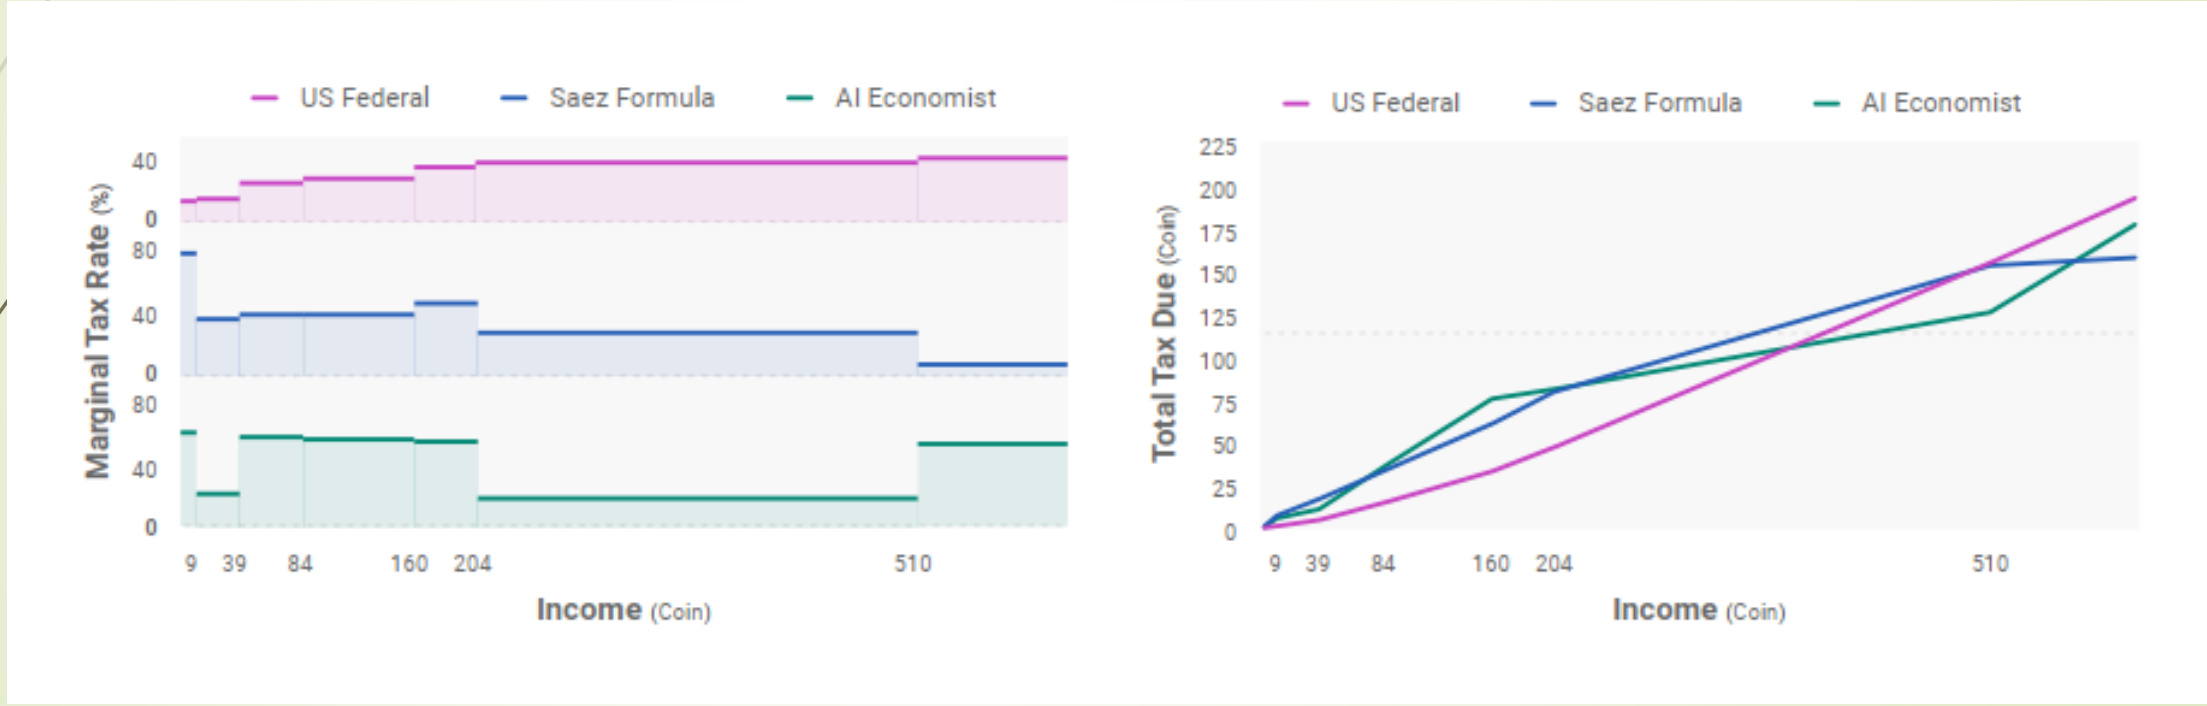

Results (human participants)

- Achieves better equality*productivity in the cost of imposing higher taxes on the middle-income people.

