LOOKAHEAD AGENT ANALYSIS IN LONG-RUN FAIRNESS SIMULATION

LOOKAHEAD AGENT ANALYSIS IN LONG-RUN FAIRNESS SIMULATION

A thesis presented by
Gayatri Balasubramanian
to
Computer Science
in partial fulfillment of the honors requirements
for the degree of Bachelor of Arts

Harvard College
Cambridge, Massachusetts
April 5, 2021

Abstract

Bank lending has been used as an exploratory scenario to understand how "fair" short term lending policies fare in the long-run. Unfortunately, policies implemented in the short-run are not necessarily fair in the long-run. In this paper, we build off previous simulation studies to propose four one step lookahead lending agents. Lookahead checks the expected outcome of a lending policy before making a lending decision. We compare these lookahead policies against each other under relative improvement and active harm scenarios after multiple time steps. We find that the advantaged group and bank always benefit most from a maximum utility agent – that maximizes bank profit – while the disadvantaged group always benefits most from an unbounded equality of opportunity agent – which maximizes bank profit under lending thresholds with equal true positive rate across the groups.

Acknowledgements

My foremost gratitude goes to Dr. Andrew Perrault, who eagerly accepted my proposal to explore research and who was always an email or Google Meet away. Andrew answered every question that I had and anticipated every problem that I would come across. He always made time to answer my questions and without him, I could never have done this. Thank you, Andrew, for your time and insight – for being my guide and filling it with fun and zero stress. Getting to know and learn from you was the best part of my senior year.

Next, Professor Jen Halen, who fed my interest in the arts and always tried without fail to make me a more informed person. Thank you, Professor Halen, for the times in your office and all the conversations we have had together. You inspire me with how infinitely you care for your students and the issues you teach, and how truly outstanding of a Professor you are.

I am also grateful to my readers Professor Milind Tambe and Professor Mike Smith. Professor Tambe, your class was one of the courses I most enjoyed, especially learning about all sorts of problems you tackle. And, Professor Smith, you always calmed my nerves and never said no. I am so lucky to have gotten you as my advisor during my time at Harvard.

Thank you, Dr. Yoni Halpern, for being a resource to me. I reached out blindly to you, and you were so eager to answer my questions and help me understand your work in `ML-fairness-gym`.

And last but most importantly, thank you to my mom for her unquestionable support and love.

Contents

1 Introduction

This paper builds on recent work from "Fairness Is Not Static: Deeper Understanding of Long Term Fairness via Simulation Studies," in which the authors D'Amour and Srinivasan et al. advocate for the use of simulation studies to better understand the long-term behavior of systems which deploy machine learning-based decision making.

Fairness metrics such as equality of opportunity are evaluated based on their ability to be applied and produce a fair outcome one timestep later. The logic is that if each decision made is fair, then the outcome several timesteps later will also be fair. However, it is not true that single-step fairness analyses necessarily hold out in the long run. Not only do simulation studies reveal this fact, but they also give greater insight into what dynamics can arise in a machine learning decision-based systems.

D'Amour and Srinivasan et al. sought to answer whether fair decisions in the short term hold out for the long run. In their follow-on section, they raised the question of designing short term policies that produce a fair decision in the long run. We present some beginning work at addressing this question built around a predictive step: proposing lending agents that use the expected outcome of their policy to arrive at a lending decision. We consider `max-util`, `GB`, `Bounded EO`, and `Unbounded EO` agents.

2 Background

D'Amour and Srinivasan's lending scenario and implementation comes from another paper, "Delayed Impact of Fair Machine Learning," by Hardt et al.. Hardt et al. define a dynamic bank lending scenario where there are two fixed, finite-sized groups. Group 0 is the disadvantaged group and Group 1 is the advantaged group. The groups are closed so no individual leaves the group and no individual enters the group. Each individual has an observable membership variable and a discrete credit score $1, 2, \cdots c_{max}$. The maximum credit score decile is $c_{max} = 7$. The groups start out with some credit score distribution, `pi_0` and `pi_1`. Applicant sampling happens with replacement from the group, and, at each time step, the bank either approves or denies the loan. If the bank approves the loan and the applicant defaults, the applicant's credit score decreases by 1 and the bank's profit decreases by r. If the applicant pays back the loan, the applicant's credit score increases by 1 and the bank's profit increases by r. The bank's decision to approve or deny the loan is based solely on the discrete credit score of the applicant. Lastly the probability any applicant repays is a function solely of fixed credit score repayment certainties, `certainty_0` and `certainty_1`.

With this model, Hardt et al. explicitly focus on threshold-policies – where the bank fixes a threshold within each group to offer loans. Any individual with credit score above the threshold receives a loan, and any individual with credit score below the threshold is denied a loan. (Again, probability of repayment is a function of credit score that is independent of group-membership.) They consider three threshold-based bank agents. First, a max-utility bank which seeks to maximize profit: this bank chooses a single threshold at or above their break-even point. Second, a "fair" bank abiding by demographic parity, which chooses a different threshold in each group so as to lend to both groups at equal rates. Lastly, a "fair" bank abiding by equality of opportunity, which chooses a different threshold in each group so as to lend to equal rates among individuals who repay their loans.

Hardt plot group mean score as a function of selection rate for each policy, and they define a policy causes "active harm" if it decreases the mean credit score of the group. "Relative harm" is if the policy causes a mean score increase less that the mean score increase of the utility-maximizing bank agent. A policy causes "relative improvement" if it increases the mean credit score more than the max-utility bank agent will. Figure 1 is taken from the Hardt paper, showing the Outcome Curve.

2.1 Related Work

Hardt's model is beneficial in that it takes into account the inherent risk that banks take, as the bank's lending risk is greater than the individual's application risk. And they formulate a "one-step feedback model" – what one step of the lending space would be and how that one step would impact score distribution immediately after. They found that just because a policy is designed for fairness/equality in the bank's decision, i.e. to lend at equal rates, that does not necessarily promote a fair outcome on the credit score distribution one step later. They find that the max-util agent does not cause active harm, fairness criteria (demographic parity and equality of opportunity) can cause relative improvement, fairness criteria can cause active harm by being over-eager (selecting at a high rate), there are credit score distributions where one fairness criteria causes active harm where the other causes improvement, inter alias.

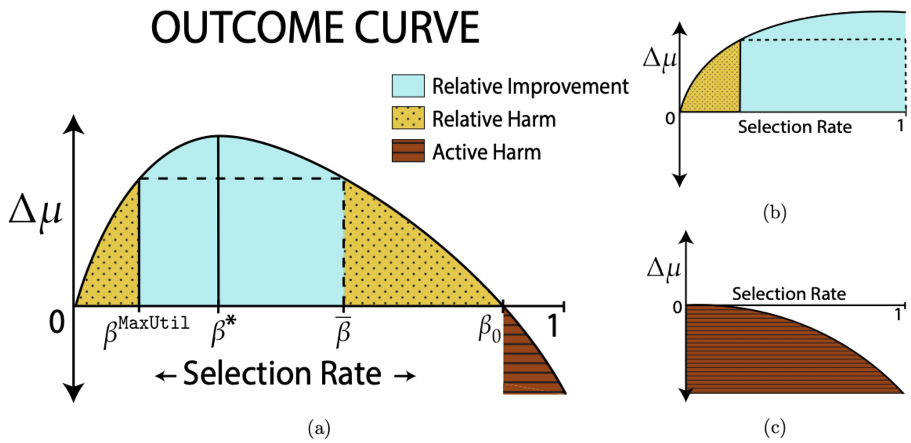

Figure 1: The above figure shows the *outcome curve*. The horizontal axis represents the selection rate for the population; the vertical axis represents the mean change in score. (a) depicts the full spectrum of outcome regimes, and colors indicate regions of active harm, relative harm, and no harm. In (b): a group that has much potential for gain, in (c): a group that has no potential for gain.

In short, Hardt et al. show that just because we take a fair step does not guarantee the impact of our decision is fair one epoch later.

D'Amour and Srinivasan et al. took Hardt et al.'s set up of the bank lending scenario and one step feedback and applt it to many epochs, creating ML-fairness-gym to run simulations. In addition to checking the mean score change in the group, they add metrics of group conditional probability of repayment, cumulative number of loans granted, and the agent's "true positive rate" (TPR – requiring the bank to lend in both groups at an equal rate among individuals who repay their loan). In their scenario, they have two agents: a max-util bank and a equality of opportunity agent. Both of these agents are threshold agents, meaning their choice to approve or deny any applicant is based on credit score and applicant group membership above or below some bin threshold cut-off. At each timestep, the agent receives a single applicant. The applicant can be from either group.

