

Artificial Intelligence/Machine Learning Economics: Transparency, Competition, and Collusion

by

Qiaochu Wang

Submitted to the

David A. Tepper School of Business

in partial fulfillment for the requirements for the degree of

DOCTOR OF PHILOSOPHY

in the field of Business Technologies

at Carnegie Mellon University

DISSERTATION COMMITTEE:

Kannan Srinivasan (co-chair)

Param Vir Singh (co-chair)

Yan Huang

Shunyuan Zhang

Carnegie Mellon University

April 2024

Copyright © 2024 Qiaochu Wang

“物有甘苦，尝之者识；道有夷险，履之者知”

“Things could be sweet or bitter, known only to those who taste them; Roads could be smooth or treacherous, understood only by those who traverse them.”

Abstract

Machine learning (ML) algorithms are being increasingly applied to decision-making processes with far-reaching impacts extending to employment, access to credit, and education. While ML algorithms have shown great predictive power in various business applications, there are rising questions about the economic and social consequences of algorithmic decision-making, and increasing calls for algorithmic transparency. In my dissertation, I study the impact of AI/ML on economic systems on several important aspects: transparency, competition, and collusion.

Chapter 1, co-authored with Professor Param Vir Singh, Yan Huang, and Stefanus Jasin, examines the economic implications of algorithmic transparency in a hiring context. Specifically, it answers the following research question: Should firms that apply machine learning algorithms in their decision-making make their algorithms transparent to the users they affect? Despite the growing calls for algorithmic transparency, most firms have kept their algorithms opaque, citing potential gaming by users that may negatively affect the algorithm's predictive power. We develop an analytical model to compare firm and user surplus with and without algorithmic transparency in the presence of strategic users. We identify a broad set of conditions under which making the algorithm transparent actually benefits the firm. By contrast, users may not always be better off under algorithmic transparency. These results hold even when the predictive power of the opaque algorithm comes largely from correlational features and the cost for users to improve them is minimal. These results suggest that firms should not always view manipulation by users as bad. Rather, they should use algorithmic transparency as a lever to motivate users to invest in more desirable features.

Chapter 2, co-authored with Professor Param Vir Singh and Yan Huang, studies strategic information revelation in the algorithmic lending space, and its impact on competition and welfare. Financial lenders' opaque use of algorithms to screen unsecured credit applicants, coupled with high borrower uncertainty and search costs, can lead to sub-optimal credit decisions by many borrowers. Although some lenders facilitate informed decision-making for borrowers by providing personalized pre-approval probabilities, not all lenders do so. In this study, we examine how competition among lenders influences their decision to disclose approval odds to borrowers via pre-approval tools. Our findings suggest that competitive pressures, particularly in cases where lenders' algorithms are accurate, can undermine disclosure incentives. Lenders strategically employ asymmetric disclosure of pre-approval outcomes to reduce competition and differentiate their products. We demonstrate that borrower surplus is maximized when both lenders provide pre-approval tools and minimized when neither lender does so. However, mandating all lenders to provide personalized pre-approval outcomes may not necessarily enhance borrower surplus.

Chapter 3, co-authored with Professor Param Vir Singh, Yan Huang, and Kannan Srinivasan, studies algorithmic pricing and its impact on competition and collusion. The increasingly popular automated pricing strategies in e-commerce can be broadly categorized into two forms: simple rule-based algorithms, such as undercutting the lowest price, and more sophisticated artificial intelligence (AI) powered algorithms, like reinforcement learning (RL). RL algorithms are particularly appealing for pricing due to their ability to autonomously learn an optimal policy and adapt to changes in competitors' strategies and market conditions. Despite the common belief that RL algorithms hold a significant advantage

over rule-based strategies, our extensive experiments, conducted under both a canonical Logit demand environment and a more realistic non-sequential search structural demand model, demonstrate that when competing against RL pricing algorithms, simple rule-based algorithms can lead to higher prices and benefit all sellers, compared to scenarios where multiple RL algorithms compete against each other. Theoretical analysis in a simplified setting yields consistent results. Our research sheds new light on the effectiveness of automated pricing algorithms and their interactions in competitive markets, providing practical insights for retailers in selecting appropriate pricing strategies.

Acknowledgments

First and foremost, I want to express my deepest gratitude to my advisors, Professor Yan Huang, Kannan Srinivasan, and Param Vir Singh. It has been an incredible honor to have your guidance as I embarked on my academic journey. I am particularly grateful to Professor Yan Huang, from whom I have learned all the essential qualities of being an outstanding researcher. I also want to extend my heartfelt thanks to Professor Kannan Srinivasan, whose inspiration, perspective, and insights have had a profound impact on me, expanding my understanding of the field and pushing me to think more critically. I also want to extend my special gratitude to Professor Param Vir Singh, who has invested countless hours in nurturing my knowledge, confidence, skills, and way of thinking from the ground up, transforming me from a novice master's student into a real marketing researcher I always dreamed to be.

I also want to extend my sincere gratitude to my dissertation committee member, Professor Shunyuan Zhang, who has been my role model since the very first day of my Ph.D. journey, and her unwavering support has been a constant source of inspiration and encouragement. Furthermore, I want to express my heartfelt appreciation to my co-author, Professor Stefanus Jasin. Your expertise has been instrumental in refining my modeling and writing skills. Working alongside you greatly enhanced the quality of my research.

I would like to extend my special gratitude to Professors Nitin Mehta, Xiao Liu, and Runshan Fu for their invaluable mentorship and unwavering support throughout my doctoral journey, particularly during the stressful period of the job market. Additionally, I am deeply thankful to Professors Zoey Jiang, Mohsen Foroughifar, Alan Montgomery, Joy Lu, Christopher Olivola, Nikhil Malik, Aniko Oery, Yong Tan, Elina Hwang, Maryam Saeedi, Minkyung Kim, Evelyn Gong,

Karan Singh, Ananya Sen, BeiBei Li, Vineet Kumar, for their insightful advice and constructive feedback on my dissertation and job interview preparations.

My gratitude extends to Professors Tridas Mukhopadhyay, Manmohan Aseri, Sridhar Tayur, Tim Derdenger, Peter Boatwright, Jeffrey Galak, James Best, Ali Shourideh, Charles Zheng, Robert Miller, George Loewenstein, Anindya Ghose, David Childers, Aislinn Bohren, Peter Stuetzgen, Ariel Zetlin-Jones, June Shi, and Serim Hwang for the significant impact they have had on my journey through my Ph.D. Their wisdom and insights have been pivotal in shaping my academic and personal growth.

I must also express my heartfelt appreciation to my undergraduate advisor, Professor Jing Yang, and my mentors and letter writers from my master's program, Professors Irving Oppenheim, Jacobo Bielak, Matteo Pozzi, and Sivaraman Balakrishnan. A special note of gratitude is due to Professor Alexandra Chouldechova, who introduced and encouraged me to apply to this Ph.D. program, which has been a turning point in my academic career.

I am particularly grateful to Laila Lee and Lawrence Rapp for their dedicated care of all the Ph.D. students over the years, and for their tireless assistance with all the administrative tasks.

I am deeply thankful for the mutual support shared with my friends and fellow students, both within the realms of academia and in our personal lives. In the Business Technologies and Marketing programs: Samuel Levy, Behnam Mohammadi, Liying Qiu, Flora Feng, Julie Wang, Yuan Yuan, Jaymo Kim, Angela Xiao, Michelle Min, Jiaming Jiang, and Ziqian Ding. In other disciplines: Nikhil George, Zhaoqi Cheng, Duyu Chen, Jaepil Lee, Jason Gates, Martin Michelini, Shuoqi Sun, Rudy Zhou, Savahna Tang, Yikang Shen, Leo Li, Xiaonan Hong, Hoang-Anh Nguyen, Jenny Oh, Yuanchun Ye, Ruilong Zhang, Xiyang Hu, Yue Zhao, Jingwei Dai, and Yuan An.

My heartfelt gratitude goes to my parents, Yunhe Wang and Yuhong Sun, whose unwavering love and support have been my foundation for many years. To my extended family, Jingyi Cheng, Yumei Sun, Chongbiao Cheng, and Xiuyun Ni, your backing and encouragement have been a constant source of strength. Finally, a special word of thanks to my partner, Lavender Yang. Together, we have faced life's challenges and celebrated its joys, side by side, hand in hand. Graduating alongside you stands as one of the most cherished milestones in my life.

