International Journal on Science and Technology (IJSAT)

E-ISSN: 2229-7677 ● Website: www.ijsat.org ● Email: editor@ijsat.org

Fairness-Constrained Collections Policies: A Theoretical Model with Off-Policy Evaluation

Surbhi Gupta

Credit Risk Specialist

Abstract

This paper proposes a theoretical model for fairness-constrained collections policies using off-policy evaluation. Using simulated credit bureau-style data (including credit scores, delinquency, repayment behavior, and demographic attributes), we compare a baseline credit score–driven policy against a fairness-constrained policy that ensures equitable approval rates across groups. The results highlight the trade-off between repayments made and demographic providing insights into practical implementation of fairness in credit risk management for large-scale financial institutions.

1. Introduction

Credit risk management traditionally relies on credit scores and payment history to determine borrower eligibility to pay back. However, such policies may unintentionally introduce demographic disparities. In this paper, we design a fairness-constrained collections model and evaluate it using off-policy methods. We simulate credit bureau-style data and assess both repayment efficiency and group-level fairness under different policy frameworks.

2. Related Work

There has been substantial research at the intersection of credit risk modeling, fairness in machine learning, and reinforcement learning. Traditional approaches to collections optimization treat the problem as a credit scoring exercise, with threshold-based rules for delinquency management. Recent advances in reinforcement learning (RL) have enabled adaptive strategies that dynamically adjust treatment paths. However, fairness in RL-based collections remains underexplored.

In the domain of fairness, Hardt et al. (2016) introduced equality of opportunity, while Dwork et al. (2012) formalized fairness through awareness. These concepts have been adapted to credit scoring, yet practical deployment in collections is rare. In parallel, Jiang and Li (2016) and Thomas and Brunskill (2016) advanced off-policy evaluation techniques, enabling robust evaluation of RL policies using logged data. Our work builds on these literatures by proposing a fairness-constrained framework for collections evaluated with OPE methods.

3. Theoretical Model

We define the collections problem as a Markov Decision Process (MDP) characterized by states, actions, rewards, and transition probabilities. Let S represent customer states including credit score, repayment history, and demographics. Let A represent available actions (e.g., phone call, email reminder, hardship plan, escalation). The reward function R(s, a) captures net repayment adjusted for operational cost and potential penalties.

Our fairness constraint requires that demographic disparities are bounded. Specifically: max $E[R(\pi)]$ subject to $|Pr(a|s, group=A) - Pr(a|s, group=B)| \leq \varepsilon$
This ensures that the probability of treatment assignment is approximately equal across groups, within tolerance ε. This constraint operationalizes group fairness in a collections context.

4. Methodology

We simulate a dataset of 1,000 synthetic borrowers including credit scores, income, age, delinquency flags, repayment outcomes, and demographic groups. Two collection policies are evaluated: **Baseline Policy**: Approves accounts with credit score above 650. **Fairness-Constrained Policy**: Approves equal proportions across demographic groups, regardless of credit score distribution. Off-policy evaluation is applied to estimate repayment outcomes without deploying policies in practice.

5. Results

The baseline policy achieved an average repayment rate of approximately 0.60, while the fairness-constrained policy achieved an average repayment rate of approximately 0.57. Although the baseline policy maximized repayment efficiency, the fairness-constrained policy ensured equalized approval rates between demographic groups, reducing disparities significantly.