Not only did the multi-step analysis of D'Amour and Srinivasan et al. also conclude that a fair policy decision does not guarantee a desired outcome, but their results also diverged from Hardt et al.'s one step analysis in some places. For instance, while they found that equality of opportunity policies over-lend to the disadvantaged group, they found the impact of fairness criteria more ambiguous. And, most interestingly, that the average credit score of both groups drops under the max-util bank agent, and that equality of opportunity applied at each time step does not succeed at equalizing the TPR across the full simulation.

2.2 Discussion

While Hardt et al.'s lending scenario provides a canonical and simple model for analyzing the one step and multi-step feedback of fairness criteria and max-util behavior, there are several drawbacks to the model. Hardt et al. write that the outcome-based solution corresponds to giving out loans to the protected group in a way that reduces profit for the bank compared to unconstrained profit maximization, but avoids loaning to those who are unlikely to benefit, resulting in a maximally improved group average credit score. Indeed this is consistent with the findings in D'Amour and Srinivasan et al.

The imposition of fairness criteria thru a threshold classifier agent and the model's assumption that not getting a loan has no negative effect on credit score of an individual could compound to cause credit score distribution stagnation below the threshold and bimodal behavior above the threshold. In similar words, the "rich get richer" which the poor go nowhere. Also, as D'Amour and Srinivasan et al.'s results show, the focus on economic mechanism means that under fairness constraint, the mean credit score of the group can decrease. We can be fair towards all applicants and make sure the bank at least breaks even, but hurt everyone fairly. Moreover, the focus on group credit score as the metric of interest means that the side-effects of a threshold policy are masked.

We will track the Earth mover's distance on top of the group mean credit score to make better sense of the four lookahead policies that we propose.

3 Model

We model our bank lending scenario like in Hardt et al., and we run our tests with the same initial conditions as used in `ML-fairness-gym`'s `lending_params.py` and `lending_experiments_main.py` – with the exception of interest rate. D'Amour and Srinivasan use an interest rate of 1.0 while an interest rate of 2.0 is used here.

This bank model is set up that that the bank either makes `interest_rate * loan_amount` if an individual repays or loses the loan amount in full if they default. With an interest rate of 1.0, the bank was breaking even on loans, meaning if one person repaid and one person defaulted, the bank had a net of zero profit. Multi-step simulation algorithms highly dependent on randomness, i.e. frequent use of `numpy.random`, such as in determining an individual's actual success or repayment, means that when the bank is breaking even on its loan repayment with an interest rate of 1.0, even instances where the individual comes from credits score bins with high repayment certainty, `numpy.random` can still choose the individual to default. If we have enough randomly chosen defaults at the 1.0 interest rate compared to randomly chosen repayments, then the bank has a loss. Hence, working with a 2.0 interest rate gets us away from this scenario.

Just like described in the first paragraph of "Relevant Work," at each timestep, only one individual walks into the door and requests a loan – they can be from either group. They ask for the loan amount, always the same loan amount. The bank looks at their credit score and approves or denies their loan.

If the bank approves their loan and they repay, the `bank_cash` increases by the `loan_amount * interest_rate` and the credit score of the individual bumps up one bin, potentially improving the mean credit score of that individual's group. If the bank approves and the individual defaults, the bank loses all the loaned money and the individual gets bumped down one credit score bin, potentially decreasing the mean credit score of their group.

3.1 Agents

There are four agents. As per the model, the agent has full knowledge of the two group's credit score distributions and repayment certainties per bin. The agent receives an individual's credit score decile and group affiliation, and determines:

1. `max util`: a loan should be granted if expected bank profit does not decrease
2. `GB`: a loan should be granted if expected bank cash does not dip below bank minimum and group's expected credit score does not decrease
3. `equality of opportunity with lower limit`: a lending threshold in the disadvantaged group at least at a 50% repayment certainty and the corresponding threshold in the advantaged group to equalize TPR and maximum bank's expected profit
4. `equality of opportunity no limit`: a lending threshold in the disadvantaged group and the corresponding threshold in the advantaged group to equalize TPR and maximum bank's expected profit

Note that just as in "Fairness Is Not Static: Deeper Understanding of Long Term Fairness via Simulation Studies," because of the discrete scores in this setting, we cannot equalize TPR with a deterministic decision policy. We generate the ROC curves for both groups and determine continuous threshold values across both groups.

3.2 Metrics

We track the following metrics:

1. bank profit
2. Earth mover's distance between the final distribution of the advantaged group and the final distribution of the disadvantaged group: checking for how the distance changes from the initial distributions' distance
3. Earth mover's distance between the group's final distribution and its initial distribution: to see how far the group has moved from where it began – beyond just looking at the group's mean score change, for both disadvantaged and advantaged groups separately
4. Difference between average credit score of the group's final and initial distributions: to check whether the overall credit score outcome of the disadvantaged group has improved; and in the advantaged group to check whether the advantaged group experiences leveling down
5. Rate of successful loan (total repayments / total loans granted): to tell how good the agent is at determining who to loan to, for both disadvantaged and advantaged groups separately
6. Loan rate (total loans granted / total loans applicants): to tell how generous the agent is at giving out loans, for both disadvantaged and advantaged groups separately

Each timestep represents one individual appling for a loan. We graph these metrics in two cases: relative improvement and active harm, at timesteps 25, 50, 100, 250, 500, 1000, 2000.

3.3 Algorithm

The algorithm has the following parts:

- `iterate` method: calls a specific agent's `one_step` method the specified number of iterations, in other words run the simulation for the desired number of timesteps (where 1 timestep represents one person) for a particular agent
- `one_step` methods: calls `actual_update` method and `expected_update` agent methods; loans based on agent's decisions and updates distributions with actual outcome is loan was granted
- `actual_update` method: calls `get_person` method and determines the altered outcomes
- `get_person` method: returns which group and credit score grouping an individual is selected (with replacement) from; and whether loan was repaid or not

To run the program, see the Appendix.

3.4 Methodology

Each timestep represents one individual applying for a loan. We graph these in two cases: relative improvement and active harm.

We seed `numpy.random` once and run each agent 30 times each at timesteps 25, 50, 100, 250, 500, 1000, 2000. The metrics are averaged over the 30 repeated trials at each timestep.

4 Results

Our results are presented on pages 11 – 17. Below we discuss the long-run simulation results for each of the agents: `max util`, `GB`, `Bounded EO`, and `Unbounded EO`.

Relative Improvement. The bank makes the most profit under a `max util` policy. EO policies follow close behind with `Unbounded EO` slightly more profitable than an EO agent that does not loan below the disadvantaged group's 50% repayment certainty threshold. `GB` falls behind all other policies.

Regarding relative improvement, which is a positive change in mean credit score, we see that `Unbounded EO` produces the largest change in the disadvantaged group's mean credit score. `Bounded EO` and `max` util fallow close behind, while GB trails. This behaviour is confirmed by the Earth Mover's distance between the disadvantaged groups final and initial distributions, which show EO and max util policies change the disadvantaged group's resulting credit score distribution, with GB training behind.

On the other hand, the advantaged group's mean credit score changes most under max util, with EO policies behaving similarly, and GB trailing behind. The advantaged group's Earth Mover's distance between its final and initial distributions looks almost exactly like the graph of its mean credit score change.

Interestingly, the Earth Mover's distance between the advantaged group's final distribution and the disadvantaged group's final distribution shows that after some time steps of the policy acting, EO policies actually bring the distributions closer than they started, with Unbounded EO having a much strong effect bringing the distributions closer than Bounded EO. The GB policy keeps the distributions about as far apart as they started, and max util actually pulls them further apart than they started. This might show that a max util policy has some effect that make the "rich get richer," while EO policies "level the playing field."

Overall, however, the disadvantaged group's mean credit scores changes about the same as the advantaged group's credit score, while EO and max util policies move the disadvantaged group further than the advantaged group.

With respect the ability for the agent to discern well who will repay or not, given by the successful repayment rate metric (rate of successful loan), we observe that in the disadvantaged and advantaged group EO and max util policies are equally good at figuring out who will repay, and much better than the respective GB policy outcomes. However, interestingly, EO and max util policies are better at doing so in the disadvantaged group over the advantaged group.

Lastly, we look at how generous each of the policies are in granting loans, the loan rate. We find that EO and max util policies have similar loan rates in the disadvantaged group. In the advantaged group, max util loans more generously than EO policies. In both groups, GB is the stingiest loaner. Overall, loans are offered at a higher rate to the disadvantaged group than to the advantaged group. A GB policy is never good to follow, likely because it is too restrictive in terms of giving loans

Active Harm. Under active harm, we find that max util increases bank profit most, EO policies behave similarly with Unbounded EO performing slightly better than Bounded EO, and GB trails behind.