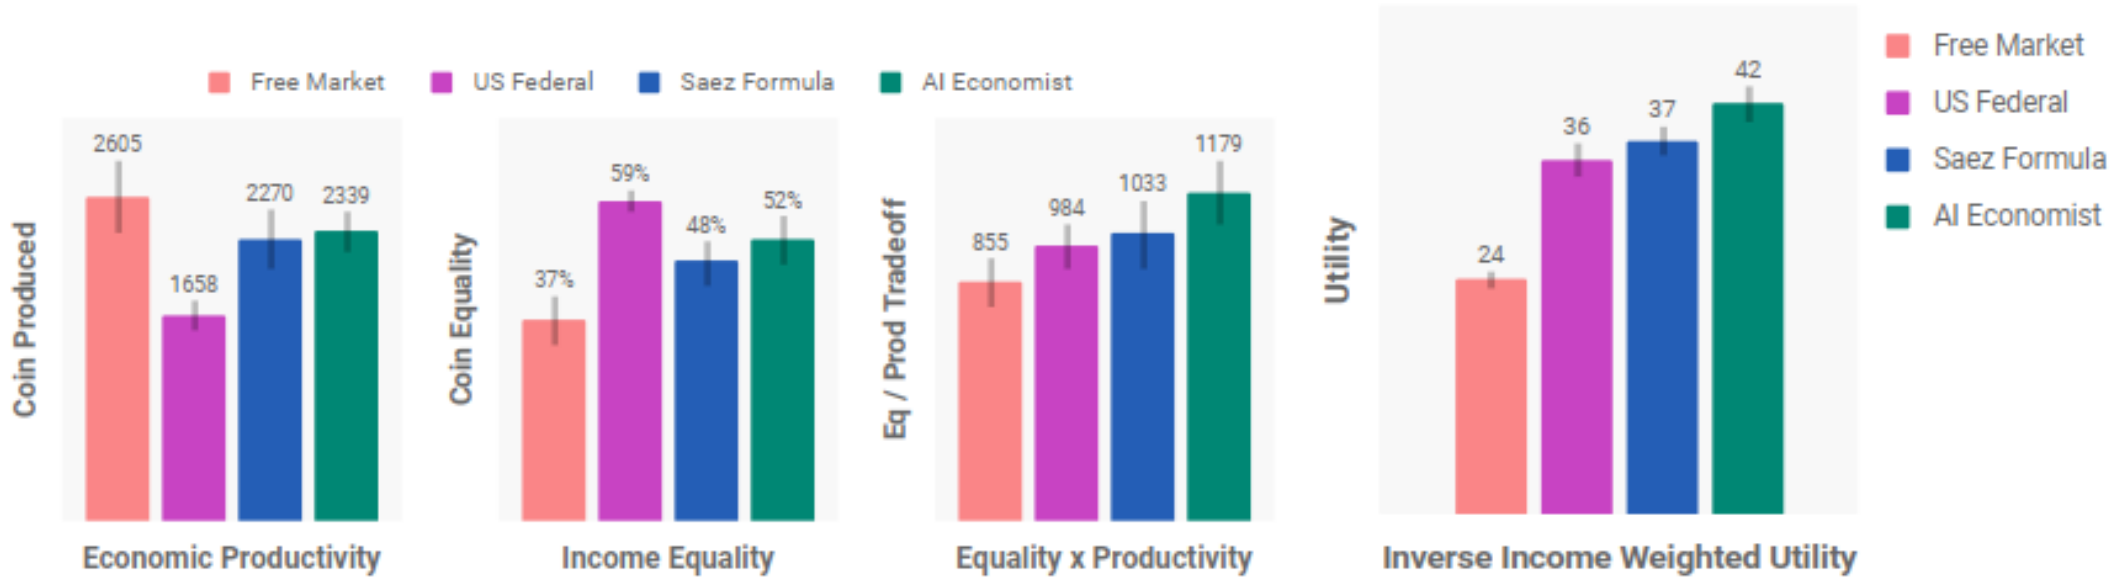

Results (human participants)

- Achieves better equality*productivity in the cost of imposing higher taxes on the middle-income people.