Contents

- 1 Algorithmic Transparency With Strategic Users 1**
- 1.1 Introduction 1
- 1.2 Literature 10
- 1.3 Model 15
 - 1.3.1 The Firm’s Utility 18
 - 1.3.2 The Agents’ Utility 19
 - 1.3.3 Game Sequence and Equilibrium Concept 20
 - 1.3.4 Additional Parametric Assumptions 21
- 1.4 Analysis 22
 - 1.4.1 Opaque Scenario 23
 - 1.4.2 Transparent Scenario 25
 - 1.4.3 The Firm’s Decision on Algorithmic Transparency 27
 - 1.4.4 The Effect of the Predictive Power of the Correlational Feature (λ) and the Fraction of High-Talent Agents (θ) 34
 - 1.4.5 Agents’ Welfare 36
- 1.5 The Stackelberg Model 41
- 1.6 Extensions 45
- 1.7 Conclusion 46
 - 1.7.1 Summary of Results 46
 - 1.7.2 Implications for Managers 46

1.7.3	Implications for Public Policy	49
1.7.4	Generalizability of the Results	49
1.7.5	Future Research Directions	50
2	Algorithmic Lending, Competition, and Strategic Information Disclosure	51
2.1	Introduction	51
2.2	Contributions to Literature	55
2.3	Model	57
2.3.1	Borrowers	57
2.3.2	Lenders	58
2.3.3	Borrowers' Belief	59
2.3.4	Game Sequence	60
2.3.5	Illustration	61
2.4	Analysis	63
2.4.1	Additional Assumptions	63
2.4.2	Analysis: Approving Stage	64
2.4.3	Analysis: Applying Stage	64
2.4.4	Analysis: Pricing Stage	64
2.4.5	Analysis: Pre-approval Revealing Stage	77
2.4.6	Borrower Surplus and Social Welfare	81
2.5	Mandating Algorithm Revelation	83
2.6	Conclusions	87
2.6.1	Summary of Results	87
2.6.2	Managerial Implications	88
2.6.3	Implications for Public Policy	89
2.6.4	Limitations	89
2.6.5	Funding and Competing Interests	90

3	Algorithms, Artificial Intelligence and Simple Rule-Based Pricing	91
3.1	Introduction	91
3.2	Relevant Literature	96
3.3	Competition With Pricing Algorithms	98
3.3.1	The Market	98
3.3.2	The Algorithm	99
3.3.3	Competition	103
3.4	Market Environment 1: Logit Demand Model	103
3.4.1	Experiment Outcomes: Baseline Setting	105
3.4.2	Experiment Outcomes: Alternative Settings	114
3.5	Market Environment 2: A Structural Demand Model	120
3.6	Theoretical Analysis	124
3.7	Discussion and Conclusion	127
A	Appendix for Chapter 1	129
A.1	Notation Summary	129
A.2	Discussions of Extensions of the Main Model	129
A.2.1	Gaming is Costly	130
A.2.2	Endogenous Wage	133
A.2.3	Agents Have Incorrect Belief of λ	135
A.3	Omitted Discussions	137
A.3.1	Derivation of the Lower and Upper Bound of α and β	137
A.3.2	Discussion on the case where β does not satisfy Assumption 1.3.4	139
A.3.3	Discussion on the Effect of θ and λ on Decision on Algorithmic Transparency Under the Stackelberg Model	141
A.4	Mathematical Appendix (Proofs of All Results)	142
A.4.1	Proof of Lemma 1	142
A.4.2	Proof of Lemma 2	152

A.4.3	Proof of Lemma 3	157
A.4.4	Proof of Lemma 4 and Lemma 5	157
A.4.5	Proof of Lemma 11 and Lemma 12	158
A.4.6	Proof of Proposition 3	159
A.4.7	Proof of Proposition 11	161
A.4.8	Proof of Proposition 12	161
B	Appendix for Chapter 2	165
B.1	Notation Summary	165
B.2	Model Extension	165
B.2.1	Multiple Applications	165
B.2.2	Vertically Differentiated Products	169
B.2.3	Asymmetric Lender Accuracy	174
B.2.4	Correlated Predictions	175
B.3	Mathematical Appendix	176
B.3.1	Proof of Lemma 6	176
B.3.2	Proof of Lemma 7	177
B.3.3	Proof of Lemma 8	178
B.3.4	Proof of Proposition 4	180
B.3.5	Proof of Proposition 5	181
B.3.6	Proof of Proposition 6	182
B.3.7	Proof of Proposition 7	184
B.3.8	Proof of Proposition 8	185
B.3.9	Proof of Proposition 9	187
C	Appendix for Chapter 3	189
C.1	Experiment Outcomes: Alternative Settings	189
C.1.1	More Advanced RL Algorithms	189

C.1.2	Multiple Players in the Market	189
C.2	Structural Demand Model Details	190
C.2.1	Data	190
C.2.2	Testing Search Models	192
C.2.3	A Non-sequential Search Structural Model	194
C.3	SKU Subset Selection Details	203
C.4	Theoretical Framework Details	204
C.4.1	Learning Dynamics of The Q-Q scenario (Proof of Lemma 9)	206
C.4.2	Learning Dynamics of The Q-R scenario (Proof of Lemma 10)	210
C.4.3	Proof of Theorem 3	211
C.5	A Hierarchical Bayes Estimation Algorithm	212

Bibliography		217
---------------------	--	------------

Chapter 1

Algorithmic Transparency With Strategic Users

1.1 Introduction

Machine learning algorithms are being increasingly applied to decision-making processes with far-reaching impacts extending to employment, access to credit, and education (Schellmann and Bellini, 2018, Fu et al., 2021a). However, firms typically keep these algorithms as closely guarded secrets, on par with KFC or Coca-Cola’s recipes. As a result, these algorithms stay opaque to the people whom they affect and lack clear explanations for the decisions they make.

Our study is motivated by the growing calls from different parts of society to require firms to make their algorithms transparent. According to American privacy law expert Marc Rotenberg: “At the intersection of law and technology – knowledge of the algorithm is a fundamental right, a human right.”¹ The European Union’s General Data Protection Regulation (GDPR) dictates that, whenever personal data is subject to automated decision making, people have “the right to obtain human intervention on the part of the controller” or

¹See “Algorithmic Transparency: End Secret Profiling,” <https://epic.org/algorithmic-transparency/>

the right to explanation.²

While making algorithms transparent is desirable, it can open a door to gaming, which potentially adversely affects the algorithms' predictive power. If strategic agents were to know the information about a classifier (i.e., how observed attributes affect the classification outcome), they may manipulate their attributes to receive a more desirable classification, hurting the predictive power of the algorithm. In financial and economic policy making, this problem is widely known as Goodhart's Law, which proclaims that "when a measure becomes a target, it ceases to be a good measure" (Goodhart, 1984). A similar notion is captured in the Lucas critique (Lucas et al., 1976). In fact, Fair Isaac Corporation keeps its exact credit scoring formula secret to make it more difficult for consumers to game the algorithm (Citron and Pasquale, 2014a). Similarly, Google continues to modify its secret search ranking algorithm to keep businesses and people from gaming the system (Segal, 2011).

Motivated by the calls for algorithmic transparency and the threat of manipulation by agents to transparent algorithms, we investigate how algorithmic transparency may affect firms and agents. First, from the perspective of the firm (the decision-maker), we ask, is there any advantage in making its algorithm transparent even when there is the potential of manipulation by the agents? Second, we ask, would the agents be better off or worse off if firms make their algorithms transparent (i.e., if the agents receive more information about the factors affecting algorithmic decisions)? Third, we ask, how are the results affected by the predictive power of those features that are more susceptible to gaming by the agents? Finally, we ask, how does the market composition in terms of desirable and undesirable agents affect these results? In this paper, we develop and analyze a game-theoretic model to answer these questions.

We explicitly model the agents as strategic and the algorithm designer (the firm) is aware of the potential for manipulation. Hence, the firm can react to gaming by the agents. For example, consider a setting where the firm collects data, trains an algorithm that maps a set

²See "Algorithmic transparency and the right to explanation: Transparency is only the first step," <https://www.apc.org/en/blog/algorithmic-transparency-and-right-explanation-transparency-only-first-step>

of observed features to a classification outcome, and publishes a decision rule. If agents desire to be positively classified, they would manipulate their values of the features to achieve it. However, the firm would be aware that the behavior of the agents has changed. It will then update the model and the decision rule. The agents' behavior would change once again. Over time, the firm will iterate to a *fixed point* – this decision rule would be the best response to the agents' strategic behavior.

More specifically, in this paper, we consider a *job hiring* scenario where a risk-neutral firm offers a fixed wage and wants to recruit only highly productive agents. There are two types of agents, High talent (H) and Low talent (L). The H type agents are more productive compared to the L type agents. While the type of each agent is fixed, it is not observed by the firm ex-ante. The firm, however, does have access to a number of observed features (i.e., the *observables*), which it uses to infer the agents' types (i.e., using historical data and an algorithm) and to determine a decision rule for hiring the agents.

We model two types of observables, *causal* and *correlational*. Typically, in machine learning models, the designer focuses primarily on model accuracy, not causality. However, any features that are captured by the machine learning model can still be classified as either *causal* or *correlational*. For simplicity, we consider only two features; one is causal and the other is correlational. There are several characteristics of these features that are important for our model. By definition, the causal feature impacts the productivity of the agent, whereas the correlational feature does not. The difference between 'causal' and 'correlational' features is similar to the difference between 'improvement' and 'gaming', two terms commonly used in the strategic classification literature (Kleinberg and Raghavan, 2019, Alon et al., 2020, Haghtalab et al., 2020). Putting an effort to change the value of a causal feature in the decision maker's favorable direction is called an 'improvement' because this is beneficial to the decision maker or the society. By contrast, putting effort on to change the value of a correlational feature does not benefit the firm and is a waste from the perspective of social welfare.