Next, we look at how the means and credit distributions are changing over time for the disadvantaged group. The Unbounded EO policy prevents active harm the best, with max util trailing closely behind. Bounded EO performs worse than max util, and GB harms the disadvantaged group the most by bringing down their mean credit score the most. We can confirm these changes by looking at the Earth mover's distance between the final and initial credit score distributions for the disadvantaged group. Indeed, we see that under GB, the final distribution for the disadvantaged group is the furthest from where it started. Grouped closer are the EO and max util policies with Bounded EO moving the distribution further than max util, and Unbounded EO moving the distribution least.

For the advantaged group's final to starting distribution, we see that max util is still able to prevent active harm entirely. The EO policies cause slight harm over time, and GB causes harm the entire time. We confirm this with the Earth mover's distance, which shows max util allowing the advantaged group's distribution to change the most from where it stared. EO and GB policies cause the distribution to change, but not much, from where disadvantaged group started. It is interesting that for the advantaged group max util moves the distribution about as much as EO policies, but brings down the mean credit score much further down in the same distance of change.

Comparing across the two groups, the advantaged group wants to pursue a max util policy while the disadvantaged group wants to pursue the Unbounded EO or max util policy. This is a slight discrepancy because depending on what policy is pursied, the other group can experience relative/active harm.

Next, we compare how well the agents discern who to lend to, the successful loan rate (rate of successful loan). Under active harm, max util and EO policies perform about the same for the disadvantaged group while GB gets a smaller fraction of successful repayments. For the advantaged group, GB discerns the best in the long run, with EO policies behaving similarly and Bounded EO performing slightly better than `Unbounded EO`, and `max util` discerning which individuals to lend to the worst. Across the two groups, `max util` discerns better for the disadvantaged group than the advantaged group, while `GB` discerns better for the advantaged group than the disadvantaged group. The `EO` policies perform similarly across the two groups.

Lastly, we look at how generously the different policies grant loans in the different groups. In the disadvantaged group, `EO` and `max util` policies grant loans at about the same rate, with `GB` lagging slightly behind. In the advantaged group, `max util` grants loans more generously than the `EO` policies, with `GB` the most stingy. Across the two groups, all the agents loan to the advantaged group significantly more than they do the disadvantaged group.

4.1 Analysis

Regardless of what the credit score distributions of the groups are, the bank always wants to pursue a `max util` policy. This is the outcome which allows the bank to make the most profit. This makes sense because the `max util` agent lends only if it expects to not lose profit.

If the bank must pursue a fairness policy, it should pursue `Unbounded EO`. Both of the `EO` policies optimize by choosing the threshold that maximize bank profit while having the an equal TPR across groups. The reason `Unbounded EO` performs better than `Bounded EO` is likely that it is able to have a lower threshold and loan to more individuals in the advantaged group that it would under `Bounded EO`.

Regardless of harm, the disadvantaged group always wants the bank to pursue an `Unbounded EO` policy. This policy harms the group's mean credit score least under active harm and improves their mean credit score most under relative improvement.

On the other hand, the advantaged group always wants the bank to pursue a `max util` policy. This policy increases their mean credit score most under both active harm and relative improvement.

The advantaged group would not want to pursue an `EO` policy. So the best compromise between the two groups is for the bank to pursue a `max util` policy. While this does not improve the disadvantaged group's mean credit score most under relative improvement, the `max util policy` is roughly as effective as a `Bounded EO` policy. Under active harm, the disadvantaged group would rather pursue a `max util` policy than a `Bounded EO` policy, if `Unbounded EO` is out of the picture.

In general `Unbounded EO` does performs slightly better at all the metrics than `Bounded EO`. It has a lower threshold in the disadvantaged group, and thus a lower threshold in the advantaged group, allowing the agent to lend to more individuals in the advantaged group than it would when the disadvantaged group lower bound is restricted to a 50% repayment certainty. The exception to this is that `Bounded EO` is very slightly (yet within the standard error) better than `Unbounded EO` at discerning who to lend to, the successful loan rate (rate of successful loan). This could be the case because the disadvantaged group threshold is lower, the the an individual can be chosen from a bin with lower repayment certainty in the disadvantaged group, defaulting.

4.2 Contributions

While we cannot directly compare our results with D'Amour and Srinivasan, we can draw analogious that may be helpful. For D'Amour and Srinivasan, the group conditional probability or repayment was lower for `EO` than `max util`, and it decreased steadily over time. Adding the lookahead, our `EO` has a higher probability of repayment than our `max util` in both groups under both active harm and relative improvemnt, and this probability increased over time.

Secondly, D'Amour and Srinivasan find that `EO` widens the credit-gap between groups as compared to `max util`, we find that under relative improvement, our `EO` agents actually narrows the credit-gap between the groups compared to `max util` and their starting distances, while under active harm `EO` agents narrow the credit-gap between the groups compared to `max util`. In fact, under active harm the disadvantaged group's average credit score decreases less under `EO` than under `max util` while both lending more and lending more successfully, suggesting that under `Unbounded EO` the disadvantaged group may actually be better off.

4.3 Future Work

A problem mentioned earlier with this algorithm and model is high reliance on `numpy.random`, meaning that different seeds caused different outcomes and high variance among the trials. To address this and get cleaner results, it would help to average over many different seeds in order to better understand the performance of each of these agents.

We would also plot different metrics such as aggregate TPR over time to see if a lookahead may perform better at equalizing TPR in the long run than policies that just apply `EO` at every step without lookahead.

Lastly, we seek to compare non-lookahead agents to lookahead agents diretly to see if checking the expectation of a loan decision at each step does indeed create better outcomes than just imposing the lending constraint at each step.

5 Conclusion

In this paper, we imposed lookahead on `max util`, `GB`, `Bounded EO`, and `Unbounded EO` agents – these calculate expected outcomes of loaning under that policy in order to decide loan decisions for individual applicants. We found that our `Unbounded EO` agent consistently produces better outcomes than our `max util` agent for the disadvantaged group under both relative improvement and active harm. For the advantaged group, `max util` was always the best performing policy. For the bank, `max util` was always the best policy.

Under `EO` policies in relative improvement, especially an `Unbounded EO`, the credit-gap between the advantaged and disadvantaged groups can actually decrease.

Given that even under lookahead agents, the bank and advantaged group seek a `max util`, and the disadvantaged group does not lose much with a `max util` policy over an `Unbounded EO` policy – except from the credit-gap widening under active harm – it appears that a `max util` policy may be the best policy for both groups and the bank.

6 References

[1] Liu, L., Dean, S., Rolf, E., Simchowitz, M., and Hardt, M. Delayed impact of fair machine learning. In International Conference on Machine Learning, pp. 3156–3164, 2018.

[2] D'Amour, A., Srinivasan, H., Atwood, J., Baljekar, P., Scul-ley, D., and Halpern, Y. Fairness is not static: deeperunderstanding of long term fairness via simulation stud-ies. InProceedings of the 2020 Conference on Fairness,Accountability, and Transparency, pp. 525–534, 2020.

7 Appendix

To replicate tests and run the class, please follow these directions:

1. Clone this repo: https://github.com/gayatri-a-b/thesis_code.
2. The `README.md` contain instructions on how to use the file and the parameters used in the paper.
3. Tests are commented at the bottom of the file.
4. Run `lending_environment.py` with the command `python lending_environment.py`.

You will need to define the Initial Conditions:

1. `pi_0`: distribution of credit scores for disadvantaged group
2. `pi_1`: distribution of credit scores for the advantaged group
3. `certainty_0` =`certainty_1`: probabilities of repaying for each of the credit scores bins, same across both groups

4. `group_chance`: probability of either group being selected
5. `loan_amount`: size of the any loan request made
6. `interest_rate`: interest rate on any loan
7. `bank_cash`: starting amount of money in the bank
8. `bank_minimum`: minimum amount of money bank should have

Continue onto the next pages to see the metrics' graphs discussed in "Results."