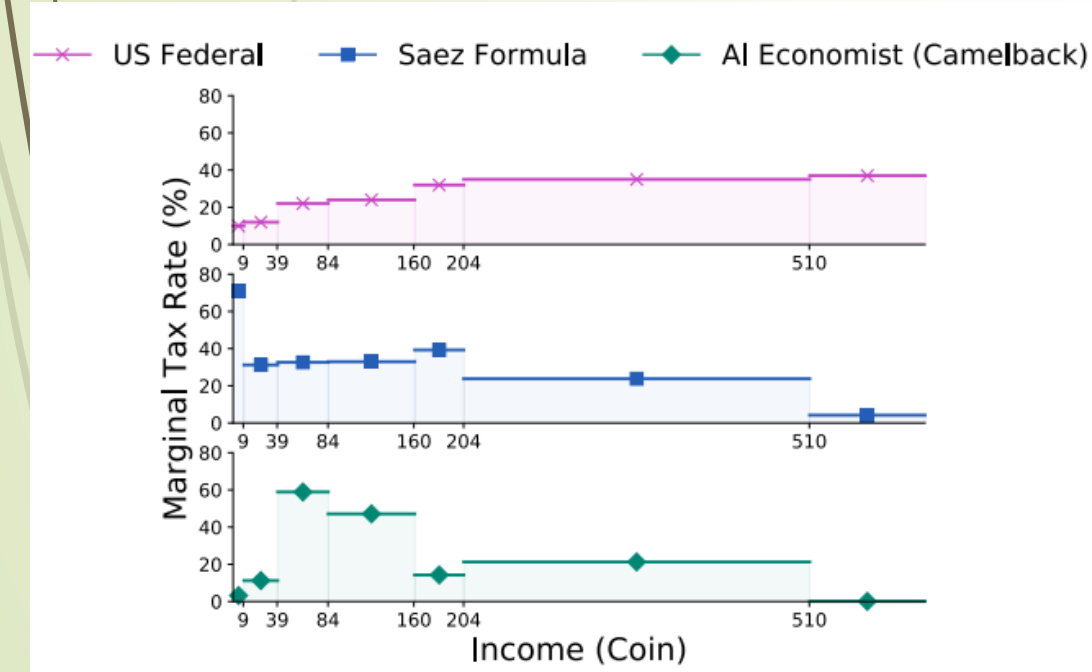

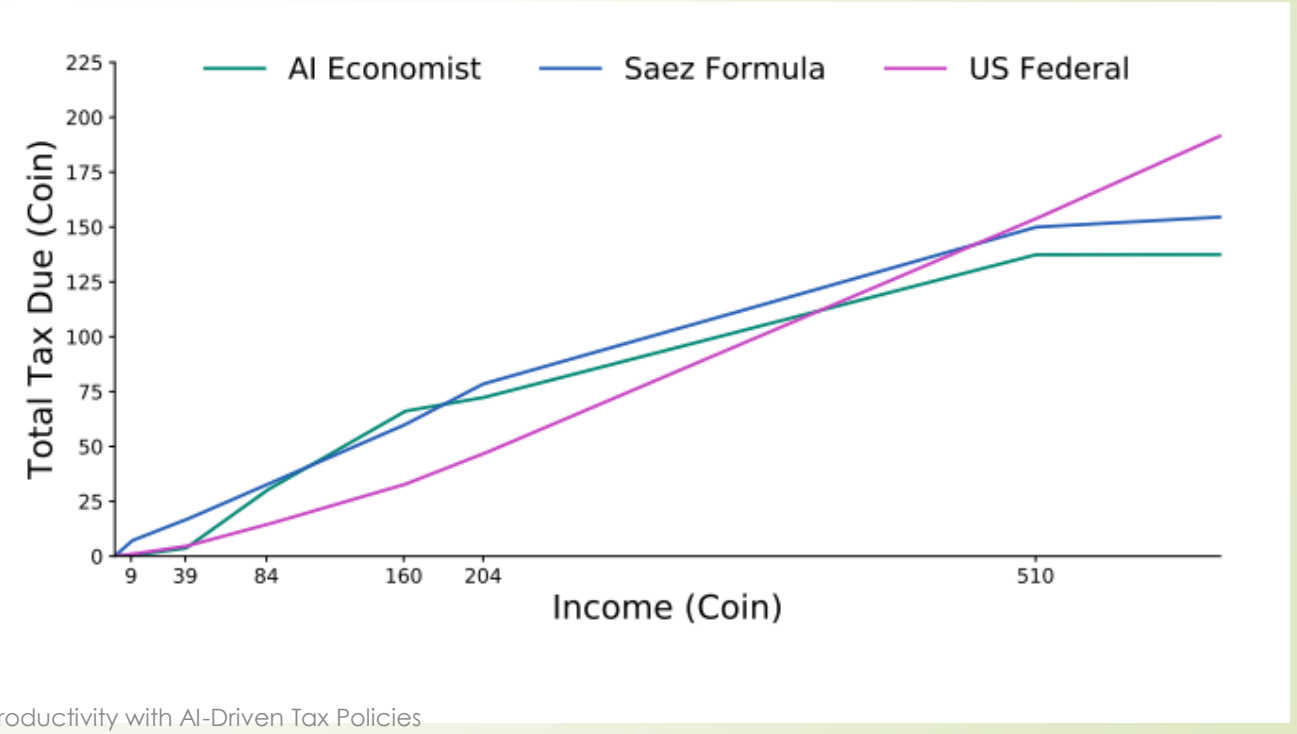

Thank You!

Contact: abdullahal.mamun1@wsu.edu
Read the paper here