The agents can game (alter) both features by incurring a cost. As is typically assumed in most signaling game models (Spence, 1973), the H type agents have a cost advantage on the causal feature over the L type agents. As for the cost of improving the correlational feature, in the main model, we assume that this cost is independent of the agents' type and its value is marginally above zero.³

The assumptions behind the cost structures of the causal and correlational features warrant some discussion. When an H type agent has a significant cost advantage over an L type agent on the causal feature, it will trivially lead to a separating equilibrium where only the H type agents get high values on the causal feature. In the case where the cost advantage of the H type agents on the causal feature is not too large, the firm would want to include the correlational feature in the machine learning model and in the decision rule. In this case, the correlational feature would provide an additional value in separating the two types. It is easy to see that the correlational feature is the one that is more susceptible to gaming. If the agents game it, this feature can lose its predictive power. This is precisely the reason that is typically purported for opposing algorithmic transparency. If the cost to alter the correlational feature is very high, or if the H type agents have an advantage on it, either gaming will not happen or gaming will be more favorable to the H type agents. In such a case, making the algorithm transparent would be either better or at least as good as keeping it opaque for the firm. Our interest is in investigating whether algorithmic transparency can be better for the firm as opposed to keeping the algorithm opaque even when the H type agents have no cost advantage on the correlational feature and the cost to manipulate the correlational feature is minimal, which makes the algorithm particularly highly susceptible to gaming.

We solve the game between the firm and the agents in two scenarios, an *opaque* algorithm scenario and a *transparent* algorithm scenario. In both scenarios, we use *perfect Bayesian equilibrium* (PBE) as the solution concept. In the opaque scenario, the agents move first.

³For completeness, we also discuss the setting where the cost is significantly larger than zero in Appendix A.2.1.

In this scenario, we assume that the agents are aware of the causal feature but have little knowledge of the correlational feature. Consequently, the agents can only improve the causal feature. The agents know that the firm uses a correlational feature, but they do not know what that feature is. However, it is common knowledge that the feature has predictive power that can help the firm separate the H type agents from the L type agents. This assumption is reasonable: If a feature has no predictive power in separating the two types of agents, the machine learning algorithm would simply discard it.

In the transparent scenario, the firm moves first and publishes its algorithm. The agents observe this algorithm and know what features are considered by the algorithm including their respective weights. In this scenario, the agents can game both the causal and the correlational features. They also know how the correlational feature correlates with their types before any gaming occurring. Finally, the firm decides who to hire based on the agents' ending feature values.

It is worth noting that although we assume the agents and the firm move sequentially in both the opaque and transparent scenarios, we do not model these processes as Stackelberg leadership models in our main analysis. The reasons are as follows. In a Stackelberg model, the first mover typically anticipates that any deviation from the current strategy will trigger the second mover to react accordingly. However, when the first mover consists of many uncoordinated agents, this condition will be violated. Specifically, in the opaque scenario, given that the agents are non-cooperative, the unilateral deviation of a single agent will not change the follower's (i.e., the firm's) strategy.⁴ As for the transparent scenario, there is a single first mover (i.e., the firm), and it is possible for the firm to commit to the published algorithm. Consequently, a Stackelberg model is a valid model in this scenario. However, a Stackelberg model provides an advantage to the first mover. As a result, the firm may prefer the transparent algorithm over the opaque algorithm simply because of the first mover

⁴Similar assumptions have been made in many papers in the job market signaling literature. When constructing the equilibrium, those papers assume that there is no profitable deviation for any single worker (agent) given the firm's strategy, i.e., a single worker's deviation in education will not change the firm's belief and wage menu (Spence, 1973, Alós-Ferrer and Prat, 2012, Daley and Green, 2014).

advantage that the transparent algorithm provides. To clearly explain the mechanisms of the firm’s preference between algorithmic transparency and opacity and show that the firm could still prefer the transparent algorithm over the opaque algorithm even in absence of the first mover advantage, we will first focus our analysis in this paper on the case where the firm either does not have the commitment power or the firm only reveals which features it uses but not the full details of its hiring strategy (akin to ‘partial transparency’). Later, in Section 1.5, we will discuss the case where the firm has commitment power and reveals the full details to the agents (akin to ‘full transparency’). Overall, we will show that the key insights from the non-Stackelberg model will only be strengthened in the Stackelberg model.

In addition to analyzing the models discussed above, as briefly described in Section 1.6 and further elaborated in Appendix A.2, we also analyze three extensions by relaxing some of the assumptions made in the main model. These extensions include: (1) the case where the cost of improving the correlational feature is significantly larger than zero, (2) the case where the wage is endogenously chosen by the firm, and (3) the case where the agents have incorrect beliefs about the predictive power of the correlational feature. Our analysis for these three extensions shows that the results from the main analysis are robust to these alternative assumptions and modeling choices. Specifically, the main results of our paper will not change qualitatively in extension (3), and they will only be strengthened in extensions (1) and (2).

Key results and insights

Our first result in this paper challenges the conventional wisdom that making algorithms transparent will always hurt the firm economically. We identify a broad set of conditions under which making the algorithm transparent is actually beneficial for the firm. The key intuition behind this result is driven by how the H type and L type agents respond to algorithmic transparency. Since investment on the causal feature is costly and since the H type agents have an advantage on the correlational feature (we assume that an H type agent has a higher probability of being ‘high’ on the correlational feature, otherwise the firm will have no incentive to use this feature in the first place), the H type agents would invest in improving

the causal feature only to the extent that it, along with the correlational feature, helps them separate themselves from the L type agents. When the algorithm is made transparent, the L type agents game the correlational feature. As a result, the L type agents become similar to the H type agents on that feature, and the predictive power of the correlational feature decreases. Hence, the H type agents have to invest more in the causal feature to separate themselves from the L type agents. When the H type agents have a significant cost advantage over the L type agents on the causal feature, this leads to full separation, which benefits the firm. When the H type agents only have a marginal cost advantage over the L type agents on the causal feature, the L type agents will also invest significantly in the causal feature. In this case, both the H type and L type agents become more productive because of their higher investment in the causal feature. Although the firm cannot separate the two types of agents in this case, when the impact of the causal feature on productivity is above a certain threshold, the average productivity of the hired agents is significantly higher than in the opaque scenario. In other words, making the algorithm transparent allows the firm to motivate the agents to invest in improving features that are valuable to the firm.

Our second result in this paper is that the agents are not always better off under the transparent scenario. This might appear counter-intuitive at first: Since the agents would have access to more information under the transparent scenario, one would think that they should be better off under the transparent scenario. However, we show that there are conditions under which the agents are worse off in the transparent scenario. Interestingly, in most cases where the firm prefers the transparent scenario, the agents would prefer the opaque scenario, and vice versa. But we also identify a set of conditions where both the firm and the agents prefer the transparent scenario. The intuition for our second result is similar to that for the first result. The firm prefers the transparent scenario in situations where transparency motivates the agents to invest highly in the causal feature. When the cost of investment is high and the H type agents have a significant advantage, only the H type agents will invest in the causal feature. In this situation, although only the H type agents are hired, they are

worse off due to the high cost that they incur. When the investment cost is moderate and the H type agents have a marginal cost advantage over the L type agents, both types of agents invest in the causal feature, and both are hired. The agents are better off because they have to incur only a moderate cost for being hired. Simultaneously, the firm is also better off since the average productivity of the hired agents is higher when the impact of the causal feature on productivity is above a certain threshold.

Our third result shows that it is possible for the firm to prefer algorithmic transparency when the correlational feature has high predictive power and prefers opaque algorithm when the correlational feature has low predictive power. This result also appears to be counter-intuitive since one would naturally expect that, as the predictive power of the correlational feature increases, the firm would be better off keeping the algorithm opaque. We find that this is not always the case. The key intuition behind this result is as follows: When the correlational feature is good at separating the H type agents from the L type agents, the H type agents have little incentive to invest in the causal feature in the opaque scenario. However, under transparency, the H type agents lose this advantage and have to invest in the causal feature to separate themselves from the L type agents. As a result, the firm can hire more productive agents under transparency.

Our fourth result is that, when the fraction of the H type agents on the market is higher, the firm may have a stronger preference for the transparent algorithm under certain conditions. More specifically, the fraction of H type agents affects the firm's surplus significantly but does not have a large enough impact to alter the firm's decision for/against transparency when the cost for improving the causal feature is either too high or too low, or the H type agents have a large cost advantage over the L type agents. However, when the cost for improving the causal feature and the cost advantage of the H type agents are both moderate, the firm has a stronger preference for algorithmic transparency as the fraction of the H type agents on the market increases. In this cost range, both the H type and L type agents would improve the causal feature. While the firm is unable to separate the two types and hires both types,

the number of the L type agents that are hired becomes smaller as the fraction of the H type agents on the market increases.