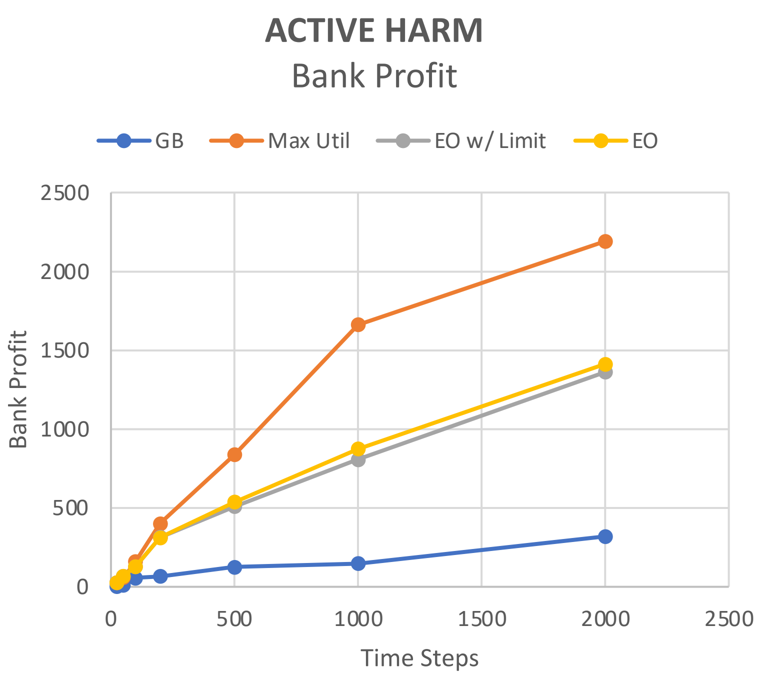
ACTIVE HARM
Bank Profit
GB
Max Util
EO w/ Limit
EO
2500
2000
1500
1000
500
0
Bank Profit
0
500
1000
1500
2000
2500
Time Steps


RELATIVE IMPROVEMENT
Bank Profit
GB
Max Util
EO w/ Limit
EO
7000
6000
5000
4000
3000
2000
1000
0
Bank Profit
0
500
1000
1500
2000
2500
Time Steps

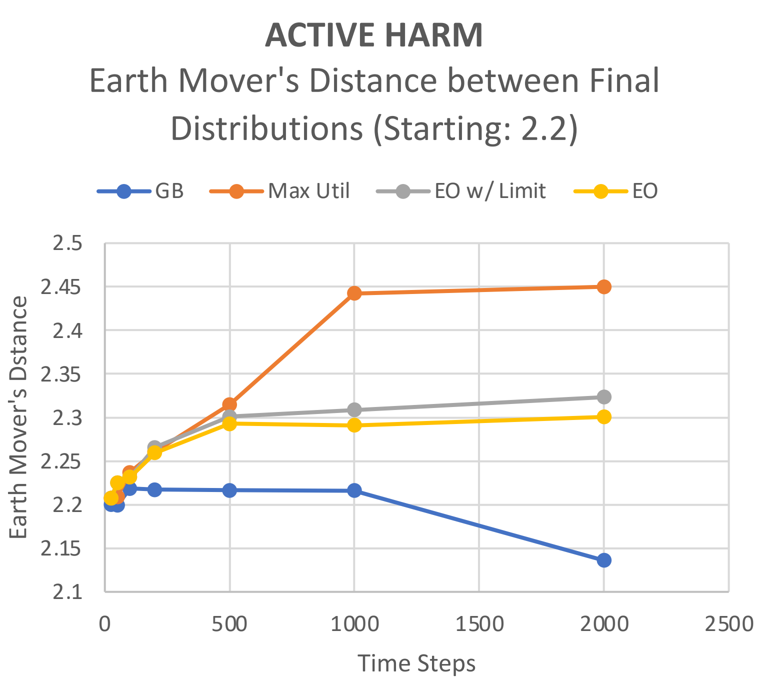
ACTIVE HARM
Earth Mover's Distance between Final Distributions (Starting: 2.2)
GB
Max Util
EO w/ Limit
EO
2.5
2.45
2.4
2.35
2.3
2.25
2.2
2.15
2.1
Earth Mover's Dstance
0
500
1000
1500
2000
2500
Time Steps

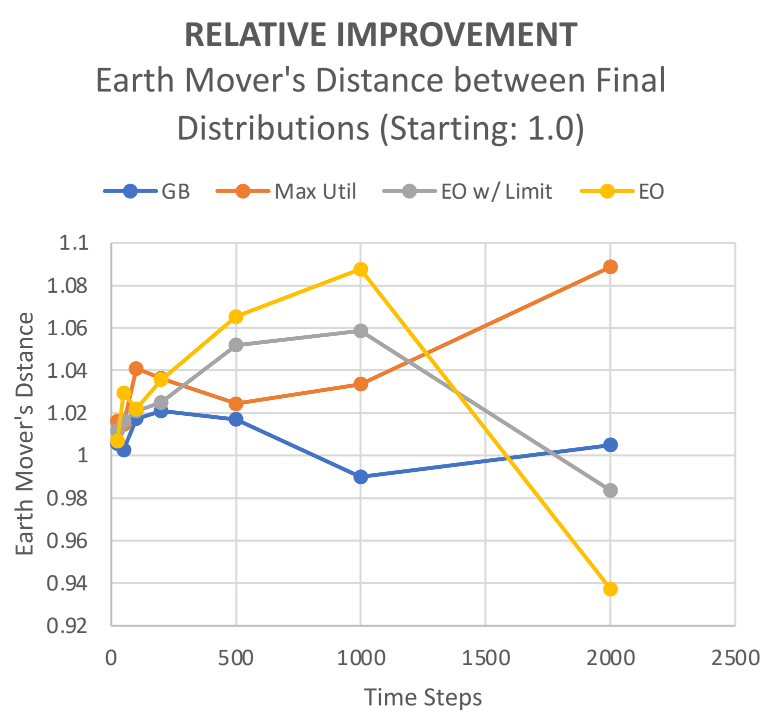
RELATIVE IMPROVEMENT
Earth Mover's Distance between Final Distributions (Starting: 1.0)
GB
Max Util
EO w/ Limit
EO
1.1
1.08
1.06
1.04
1.02
1
0.98
0.96
0.94
0.92
Earth Mover's Dstance
0
500
1000
1500
2000
2500
Time Steps

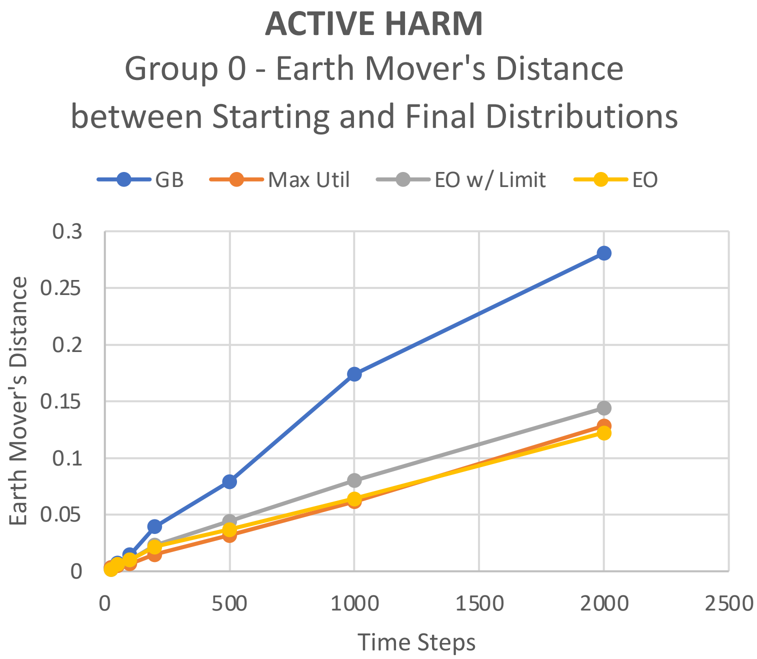
ACTIVE HARM
Group 0 - Earth Mover's Distance between Starting and Final Distributions
GB
Max Util
EO w/ Limit
EO
Earth Mover's Distance
Time Steps

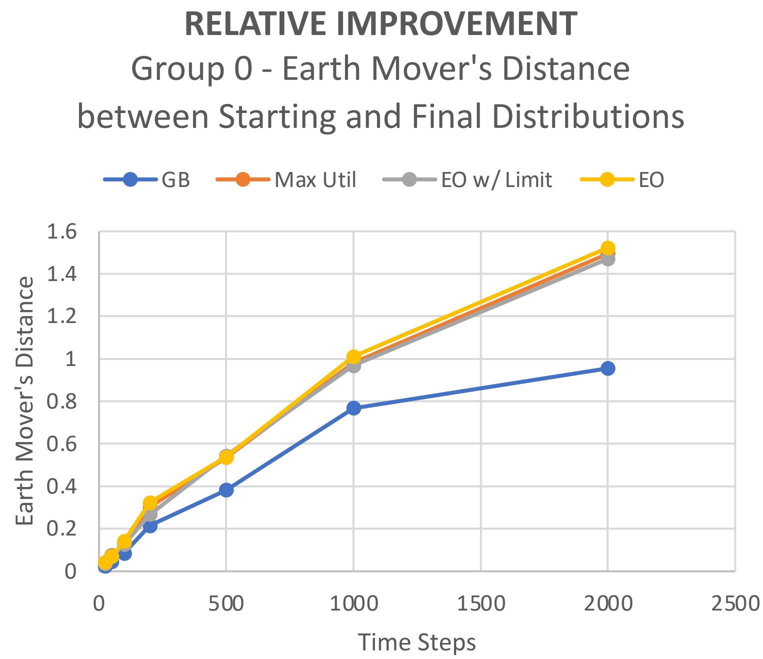
RELATIVE IMPROVEMENT
Group 0 - Earth Mover's Distance between Starting and Final Distributions
GB
Max Util
EO w/ Limit
EO
Earth Mover's Distance
Time Steps