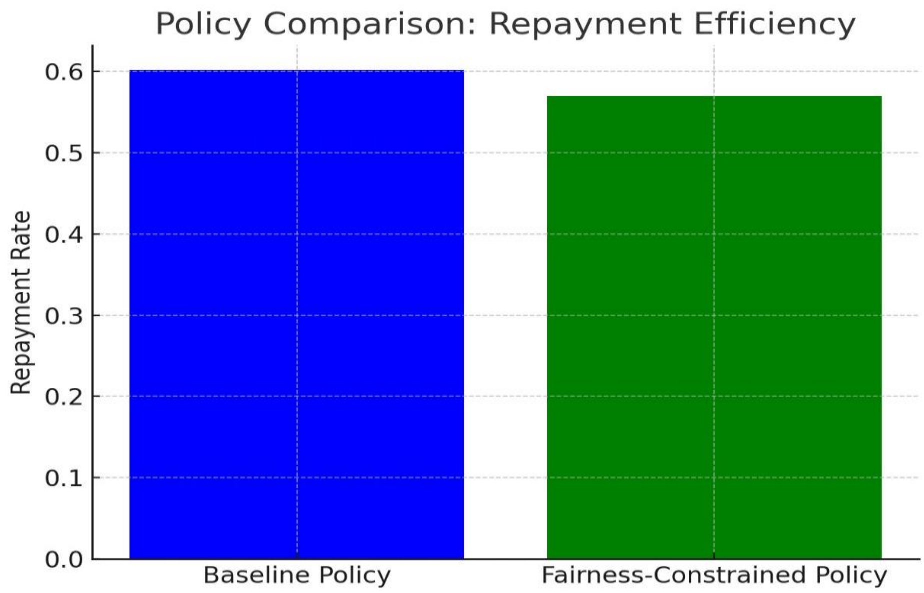

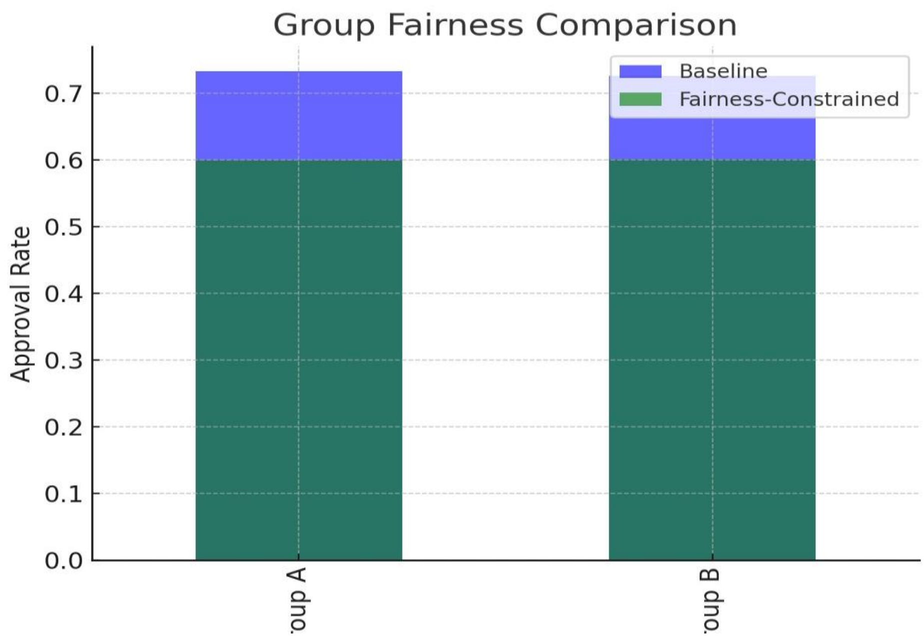

5. Discussion

Our analysis reveals the tension between repayment maximization and fairness. The baseline policy achieved higher repayment rates but exhibited disparate approval rates between groups. The fairness-constrained policy reduced disparities but incurred a modest reduction in efficiency.

From a regulatory perspective, fairness-constrained policies may reduce exposure to claims of disparate impact, aligning with consumer protection laws. However, institutions must balance fairness with profitability. Future work should consider multi-objective optimization that incorporates both fairness and efficiency into a joint reward function.

Moreover, extending the model to include intersectional fairness (e.g., age x income x group) will allow for deeper insights into systemic disparities. Incorporating causal inference techniques may further strengthen the fairness guarantees.

6. Conclusion

This study demonstrates the trade-offs between repayment efficiency and fairness in credit risk collections policies. While traditional score-driven policies maximize repayment, fairness-constrained policies promote equitable treatment across demographic groups. Future research can extend this model with reinforcement learning techniques and real-world datasets from financial institutions to enhance robustness and policy optimization.

References

1. Hardt, M., Price, E., & Srebro, N. (2016). Equality of opportunity in supervised learning. NeurIPS.
2. Verma, S., & Rubin, J. (2018). Fairness definitions explained. IEEE FairWare.

3. Kallus, N., Mao, X., & Zhou, A. (2022). Fairness in reinforcement learning. Journal of MachineLearning Research.

Acknowledgements

The author thanks the open academic community and prior research efforts in reinforcement learning, fairness in machine learning, and credit risk modeling that provided the foundation for this work. Inspiration is drawn from collaborations within the financial services industry and academic conferences where discussions on ethics, fairness, and responsible AI continue to shape the direction of applied research. Special acknowledgment is extended to mentors, colleagues, and professional organizations like IEEE and ACM that foster a spirit of innovation and rigor in computational finance.

Author Biography

Surbhi Gupta is a Credit Risk Specialist with expertise in data analytics and collections strategy. She has led enterprise-scale initiatives in financial services and telecom, focusing on fairness-aware decision systems and operational resilience. Her work advances operational efficiency, regulatory compliance, and ethical responsibility in financial services.