This paper makes several contributions. It is one of the first to provide an analytical model that systematically compares a firm’s decision for algorithmic transparency *versus* opacity in the presence of strategic agents. We show that, counter to conventional wisdom, the firm can be better off under algorithmic transparency. Moreover, in most cases where the firm prefers algorithmic transparency, the agents will be worse off. Agents underinvest in the causal feature when the algorithm is opaque. Consequently, the firm depends heavily upon the correlational feature to separate different types of agents from one another. Our analysis and results show that the firm should not always worry about the potential loss of the predictive power of the correlational feature in its machine learning model under transparency. Rather, it can use algorithmic transparency as a lever to motivate the agents to invest more in the causal feature. The firm would typically be reluctant to adopt algorithmic transparency when their machine learning model derives large predictive power from the correlational feature. However, we show that the firm should recognize that investment in the causal feature by agents is endogenous. The H type agents are less likely to invest in the causal feature when the correlational feature is sufficient to separate them from the L type. This is the situation where the firm should be willing to sacrifice the predictive power of the correlational feature. We have demonstrated our results in the setting where (1) the firm is certain it will lose the predictive power of the correlational feature and (2) the firm does not have a first mover advantage under algorithmic transparency. Intuitively, one would think that algorithmic transparency would be bad for the firm when these two conditions hold. However, we show that, once we consider the endogenous investment in the causal feature, the firm would be better off making the algorithm transparent.

Organization of the paper

The rest of the paper is organized as follows. We discuss how we build upon and contribute to the literature in Section 1.2. The details of our main model are presented in Section 1.3.

In Section 1.4, we present the analysis of the main model. In Section 1.5, we discuss the model of Stackelberg competition. In Section 1.6, we describe three extensions of the main model. We conclude the paper in Section 1.7. Unless otherwise noted, all the proofs can be found in Appendix A.4.

1.2 Literature

Algorithmic transparency is a relatively new topic, but it is closely related to the literature on information asymmetry. Following the canonical job market signaling model developed by Spence (1973), a rich stream of research has focused on the interaction between a decision maker and strategic agents under asymmetric information. Some of this research is focused on the agents' side and study how agents strategically reveal their type under various market conditions (e.g., the stream of signaling game literature). Other research is focused on the decision makers' side, studying how they can design an optimal algorithm to extract agents' private information (e.g., the stream of strategic classification literature). Our work on algorithmic transparency is built upon and contributes to both these streams of literature.

In both the opaque and transparent scenarios, the interaction between the firm and the agents can be adapted into the signaling game framework, where individuals (senders) first send signals, and the firm (receiver) then makes hiring decisions based on the observed signals (Spence, 1973, Daley and Green, 2014, Engers, 1987, Weiss, 1983). Signaling game models often focus on specifying the equilibrium outcome under various market conditions. Receivers are assumed to be in a competitive environment and earns zero profit in the equilibrium. All the surplus is extracted by senders. The equilibrium concept typically used in signaling game models is the perfect Bayesian equilibrium (PBE), where three conditions are satisfied: senders use optimal strategies facing the wage offer, receivers give wage offers such that they will obtain zero profit, and the receivers' beliefs about the senders' type given the signal are consistent with the truth. Similar to the signaling games, we also specify the equilibrium outcome under various market conditions in both the transparent and opaque scenarios.

However, in our model, the firm offers a fixed wage and focuses on designing the algorithm to increase its chances of hiring the most productive agents, which contrasts with the signaling games where all agents are hired and the firm’s objective is to decide how much salary to offer to each agent.

On the firm side, our model setup bears more similarity to the strategic classification problems – offering a fixed wage, the firm decides whether to hire the agents based on their signal (Kleinberg and Raghavan, 2019, Frankel and Kartik, 2019, Bonatti and Cisternas, 2019). In other words, the firm is trying to classify agents based on whether their expected productivity exceeds the wage or not. This classification setting makes it possible for us to analyze the economic impact of algorithmic transparency on the decision maker (firm) by comparing its equilibrium payoff in the opaque and transparent scenarios. Moreover, we do not assume the signal (education level for example) to always be pure money-burning. Instead, we allow the causal feature to positively impact productivity and specify conditions regarding this positive effect under which algorithmic transparency benefits the decision maker. In that aspect, our work is also closely related to the strategic classification literature that assumes the existence of both causal and non-causal features.

Signaling Games. The signaling game literature studies how agents strategically reveal their type to a principle in a situation of information asymmetry. Traditional signaling models typically assume that costly actions are the only channels through which agents can signal their type (Spence, 1973). In these models, standard assumptions such as the Spence-Mirrlees single-crossing condition ensure the existence of separating equilibria: equilibria that fully reveal agents’ private information. While the machine learning models are trying to solve the same problem (i.e., trying to identify the type of agents under information asymmetry), they differ from decision makers in the classical signaling models in the following way. A machine learning model uses multiple features to learn an agent’s type. Each feature is essentially an action taken by the agent that signals her type. Some of these features are costly to improve, while others are not.

Our paper is related to recent signaling game papers that have also considered multiple actions as channels through which an agent can signal her type (Alós-Ferrer and Prat, 2012, Daley and Green, 2014, Engers, 1987, Frankel and Kartik, 2019). In these papers, the agents are always aware of the actions that are used as signals by the decision maker. In contrast, in our model’s opaque scenario, the agents know that a correlational feature is being used by the firm, but they do not know exactly what feature that is.

In our model, agents can use causal and/or correlational features to signal their type. The causal feature is similar to the costly signal typically captured in the traditional signaling game literature. A key difference is that we allow it to not only act as a signal of an agent’s type but also have an impact on the agent’s productivity, similar to Weiss (1983). For example, if agents of the same type have different levels of education, in our model, the firm would receive different payoffs from hiring them. The correlational feature that we model bears some similarities to the information in cheap talk games (Crawford and Sobel, 1982). This feature is almost costless to share, and it affects the eventual payoff of both the firm and the agents where their incentives are not perfectly aligned. In cheap talk games, the agent strategically manipulates this information, whereas, in our model, the agent does not know about this feature and cannot manipulate it in the opaque scenario.

Similar to our paper, a few recent papers have modeled the tradeoffs that an agent faces in the presence of multiple signals. For example, Daley and Green (2014) modeled a scenario where a student can send a costly signal (e.g., joint degree completion) to the recruiter or rely on a type-correlated noisy signal (e.g., grades). They characterized the results based on the informativeness of the noisy signal. The noisy signal is similar to the correlated feature in our model, and its informativeness is also modeled similarly to how we capture the predictive power of the correlated feature. A key finding of the paper is that, when grades are informative, H type individuals are less eager to send costly signals because they can now rely on grades to signal their type, while L type individuals are more willing to send a costly signal to de-emphasize the ‘grades’ dimension. Consequently, a separating equilibrium on the

costly signal dimension is harder to sustain. Our model also shares some similarities with this paper in that we also consider the possibility that the existence of an extra noisy signal will change individuals’ decision on the more costly signal. However, there are key differences in our model’s assumptions and results. Unlike the ‘grades’ dimension, whose value is impossible to manipulate, in our model, the firm can give individuals the opportunity to game the correlational feature by making the algorithm transparent. On the one hand, manipulation will make this dimension less informative. On the other hand, individuals’ behavior on the more costly signaling dimension will also change and could lead to a separating equilibrium in many cases.

Strategic Classification. Strategic classification literature considers the problem of designing optimal classification algorithms when facing strategic users who may manipulate the input to the system at a cost (Hardt et al., 2016). Canonical strategic classification models deem that the user’s manipulation always hurts the decision maker. Guided by this belief, a large stream of research on strategic classification is focused on developing algorithms that are robust to gaming (Meir et al., 2012, Cummings et al., 2015). Recently, several papers have argued that this gaming itself can be beneficial to the decision maker; thus, instead of focusing on manipulation-proof algorithms, these papers focus on designing algorithms that incentivize individuals to invest in desirable features (Kleinberg and Raghavan, 2019, Alon et al., 2020, Haghtalab et al., 2020). These papers are the ones we want to highlight since our paper also points out the difference between ‘gaming’ and ‘improvement’: gaming is bad for the decision maker because it deteriorates the information contained in the relevant features, but ‘improvement’ could be beneficial to the decision maker since it will causally impact the target variable.

Kleinberg and Raghavan (2019) studied the principle-agent problem where the agents’ features (e.g., final exam score) can be improved in two ways: by investing effort in a desirable way (e.g., spending time on course material) or by investing effort in an undesirable way (e.g., cheating). The effectiveness of each kind of effort on the feature is called the *effort profile*.

The decision maker can observe the agents’ performance on the features but cannot observe in which way the agents achieve their scores. Alon et al. (2020) examined a similar setting but extended Kleinberg and Raghavan (2019)’s model into a multi-agent scenario. Instead of assuming all individuals share the same *effort profile*, they focused on designing optimal algorithms that can work for different groups of individuals who may have different effort profiles. Our work is different from theirs with respect to one important aspect: while they assumed that a feature can be improved in either a causal or non-causal way, our model assumes that there are pure causal features and pure correlational features. Causal features (e.g., education level) can be improved only in a ‘causal’ way (such that the value of the target variable will also increase). Correlational features (e.g., whether an applicant wears glasses or not) can be improved only by gaming. If there was a ‘causal’ way to improve the correlational feature, then the firm’s willingness to publish the correlational feature would be stronger and might come from the potential productivity-enhancing effect of individuals’ improvement on the correlational feature. We show that even in the case where gaming on the pure correlational feature has no positive effect on productivity, the firm may still want to publish it. Furthermore, none of the papers above has studied how firms choose between opaque and transparent algorithms.