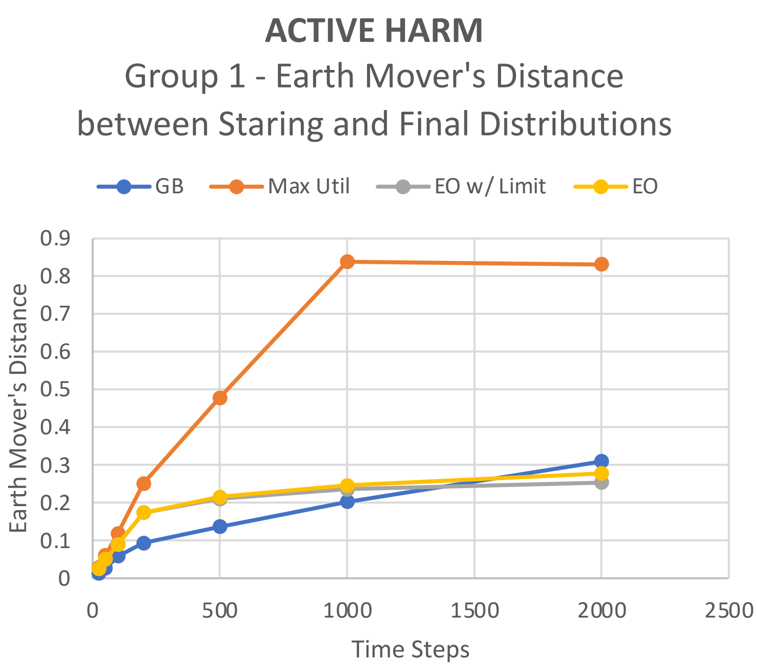
ACTIVE HARM
Group 1 - Earth Mover's Distance between Staring and Final Distributions
GB
Max Util
EO w/ Limit
EO
Earth Mover's Distance
Time Steps

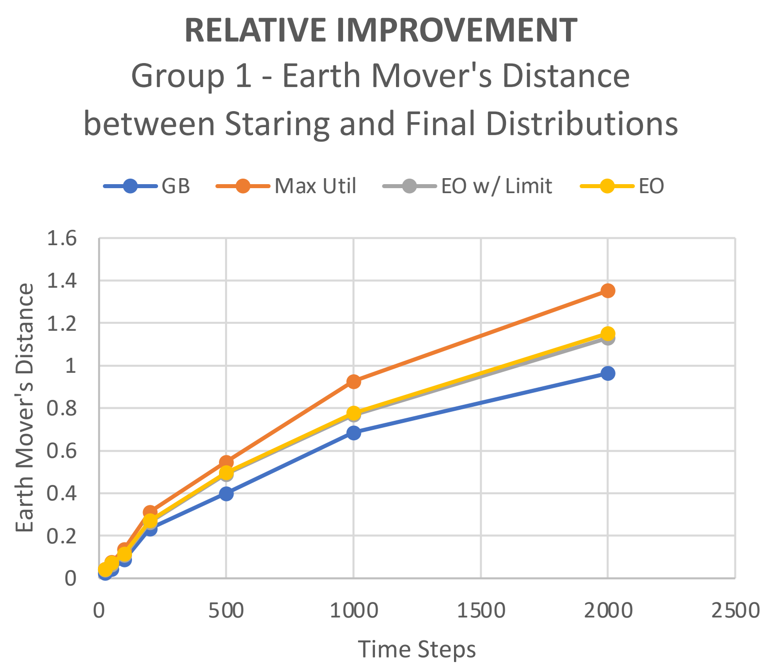
RELATIVE IMPROVEMENT
Group 1 - Earth Mover's Distance between Staring and Final Distributions
GB
Max Util
EO w/ Limit
EO
Earth Mover's Distance
Time Steps

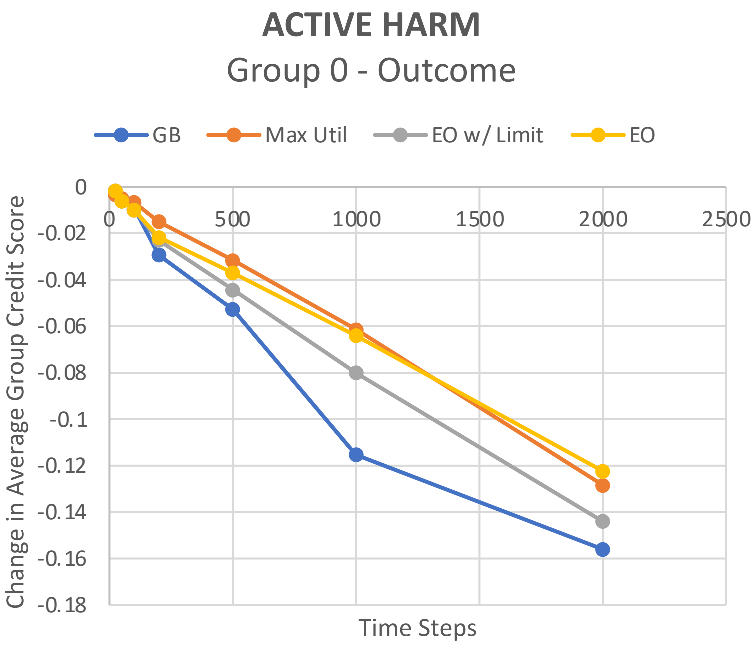
ACTIVE HARM
Group 0 - Outcome
GB
Max Util
EO w/ Limit
EO
Change in Average Group Credit Score
Time Steps
0
-0.02
-0.04
-0.06
-0.08
-0.1
-0.12
-0.14
-0.16
-0.18
0
500
1000
1500
2000
2500

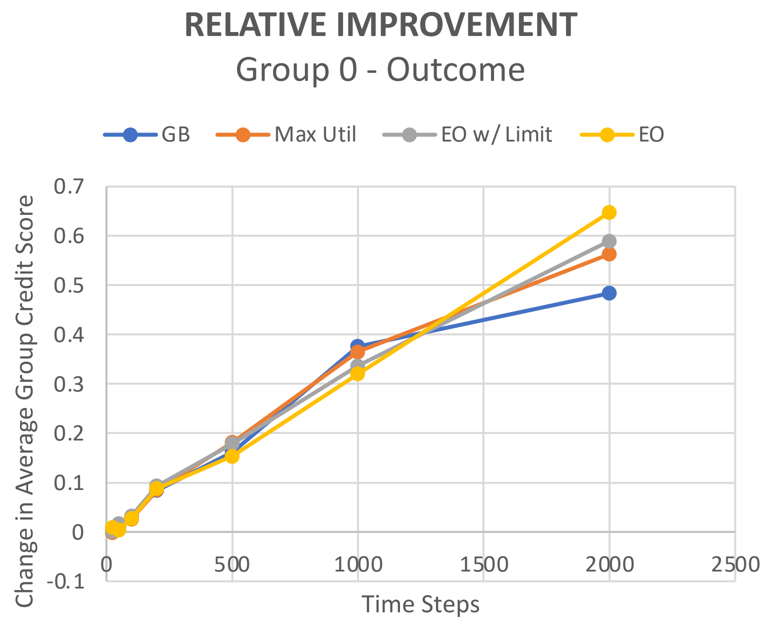
RELATIVE IMPROVEMENT
Group 0 - Outcome
GB
Max Util
EO w/ Limit
EO
Change in Average Group Credit Score
Time Steps
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
-0.1
0
500
1000
1500
2000
2500