Another two papers we want to highlight are Frankel and Kartik (2019) and Bonatti and Cisternas (2019). These two economics papers showed that the decision maker could be ex-ante better off by committing to some ex-post sub-optimal strategies such as down-weighting some relevant features. Our paper is related to these papers in the sense that ‘publishing the algorithms’ could also be seen as a way to down-weight the correlational feature. However, there are critical differences between these papers and ours in terms of mechanisms. In their papers, ex-post sub-optimal behavior such as ‘under-utilizing’ some informative features might be preferred by the decision maker because individuals may have less incentive to manipulate these features if they anticipate that the features will be down-weighted. The decision maker loses some predictive accuracy due to under-utilizing those features, but the

observed values for those features now better represent individuals' natural behavior instead of gaming behavior. Consequently, not fully exploiting the information contained in relevant features might be optimal ex-ante. Our paper, however, shows that even if such 'feature down-weighting' cannot reduce individuals' gaming, it may still benefit the decision maker. In our paper, the purpose of 'down-weighting' the correlational feature is not to reduce or eliminate gaming behavior, but rather to increase the competition intensity regarding the causal features on which the H type agents hold a cost advantage.

1.3 Model

In this section, we develop a parsimonious model that captures how the firm and agents act under opaque and transparent scenarios. We consider the hiring setting discussed above with two types of agents: high-talent (H type) agents and low-talent (L type) agents. For simplicity, we normalize the total number of agents to 1 and assume that a θ portion of them are of H type and the remaining $1 - \theta$ portion are of L type.

Talent level is directly related to job performance and, ideally, the firm would like to hire only H type agents. However, the firm cannot directly observe an agent's type until she is hired and works at the firm for a while. Consequently, the firm can only use some observable agent features to differentiate the two types of agents. We classify these features into two types: causal features and correlational features. For simplicity, we assume that the firm only uses one causal feature (which is common knowledge to the firm and the agents, e.g., education level) and one correlational feature (which is unknown to the agents unless the firm decides to reveal it). Both features take on a discrete value of 0 (low) or 1 (high). Each agent can be characterized by one of four possible combinations (or states) in the two-dimensional feature space:

- State A (low causal, high correlational);
- State B (high causal, high correlational);
- State C (low causal, low correlational); and

- State D (high causal, low correlational).

The firm’s hiring strategy can therefore be represented by four hiring probabilities for the four states. In the remainder of this paper, we will refer to these probabilities as P_A , P_B , P_C , and P_D , respectively.

We assume that it is costly for the agents to improve the common-knowledge causal feature, and that H type agents have a cost advantage on this feature. Specifically, we assume that C_H (the cost of improving the causal feature for the H type agents) is smaller than C_L (the cost of improving the causal feature for the L type agents). In contrast, the cost to improve the correlational feature is assumed to be the same for the H type and L type agents and is very small (i.e., marginally above zero). It is worth noting that, although the cost of improving any given correlational feature is small, there are many of them, and agents do not know which correlational feature will be used in the algorithm unless the firm decides to reveal it.

To model the situation where the firm has an incentive to include the correlational feature into its decision making, we further assume that a λ portion of the H type agents and a $1 - \lambda$ portion of the L type agents have value 1 on the firm’s chosen correlational feature. Moreover, $\lambda \in [0.5, 1]$, which indicates a positive correlation between an agent’s value on the correlational feature and her type. Such a correlation can be viewed as a result of H and L type agents working to improve their correlational features. That is, even when agents do not know which correlational feature the firm is using, they may attempt to manipulate as many correlational features as possible because they incur little cost to do so. $\lambda (1 - \lambda)$ can then be interpreted as the probability of H (L) type agents hitting the feature that is actually used by the firm. The different probabilities of hitting the right correlational feature are another key difference between the two types of agents. The true value of λ is known to the firm. We assume that the agents also have a correct belief about λ in the main model. This assumption will be relaxed in Appendix A.2.3 where we allow the agents to have an ‘incorrect belief’ about λ .

We will study two scenarios, an opaque scenario and a transparent scenario. In order to focus on the more interesting and realistic cases, we make the following assumptions regarding the strategy sets of the agents:

1. *In the opaque scenario, the agents will only focus on whether to improve their causal feature.* This assumption is motivated by the fact that, in reality, while causal features are usually common knowledge between the firm and the agents (e.g., everyone knows that education level plays an important role in hiring decisions), correlational features are less so. In the opaque scenario, the firm does not reveal which correlational feature it will use in its algorithm to the agents. Consequently, the best that an agent can do is to make a guess, and that process leads to a $\lambda (1 - \lambda)$ portion of the H (L) type agents having value 1 on the firm's chosen correlational feature, as explained above. From a theoretical perspective, we note that this assumption is without loss of generality and is useful to simplify our analysis.
2. *In the transparent scenario, all agents will improve their correlational features.*⁵ Once the firm reveals the correlational feature that it will use in its algorithm, the probability that an individual agent hits the right feature becomes 1. Since the cost of improving the correlational feature is minimal, as long as it increases an agent's probability of being hired, they will improve this feature. It is worth noting that assuming all agents will improve the correlational feature used in the algorithm does not cause a loss of generality. This is so because, under the scenario where all agents achieve a "high" state on the disclosed correlational feature, this feature completely loses its predictive power and, therefore, will drop out of the prediction algorithm. If it can be shown that the firm can still be better off by making its algorithm transparent in such an extreme case of "agent gaming," it sends a strong message that algorithmic transparency can indeed be economically beneficial. This assumption will be relaxed in Appendix A.2.1 where

⁵It's worth noting that although the correlational feature will be dropped if everyone improves it, the level of transparency actually increases. The reason is that agents know exactly the state they are in and thus are more certain about their chances of being hired.

we allow gaming cost of the correlational feature to be significantly greater than zero. Our main result is strengthened after the relaxation of this assumption.

1.3.1 The Firm's Utility

As previously mentioned, the causal feature (for consistency, hereafter, we will use education level as an example of a causal feature) has a direct influence on the agents' performance, while the correlational feature does not. Thus, an agent's performance is determined by both her type ($T \in \{H, L\}$) and her education level ($Education \in \{0, 1\}$). Here, we follow Weiss (1983) and allow education to not only act as a signal of an agent's type but also contribute to the productivity of the agent. We use α and β to denote the marginal effects of type and education, respectively. For convenience, we normalize the performance of an uneducated L type agent to 0. The mathematical expression of an agent's performance is given by:

$$W(T, Education) = \alpha \times 1(T = H) + \beta \times 1(Education = 1). \quad (1.1)$$

We can also express an agent's performance as a function of both her type and her state at the end of the game. For example, in the opaque scenario, we can write an agent's performance as a function of $T \in \{H, L\}$ and $S \in \{A, B, C, D\}$ as follows:

$$W_S^T = \alpha \times 1(T = H) + \beta \times (1(S = B) + 1(S = D)). \quad (1.2)$$

In the transparent scenario, since all agents are "high" on the correlational feature, they only differ on the causal feature (i.e., education). This means that we can reduce the number of possible states from four to two, i.e., state E (low education) and state F (high education). Using these state definitions, we can write an agent's performance as follows:

$$W_S^T = \alpha \times 1(T = H) + \beta \times 1(S = F). \quad (1.3)$$

Once the agents are hired, their performance will contribute to the firm’s payoff, and the firm will pay them a fixed reward R (i.e., job compensation). In the main model, we assume that the reward R is the same for both transparent and opaque scenarios, and that its value is exogenously given. In Appendix A.2.2, we will endogenize R and allow the firm to potentially use different rewards for the transparent and opaque scenarios. We will show that our main insights still hold.

Let n_S^T denote the number of T type agents whose final states are S , and let $n_S = n_S^H + n_S^L$ denote the total number of agents whose final states are S . n_S^H and n_S^L are determined by agents’ strategy on the causal feature, q_H and q_L , which will be defined in Section 1.3.2. Furthermore, let $\gamma_S^e = n_S^H/n_S$. The firm’s expected total payoff (for narrative convenience, we use ‘payoff’ to refer to ‘expected payoff’ hereafter) under hiring strategies (or probabilities) $P = (P_A, P_B, P_C, P_D)$ in the opaque scenario, or $P = (P_E, P_F)$ in the transparent scenario, can be mathematically expressed as

$$\begin{aligned} \Pi_{firm}(P, q_H, q_L) &= \sum_S P_S \cdot [n_S^H(W_S^H - R) + n_S^L(W_S^L - R)] \\ &= \sum_S P_S \cdot n_S \cdot [\gamma_S^e(W_S^H - R) + (1 - \gamma_S^e)(W_S^L - R)]. \end{aligned} \quad (1.4)$$

1.3.2 The Agents’ Utility

In the opaque scenario, the agents do not know which correlational feature will be used by the firm’s algorithm; therefore, they will only focus their decisions on whether to improve the causal feature (i.e., education). Since agents do not know their values on the correlational feature, agents of the same type have same information and use the same strategy. Let u_T denote the expected utility (for narrative convenience, we use ‘utility’ to refer to ‘expected utility’ hereafter) of a T type agent. u_T can be expressed as a function of the firm’s hiring strategy P and the agent’s decision on the probability of improving the causal feature

$q_T \in [0, 1]$.⁶

In the opaque scenario, we have:

$$u_H(P, q_H) = q_H \cdot (\lambda P_B R + (1 - \lambda) P_D R - C_H) + (1 - q_H) \cdot (\lambda P_A R + (1 - \lambda) P_C R) \quad (1.5)$$

$$u_L(P, q_L) = q_L \cdot ((1 - \lambda) P_B R + \lambda P_D R - C_L) + (1 - q_L) \cdot ((1 - \lambda) P_A R + \lambda P_C R) \quad (1.6)$$

In the transparent scenario, all agents will have “high” values on the correlational feature, and their decisions are the probabilities on improving the causal feature. The utility of a T type agent in the transparent scenario is:

$$u_H(P, q_H) = q_H \cdot (P_F R - C_H) + (1 - q_H) \cdot P_E R \quad (1.7)$$

$$u_L(P, q_L) = q_L \cdot (P_F R - C_L) + (1 - q_L) \cdot P_E R \quad (1.8)$$

1.3.3 Game Sequence and Equilibrium Concept

The game between the firm and the agents is played as follows. In the first stage, the firm makes a decision on transparency (i.e., opaque or transparent), and this decision is known to all agents. If the firm chooses “opaque,” the remainder of the game proceeds as follows: The agents first choose their strategies to improve their features, and then the firm makes its hiring decisions based on the observed agent features. If, on the other hand, the firm chooses “transparent,” the remainder of the game proceeds as follows: The firm first discloses the algorithm to the agents; next, the agents choose their strategies; and, finally, the firm decides whom to hire based on the observed agent features. Recall from Section 3.1 that, in the main model, we focus our analysis on the case where the firm either does not have commitment power, or reveals only the features it uses but not its hiring probabilities.⁷ We

⁶We allow the strategies to be either pure or mixed so q_T could be any value between 0 and 1

⁷In Section 1.5, we will discuss the case where the firm has commitment power and reveals both the

use perfect Bayesian equilibrium (PBE) as the equilibrium concept. The PBE is characterized by (1) a set of type-specific probability distributions $q_T^*(P) \in \arg \max_{q_T} u_T(P, q_T)$ describing the optimal (possibly mixed) strategy in improving the causal feature for each type agents given the firm uses a hiring strategy P , (2) a set of state-specific probability distribution $P^* \in \arg \max_P \Pi_{firm}(P, q_T)$, describing the firm's optimal hiring strategy, and (3) a belief system $\mu(T|S)$ describing the firm's belief of each agent's type given her state calculated via Bayes rule whenever applicable.

1.3.4 Additional Parametric Assumptions

In Section 1.4, we will solve the game using backward induction. For each combination of (C_H, C_L) , we will derive the payoff for the firm in both the opaque and transparent scenarios. We will then specify the range of values of the parameters under which the firm is better or worse off when choosing to be transparent instead of opaque. We will show our results in the C_H - C_L space.

We make the following three additional assumptions regarding the relationships among the parameters to allow us to focus on non-trivial and more interesting cases.

$$0 < \beta < R < \alpha.$$

Assumption 1 says that the performance of an individual H type agent always exceeds the salary R regardless of her education level, whereas the performance of an individual L type agent is always smaller than R . This condition ensures that the firm only wants to hire the H type agents.

$$\frac{(\theta\lambda + (1 - \theta)(1 - \lambda))R}{\theta\lambda} < \alpha < \frac{R}{\theta}.$$

Recall that α denotes the performance advantage of H type agents over L type agents.

features it uses and its hiring strategy.

Assumption 2 ensures that α falls in a certain range that guarantees that the firm will have an incentive to include the correlational feature in its algorithm even when all agents' education levels are 0. The derivations of the lower bound and the upper bound of α can be found in Appendix A.3.1. Intuitively, if α is too small, the firm will not hire anyone when all agents' education levels are 0. If, on the other hand, α is too large, the firm will hire everyone even when all agents' education levels are 0. In either case, the correlational feature is useless for the firm. We refer to the lower (upper) bound of α as α ($\bar{\alpha}$) hereafter.

$$R - \theta\alpha < \beta < R - \frac{\theta(1 - \lambda)\alpha}{\theta(1 - \lambda) + (1 - \theta)\lambda}.$$

Assumption 1.3.4 says that the marginal effect of education on performance (β) falls in a certain range that guarantees the firm will have an incentive to include the correlational feature in the hiring algorithm. The derivations and interpretations of the lower and upper bound of β can be found in Appendix A.3.1. Intuitively, if β is too small, the firm will not hire any agent even when everyone has a high level of education in the transparent scenario, which leads to trivial results. If, on the other hand, β is too large, the firm will hire everyone with a high level of education irrespective of their values on the correlational feature in the opaque scenario, which could not justify the firm's incentive to use the correlational feature in the first place. We refer to the lower (upper) bound of β as β ($\bar{\beta}$) hereafter.⁸

A summary of notations can be found in Table B.1 in Appendix A.1.

1.4 Analysis

Let γ_S^b denote the proportion that H type agents represent among all the agents in state S at the beginning of the game.⁹ Per our discussions in Section 1.3, a θ portion of agents are of H type and the remaining $1 - \theta$ portion are of L type. Moreover, a λ portion of the H type agents and a $1 - \lambda$ portion of the L type agents have value 1 on the firm's chosen

⁸We provide discussion on the relaxation of Assumption 1.3.4 in Appendix A.3.2.

⁹If there are no agents in state S , we manually set $\gamma_S^b = 0$. The same applies to γ_S^e .

correlational feature. Therefore,

$$\begin{aligned}\gamma_A^b &= \frac{\lambda\theta}{\lambda\theta + (1-\lambda)(1-\theta)}; \\ \gamma_C^b &= \frac{(1-\lambda)\theta}{(1-\lambda)\theta + \lambda(1-\theta)}; \\ \gamma_B^b &= \gamma_D^b = 0.\end{aligned}$$

Recall, per our definition in Section 1.3.1, γ_S^e denotes the proportion that H type agents represent among all the agents in state S at the end of the game.

1.4.1 Opaque Scenario

Per our discussions in Section 1.3, in the opaque scenario, agents move first and will only decide on whether to improve the causal feature (i.e., education). Given the agents' strategies, based on Equation 1.4, the firm will be indifferent between hiring and not hiring agents with a final state S if $\gamma_S^e(W_S^H - R) + (1 - \gamma_S^e)(W_S^L - R) = 0$, or equivalently,

$$\gamma_S^e = \frac{R - W_S^L}{W_S^H - W_S^L}. \quad (1.9)$$

By Equation 1.2, the above fraction equals $\frac{R}{\alpha}$ when $S \in \{A,C\}$ and equals $\frac{R-\beta}{\alpha}$ when $S \in \{B,D\}$. Let $\gamma_{th0} = \frac{R}{\alpha}$ and $\gamma_{th1} = \frac{R-\beta}{\alpha}$ (by Assumption 1, we have $0 < \gamma_{th1} < \gamma_{th0} < 1$). Quantities γ_{th0} and γ_{th1} are important in the analysis, especially in determining whether a certain outcome (i.e., a combination of the agents' strategies and the firm's strategy) can be sustained in the equilibrium.

There is a total of nine possible classes of outcomes for agents' strategies. Five of them have the potential to be sustained in an equilibrium but the other four do not. The first five cases are: case 1 (neither H type nor L type agents improve education, i.e., $q_H = q_L = 0$); case 2 (only H type agents improve education, i.e., $q_H = 1, q_L = 0$); case 3 (both H type and L type agents improve education, i.e., $q_H = q_L = 1$); case 4 (H type agents improve education

with some probability, and L type agents do not improve education, i.e., $q_H \in (0, 1)$, $q_L = 0$); and case 5 (H type agents improve education, and L type agents improve education with some probability, i.e., $q_H = 1$, $q_L \in (0, 1)$). Aside from the above five cases, there are four other cases: case 6 (only L type agents improve education, i.e., $q_H = 0$, $q_L = 1$); case 7 (L type agents improve education with some probability but H type agents do not improve education, i.e., $q_H = 0$, $q_L \in (0, 1)$); case 8 (L type agents improve education, but H type agents improve education only with some probability, i.e., $q_H \in (0, 1)$, $q_L = 1$); and case 9 (both H type and L type agents improve education with some probability, i.e., $q_H \in (0, 1)$, $q_L \in (0, 1)$). It is not difficult to see that cases 6 through 9 cannot be sustained in an equilibrium. For cases 6 through 8, the L type agents have a higher value on education than the H type agents, and the firm will have an incentive to set higher hiring probabilities in states A and C than in states B and D. However, under this hiring strategy, the L type agents will have no incentive to improve education in the first place. As for case 9, the fact that both the H type and L type agents are using mixed strategies indicates that they are both indifferent between improving and not improving education. Since the cost of improving education for the L type agents is greater than that for the H type agents, to compensate for this higher cost, the L type agents must have a higher chance of being hired by the firm than the H type agents in the equilibrium. However, the firm has no incentive to use such a hiring strategy. We conclude that although there are nine possible classes of outcomes for the agents' strategies, only five of them (cases 1 through 5) can potentially be equilibrium outcomes.