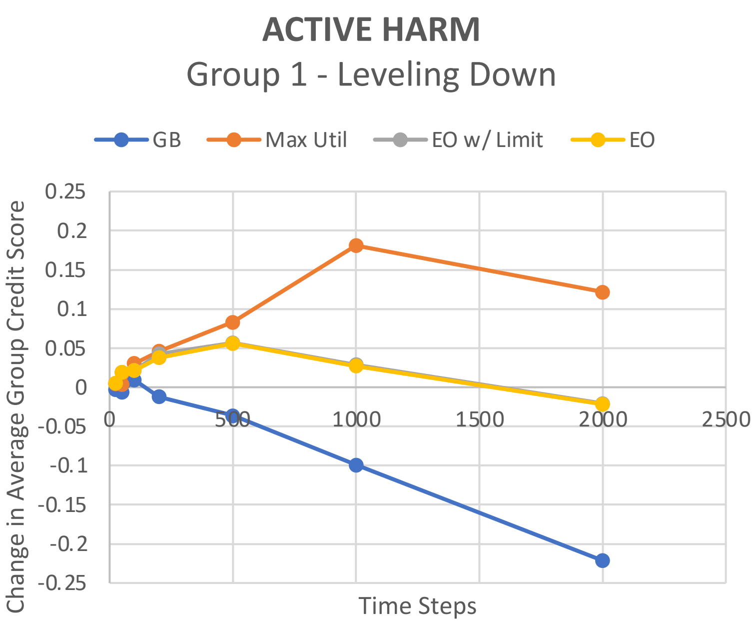
ACTIVE HARM
Group 1 - Leveling Down
GB
Max Util
EO w/ Limit
EO
Change in Average Group Credit Score
Time Steps
0.25
0.2
0.15
0.1
0.05
0
-0.05
-0.1
-0.15
-0.2
-0.25
0
500
1000
1500
2000
2500

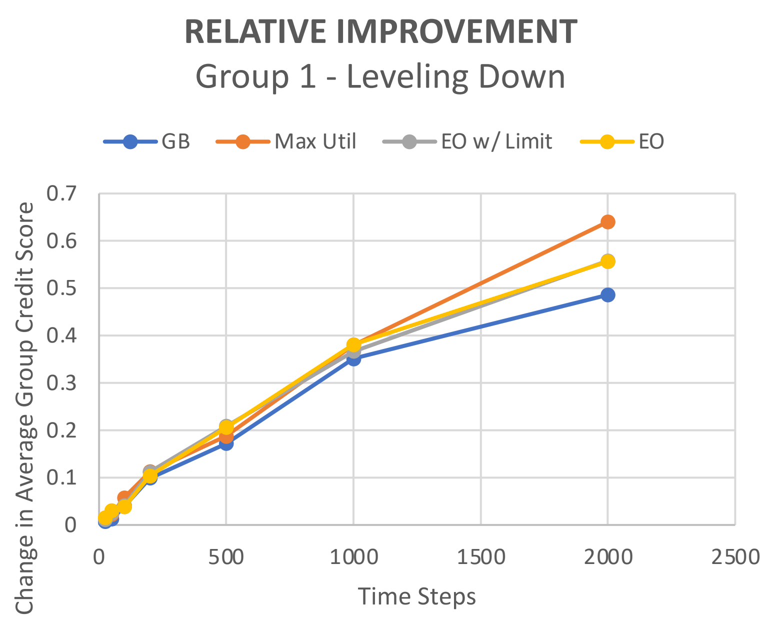
RELATIVE IMPROVEMENT
Group 1 - Leveling Down
GB
Max Util
EO w/ Limit
EO
Change in Average Group Credit Score
Time Steps
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
0
500
1000
1500
2000
2500

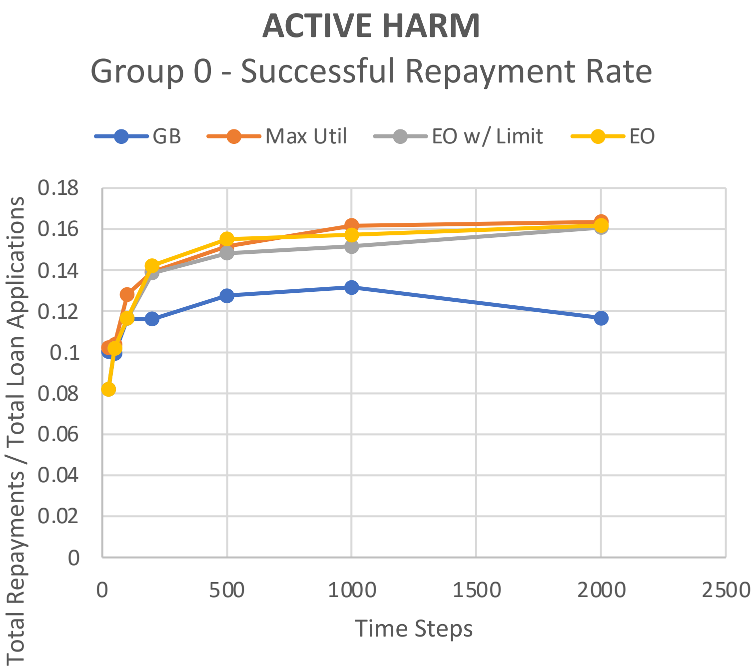
ACTIVE HARM
Group 0 - Successful Repayment Rate
GB
Max Util
EO w/ Limit
EO
Total Repayments / Total Loan Applications
Time Steps

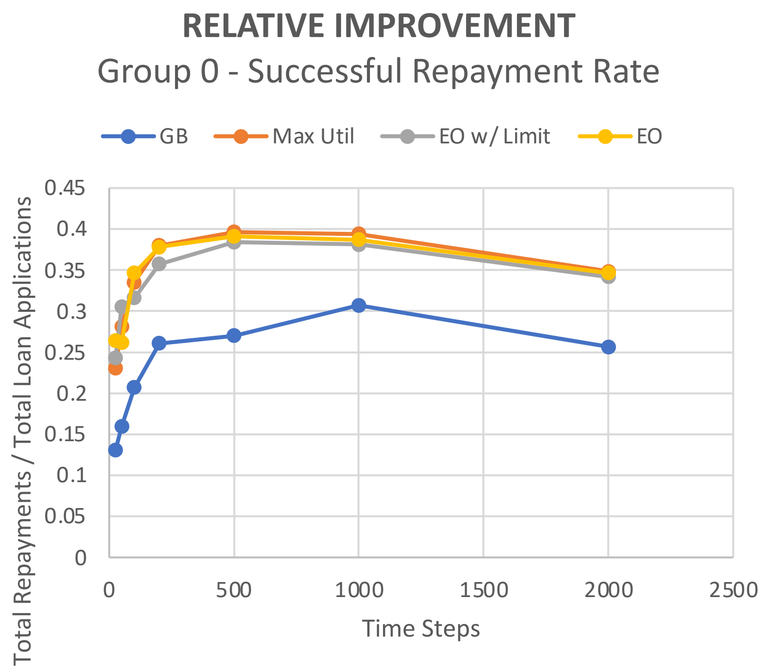
RELATIVE IMPROVEMENT
Group 0 - Successful Repayment Rate
GB
Max Util
EO w/ Limit
EO
Total Repayments / Total Loan Applications
Time Steps

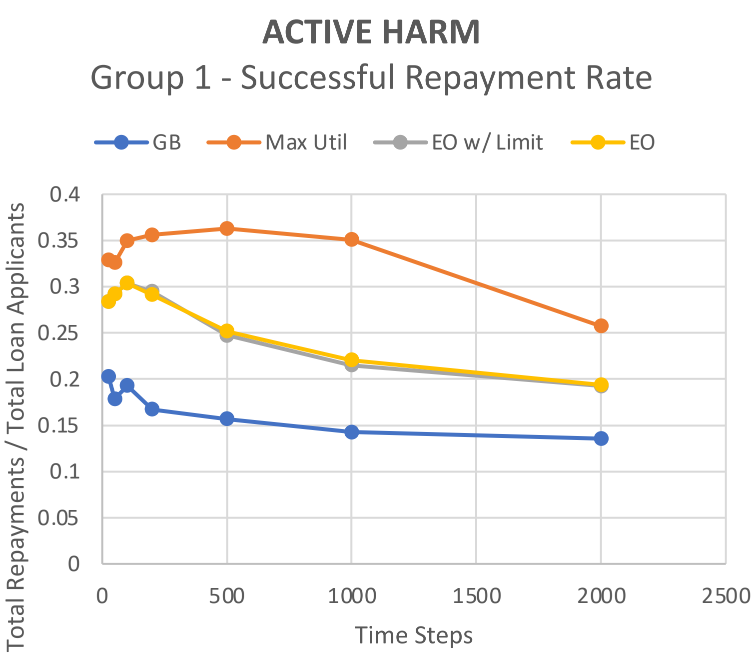
ACTIVE HARM
Group 1 - Successful Repayment Rate
GB
Max Util
EO w/ Limit
EO
Total Repayments / Total Loan Applicants
Time Steps