Out of the five feasible cases, the actual equilibrium strategies of the agents and the firm depend on the values of (C_H, C_L) . The following lemma summarizes the agents' equilibrium strategies for different values of (C_H, C_L) and the corresponding payoff for the firm. The proof can be found in Appendix A.4.1. Since we assume that the H type agents have a cost advantage to improve the causal feature (i.e., $C_H < C_L$), the region above the diagonal line in Figure 1 is infeasible.

Lemma 1 *The equilibrium outcome depends on the values of (C_H, C_L) , and this dependence*

is shown in Figure 1.1. The payoffs for the firm are given by

$$\begin{aligned}\Pi_{firm_{O1}} &= \lambda\theta\alpha - (\lambda\theta + (1 - \lambda)(1 - \theta))R \\ \Pi_{firm_{O2}} &= \theta(\alpha + \beta) - \theta R \\ \Pi_{firm_{O3}} &= \lambda\theta(\alpha + \beta) + (1 - \lambda)(1 - \theta)\beta - (\lambda\theta + (1 - \lambda)(1 - \theta))R \\ \Pi_{firm_{O4}} &= \theta(\alpha + \beta - R) \left(1 - \frac{R(1 - \theta)(1 - \lambda)}{(\alpha - R)\theta\lambda}\right) \\ \Pi_{firm_{O5}} &= \frac{2\lambda - 1}{\lambda}\theta(\alpha + \beta - R),\end{aligned}$$

where $\Pi_{firm_{O_i}}$ denotes the firm's total payoff in case i .




Figure 1.1: Equilibrium outcome in the opaque scenario

Note: In Lemma 1, we only specify agents' equilibrium strategies and the firm's corresponding equilibrium payoffs. For a full description of the PBE, including the firm's equilibrium strategy and the firm's beliefs, please refer to Appendix A.4.1.

1.4.2 Transparent Scenario

In the transparent scenario, the firm moves first by announcing both the correlational feature that it uses and the probability of hiring for each state (i.e., P_E and P_F). Similar to the opaque scenario, there is a total of nine possible outcomes for the agents' strategies (we use

the same numbering of the nine cases as in the opaque scenario). To determine whether a certain outcome can be sustained as an equilibrium, we use the fact that the firm is indifferent between hiring and not hiring agents with final state S iff $\gamma_S^e = \frac{R-W_S^L}{W_S^H-W_S^L}$. The fraction on the right-hand side of the equality equals γ_{th0} when $S = E$ and equals γ_{th1} when $S = F$.

Consistent with the opaque scenario, cases 6 through 9 cannot be sustained in an equilibrium. Moreover, according to Assumption 1.3.4, if everyone is at state F, the firm will hire all agents. In case 5, some L type agents are at state E, which gives the firm even more incentive to hire agents in state F. However, again, to be an equilibrium, the mixed strategy outcome in case 5 requires the firm to be indifferent between hiring and not hiring agents from state F. Therefore, neither case 4 nor case 5 can be sustained in an equilibrium in the transparent scenario. Altogether, this leaves us with only the first three cases as possible equilibrium outcomes. The following lemma summarizes the agents' equilibrium strategies for different values of (C_H, C_L) , as well as the corresponding payoff for the firm. The proof can be found in Appendix A.4.2.

Lemma 2 *The equilibrium outcome depends on the values of (C_H, C_L) , and this dependence is shown in Figure 1.2. The payoffs for the firm are given by*

$$\begin{aligned}\Pi_{firm_{T1}} &= 0 \\ \Pi_{firm_{T2}} &= \theta(\alpha + \beta - R) \\ \Pi_{firm_{T3}} &= \theta(\alpha + \beta) + (1 - \theta)\beta - R,\end{aligned}$$

where $\Pi_{firm_{Ti}}$ denotes the firm's total payoff in case i .

Note: In Lemma 2, we only specify agents' equilibrium strategies and the firm's corresponding equilibrium payoffs. For a full description of the PBE, including the firm's equilibrium strategy and the firm's beliefs, please refer to Appendix A.4.2.




Figure 1.2: Equilibrium outcome in the transparent scenario

1.4.3 The Firm's Decision on Algorithmic Transparency

The firm can make a decision on algorithmic transparency by comparing the payoffs in the transparent and opaque scenarios. In this subsection, we will show how the firm is *not* always worse off making its algorithm transparent instead of opaque. We divide the blue region in Figures 1.1 and 1.2 into seven smaller regions: $N1$, $N2$, and $N3$ and $C1$, $C2$, $C3$, and $C4$ (see Figure 1.3).




Figure 1.3: Comparison of agents' equilibrium behavior in the transparent and opaque scenarios

We first consider what happens in regions $N1$ through $N3$. Note that, in these regions,

agents play the same equilibrium strategies on the causal feature in the opaque and transparent scenarios. For example, region $N1$ corresponds to case 1 in both the opaque and transparent scenarios, where neither type of agents improve their education. We now discuss the payoff comparison in these regions:

- In region $N1$, agents play the strategies in case 1 in both the opaque and transparent scenarios.¹⁰ Mathematically, $\Pi_{firm_{O1}} > \Pi_{firm_{T1}} = 0$. Therefore, the firm will always prefer to be opaque in this region.
- In region $N2$, agents play the strategies in case 2 in both the opaque and transparent scenarios. Since $\Pi_{firm_{O2}} = \Pi_{firm_{T2}}$, in this region, the firm is indifferent between being opaque or being transparent.
- In region $N3$, agents play the strategies in case 3 in both the opaque and transparent scenarios. We can rewrite $\Pi_{firm_{T3}}$ as follows:

$$\begin{aligned}\Pi_{firm_{T3}} &= \theta(\alpha + \beta)(\lambda + (1 - \lambda)) + (1 - \theta)\beta(\lambda + (1 - \lambda)) \\ &\quad - R(\theta\lambda + (1 - \theta)(1 - \lambda) + (1 - \theta)\lambda + (1 - \lambda)\theta).\end{aligned}$$

Since $\Pi_{firm_{O3}} = \lambda\theta(\alpha + \beta) + (1 - \lambda)(1 - \theta)\beta - (\lambda\theta + (1 - \lambda)(1 - \theta))R$, we have:

$$\begin{aligned}\Pi_{firm_{T3}} - \Pi_{firm_{O3}} &= (1 - \lambda)\theta(\alpha + \beta - R) - \lambda(1 - \theta)(R - \beta) \\ &= \theta(1 - \lambda)\alpha - (\theta(1 - \lambda) + (1 - \theta)\lambda)(R - \beta) < 0.\end{aligned}$$

where the inequality follows Assumption 1.3.4. This means that the firm will always prefer to be opaque in this region.

We conclude that, in regions $N1$ through $N3$, being transparent is never strictly better than being opaque. This is quite intuitive since, in these regions, agents play the same strategies on the causal feature in both the opaque and transparent scenarios. Hence, the

¹⁰It is worth mentioning that although agents' strategies on the causal feature are the same in the opaque and transparent scenarios, the firm's payoff is different because of the existence of the correlational feature.

firm can only be worse off revealing its algorithm due to the loss in the predictive power of the correlational feature. Specifically, in regions $N1$ and $N3$, when the firm chooses to be opaque, the predictive power of the algorithm only comes from the correlational feature because the H type and L type agents play the same strategy on the causal feature. This suggests that the firm will incur a significant loss due to the reduction in the algorithm's prediction accuracy when the algorithm is made transparent.

We now discuss the payoff comparison in regions $C1$ through $C4$. In what follows, we first provide a summary of the payoff comparison result in each region, and then discuss the intuition.

- In region $C1$, the agents' strategies and the firm's payoff change from opaque scenario 1 to transparent scenario 3. $\Pi_{firm_{T3}} > \Pi_{firm_{O1}}$ iff

$$\beta > \beta_1 = \lambda\theta(\alpha - R) - (1 - \lambda)(1 - \theta)R + R - \theta\alpha. \quad (1.10)$$

Thus, in this region, the firm will prefer to be transparent when $\beta > \beta_1$.

- In region $C2$, the agents' strategies and the firm's payoff change from opaque scenario 4 to transparent scenario 3. $\Pi_{firm_{T3}} > \Pi_{firm_{O4}}$ iff

$$\beta > \beta_2 = R - \frac{\alpha R(1 - \lambda)}{(\alpha - R)\lambda + R(1 - \lambda)}. \quad (1.11)$$

Thus, in this region, the firm will prefer to be transparent when $\beta > \beta_2$.

- In region $C3$, the agents' strategies and the firm's payoff change from opaque scenario 5 to transparent scenario 3. $\Pi_{firm_{T3}} > \Pi_{firm_{O5}}$ iff

$$\beta > \beta_3 = \frac{\lambda\theta\alpha - \theta\alpha - 2\lambda\theta R + \theta R + \lambda R}{\lambda - 2\theta\lambda + \theta}. \quad (1.12)$$

In this region, the firm will prefer to be transparent when $\beta > \beta_3$. It is interesting

to note, however, that β_3 equals $\bar{\beta}$ defined in Assumption 1.3.4 (see below for an explanation of why this is the case). Since the value of β cannot exceed $\bar{\beta}$ (according to Assumption 1.3.4), this means that the firm will never prefer to be transparent in this region.