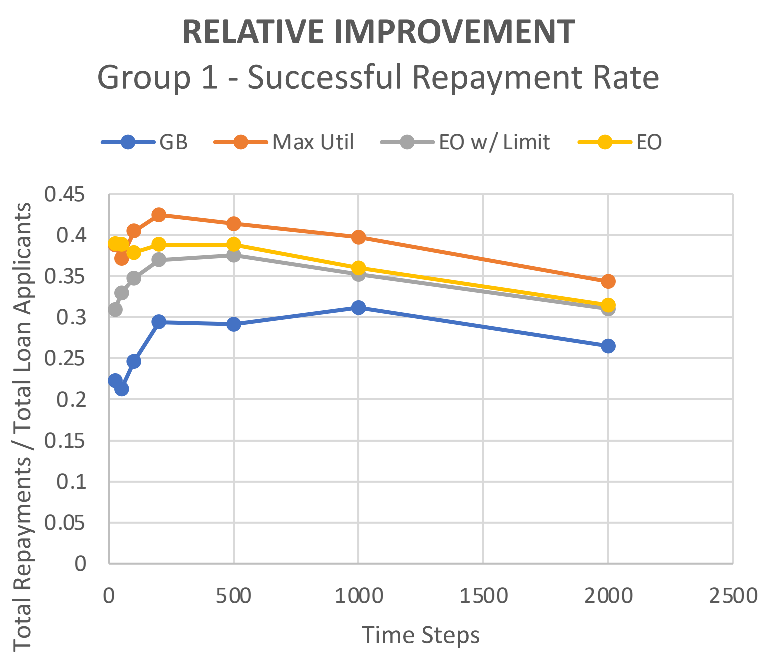
RELATIVE IMPROVEMENT
Group 1 - Successful Repayment Rate
GB
Max Util
EO w/ Limit
EO
Total Repayments / Total Loan Applicants
Time Steps

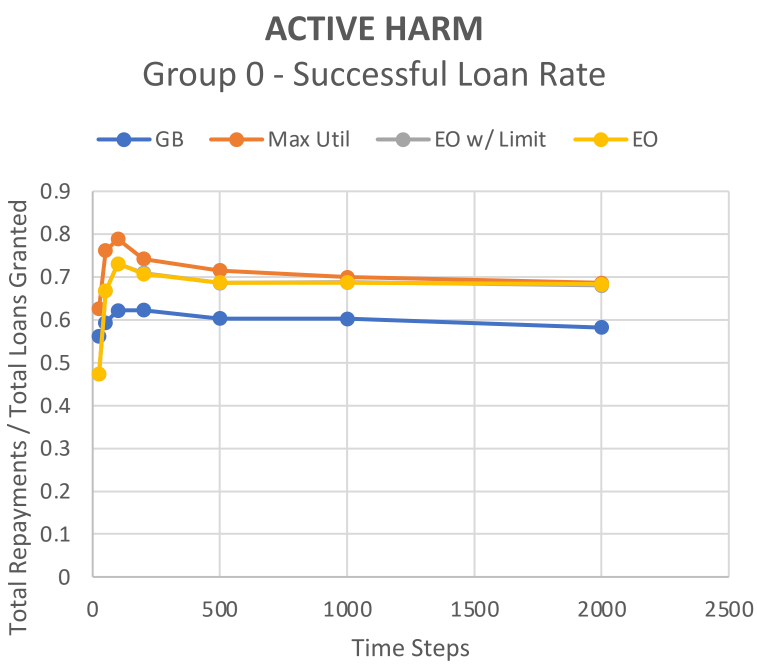
ACTIVE HARM
Group 0 - Successful Loan Rate
GB
Max Util
EO w/ Limit
EO
Total Repayments / Total Loans Granted
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
0
500
1000
1500
2000
2500
Time Steps

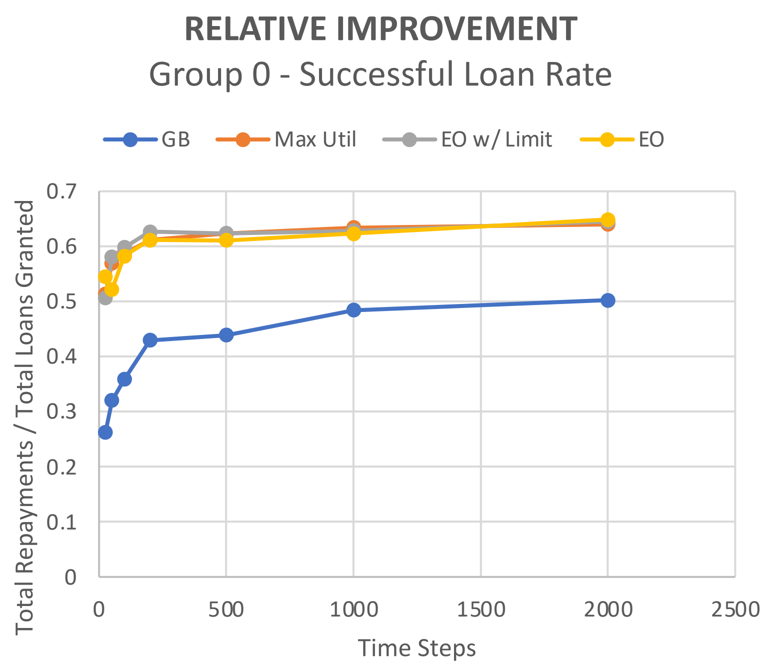
RELATIVE IMPROVEMENT
Group 0 - Successful Loan Rate
GB
Max Util
EO w/ Limit
EO
Total Repayments / Total Loans Granted
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
0
500
1000
1500
2000
2500
Time Steps

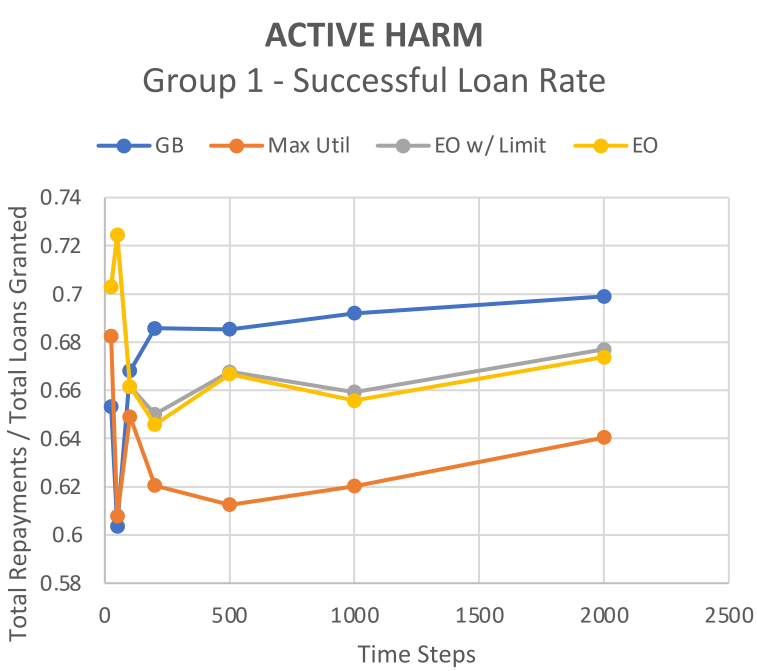
ACTIVE HARM
Group 1 - Successful Loan Rate
GB
Max Util
EO w/ Limit
EO
Total Repayments / Total Loans Granted
0.74
0.72
0.7
0.68
0.66
0.64
0.62
0.6
0.58
0
500
1000
1500
2000
2500
Time Steps

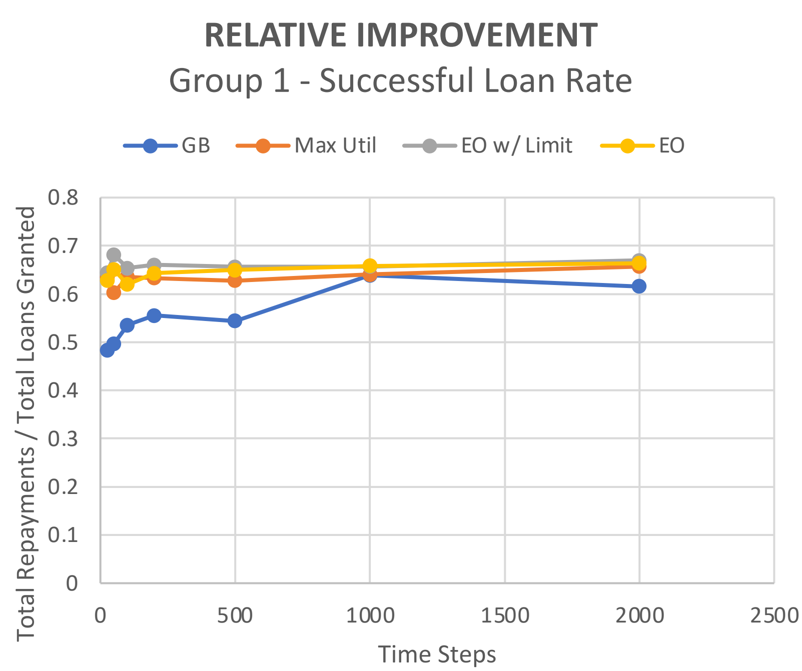
RELATIVE IMPROVEMENT
Group 1 - Successful Loan Rate
GB
Max Util
EO w/ Limit
EO
Total Repayments / Total Loans Granted
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
0
500
1000
1500
2000
2500
Time Steps

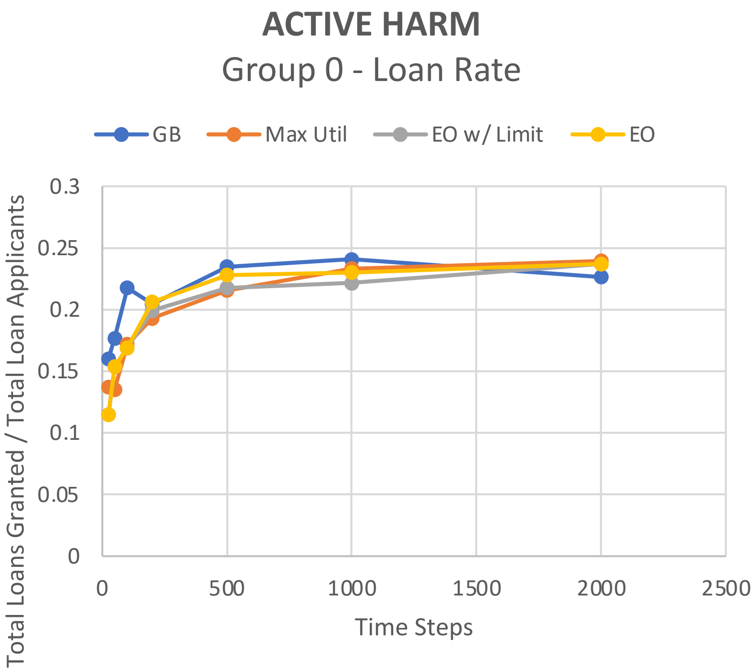
ACTIVE HARM
Group 0 - Loan Rate
GB
Max Util
EO w/ Limit
EO
Total Loans Granted / Total Loan Applicants
0.3
0.25
0.2
0.15
0.1
0.05
0
0
500
1000
1500
2000
2500
Time Steps

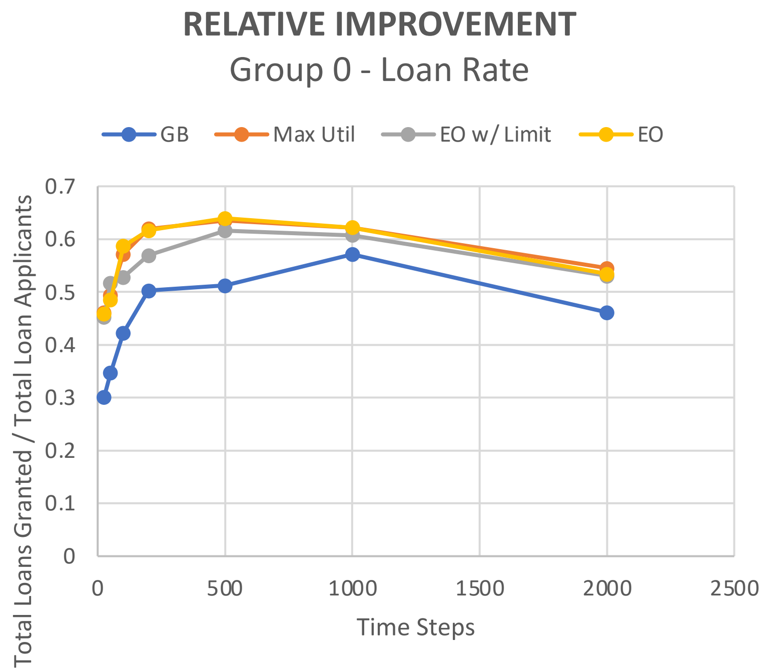
RELATIVE IMPROVEMENT
Group 0 - Loan Rate
GB
Max Util
EO w/ Limit
EO
Total Loans Granted / Total Loan Applicants
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
0
500
1000
1500
2000
2500
Time Steps

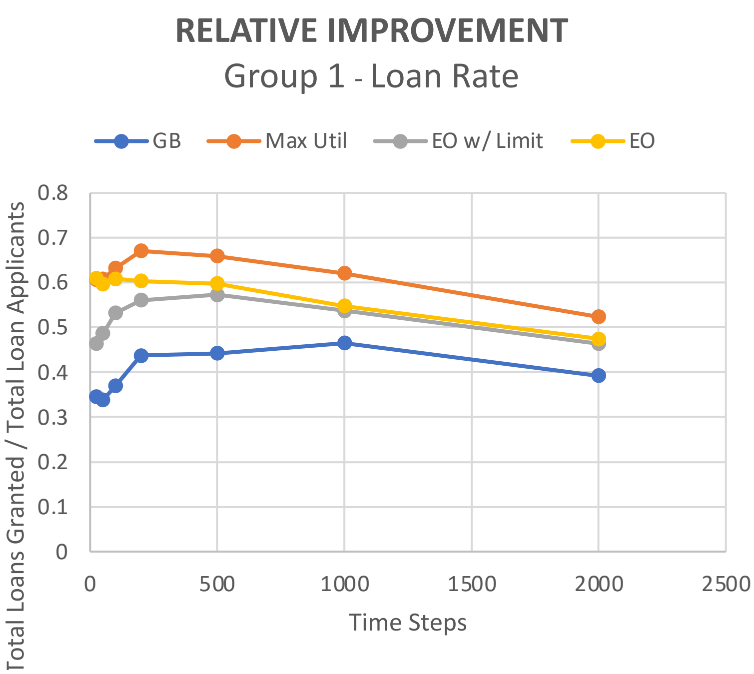
RELATIVE IMPROVEMENT
Group 1 - Loan Rate
GB
Max Util
EO w/ Limit
EO
Total Loans Granted / Total Loan Applicants
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
0
500
1000
1500
2000
2500
Time Steps

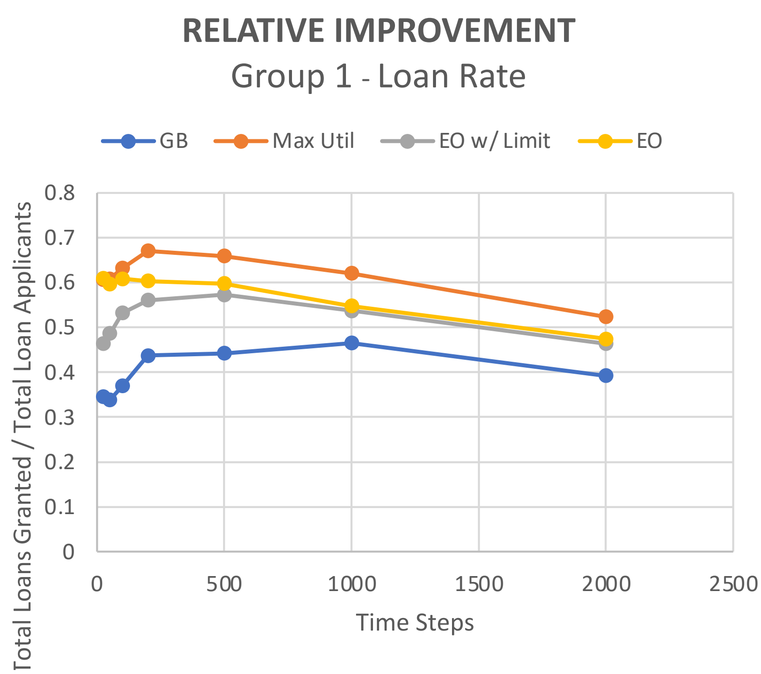
RELATIVE IMPROVEMENT
Group 1 - Loan Rate
GB
Max Util
EO w/ Limit
EO
Total Loans Granted / Total Loan Applicants
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
0
500
1000
1500
2000
2500
Time Steps