- In region $C4$, the agents' strategies and the firm's payoff change from opaque scenario 1 to transparent scenario 2. $\Pi_{firm_{T2}} > \Pi_{firm_{O1}}$ regardless of β . Thus, in this region, the firm will always prefer to be transparent.

According to Assumption 2, the value of α falls in the following interval:

$$\frac{(\theta\lambda + (1 - \theta)(1 - \lambda))R}{\theta\lambda} < \alpha < \frac{R}{\theta}.$$

Within the above interval, β_1 and β_3 are decreasing in α , and β_2 is increasing in α . Moreover, $\beta_1 = \beta_2$ when $\alpha = \alpha$ and $\beta_2 = \beta_3$ when $\alpha = \bar{\alpha}$ (α and $\bar{\alpha}$ are defined in Assumption 2). Thus, we have $\beta_1 < \beta_2 < \beta_3$. To understand why we have increasing thresholds for β as we move from regions $C1$ to $C3$ and why there is no threshold for β in region $C4$, we must look at how the firm's decision to be transparent changes the agents' strategies in different regions. To facilitate our discussions, we first define the concept of "degree of separation." Suppose that there are n_{H0} H type agents and n_{L0} L type agents who do not improve education and n_{H1} H type agent and n_{L1} L type agents who improve education. We define the degree of separation (Dos) between the H type and L type agents as follows:

$$Dos = 1 - \frac{\min(n_{H0}, n_{L0}) + \min(n_{H1}, n_{L1})}{n_{H0} + n_{L0} + n_{H1} + n_{L1}}.$$

Note that, if either all agents improve education or no one improves education, then Dos reaches its minimum value: $Dos_{min} = \max(\theta, 1 - \theta)$. If all H type agents improve education and no L type agents improve education, then Dos reaches its maximum value: $Dos_{max} = 1$. If either H type or L type agents use a mixed strategy, the value of Dos is somewhere in

between. The key observation here is that, the higher the value of Dos , the easier it is for the firm to differentiate H type agents from L type agents using the causal feature.

We now discuss how the firm's decision to be transparent changes agents' strategies in regions $C1$ through $C4$. Note that transparency intensifies agents' competition on the causal feature and that this intensified competition has the following two effects.

1. *The degree of separation between H type and L type agents on their causal feature changes.* As an illustration, consider region $C4$. In this region, when the firm switches from being opaque to being transparent, agents' strategies also switch from opaque scenario 1 to transparent scenario 2. Under opaque scenario 1, neither the H type nor the L type agents improve education, whereas under transparent scenario 2, only the H type agents improve education. This means that the H type agents are now more separated from the L type agents on the causal feature (i.e., there is a higher degree of separation). Similarly, it can also be verified that, in regions $C2$ and $C3$, the two types of agents become less separated and, in region $C1$, there is no change in the degree of separation.
2. *Agents' average value on the causal feature becomes higher and their work performance increases (according to Equation 1.2).* To see this, consider region $C1$. In this region, when the firm switches from being opaque to being transparent, agents' strategies also switch from opaque scenario 1 to transparent scenario 3. Although the change in agents' strategies does not affect the degree of separation, since both types of agents improve education, the average level of education for both agent types increases. Similarly, in region $C2$, the average level of education and, thus, the performance level of both types of agents increase. In region $C3$, only the average performance of the L type agents increases, whereas in region $C4$, only the average performance of the H type agents increases. We can see that, in regions $C1$ through $C4$, the agents' overall average performance always increases when the firm switches from being opaque to being transparent.

Since the firm always loses useful information from the correlational feature that helps it differentiate between the two types of agents when switching from the opaque to the transparent algorithm, the firm's decision on algorithmic transparency will depend on whether the above two effects (i.e., the change in the degree of separation and the increase in the agents' average performance) can offset the negative effect of information lost on the correlational feature. In region $C1$, even though the degree of separation does not change, both the H type and L type agents improve education, and the firm can benefit from the increase in the average agents' performance. Whether this benefit offsets the negative effect of the information loss on the correlational feature depends on the value of β . If β is large enough (i.e., $\beta > \beta_1$), then being transparent is preferred over being opaque. In region $C2$, the degree of separation decreases as the agents' distribution on the causal feature changes from partial separation to pooling. However, the average performance of the H type and L type agents increases, which suggests that, if β is large enough, the firm can still be better off making the algorithm transparent. The condition on β in this case is stricter than in region $C1$ (i.e., $\beta > \beta_2 > \beta_1$). This is so because the marginal effect of education must now be large enough to offset not only the previously mentioned negative effect of the loss of information on the correlational feature, but also the worse degree of separation on the causal feature.

In region $C3$, the firm's decision to be transparent affects fewer agents compared to in region $C2$. Switching from the opaque to the transparent algorithm incentivizes all L type agents and some H type agents to improve education in region $C2$, but it only incentivizes some L type agents to improve education in region $C3$. Since fewer agents are affected by the firm's switching from the opaque to the transparent algorithm in region $C3$ compared with region $C2$, and since the increase in the agents' average performance is proportional to the number of agents being affected, a larger β is needed in region 3 to achieve the same level of average performance found in region $C2$. This is why $\beta_3 > \beta_2$.

To see why $\beta_3 = \bar{\beta}$ (as defined in Assumption 1.3.4), note that in region $C3$ all H type agents have already improved education in the opaque scenario, and, only some L type

agents will switch from not improving to improving education when the algorithm is made transparent. According to Assumption 1, the individual productivity of the L type agents cannot exceed R , which means that the firm will not benefit from hiring any L type agents even if their education levels are high. Consequently, making the algorithm transparent has a negative effect on the firm's payoff since the degree of separation on the causal feature decreases. This negative effect is partially offset by some L type agents' increased investment on the causal feature. The net effect is negative and its magnitude is decreasing in β . When β reaches $\bar{\beta}$, the net effect becomes 0. To see why, consider a range of β values that are very close to $\bar{\beta}$, according to Equation A.4 in Appendix A.4.1, nearly all L type agents have improved education in the opaque scenario; thus, the degree of separation is already near zero. Therefore, making the algorithm transparent has little negative effect on the degree of separation. Additionally, the positive effect of the increased investment on the causal feature is also close to zero. When $\beta = \bar{\beta}$, both effects are zero.

In region $C4$, the degree of separation increases when the firm switches from being opaque to being transparent. In fact, the agents' distribution on the causal feature changes from pooling to perfect separation. In other words, the firm now can perfectly separate the H type agents from the L type agents based on the causal feature alone without needing the correlational feature. This effect in itself is sufficient to offset the negative effect of information lost on the correlational feature. This is the reason why, in this region, the firm prefers to be transparent regardless of the value of β .

The following theorem summarizes our findings about the firm's decision on transparency:

Theorem 1 *In regions $N1$ through $N3$, being transparent is never strictly better than being opaque. In regions $C1$ through $C4$, depending on the value of β , the firm may prefer being transparent to being opaque. Specifically, in region $C1$, the firm will prefer to be transparent if $\beta > \beta_1$; in region $C2$, the firm will prefer to be transparent if $\beta > \beta_2$; in region $C3$, the firm will never prefer to be transparent; and, in region $C4$, the firm will prefer to be transparent regardless of the value of β .*

1.4.4 The Effect of the Predictive Power of the Correlational Feature (λ) and the Fraction of High-Talent Agents (θ)

We have shown that the firm will be strictly better off making the algorithm transparent if the (C_H, C_L) pair lands either in region $C4$, or in regions $C1$ or $C2$ with some additional conditions on β . In region $C4$, the transparent algorithm is preferred regardless of β , λ , and θ since removing the correlational feature changes the agents' behavior on the causal feature from pooling to perfect separation, and perfect separation is the case where the firm receives the maximum profit. In regions $C1$ and $C2$, the main force that makes the transparent algorithm preferable is the agents' increased average level on the causal feature. As long as β exceeds the threshold $\beta_1(\beta_2)$, the increased level on the causal feature will have a large enough positive effect on work performance to make the firm better off. We now examine how the thresholds on β changes when λ or θ changes.

Taking the derivate of the expressions of β_1 and β_2 in Equations 1.10 and 1.11 with respect to λ yields the following:

$$\frac{\partial \beta_1}{\partial \lambda} = -2\theta R + \theta \alpha + R. \quad (1.13)$$

$$\frac{\partial \beta_2}{\partial \lambda} = \frac{\alpha R(\alpha - R)}{(2\lambda R - R - \alpha \lambda)^2}. \quad (1.14)$$

Both of these derivatives are greater than 0 in the parameter ranges that we consider (i.e., those given by Assumptions 1, 2, and 1.3.4). This means that, within each region, as λ becomes larger, a higher value of β is needed to make the transparent algorithm preferable. This is because a larger λ implies that more information is contained in the correlational feature, so a higher causal effect is needed to offset the loss of information from the correlational feature under the transparent algorithm. However, the effect of λ on algorithmic transparency is not this straightforward since, apart from the equilibrium payoff, λ can also determine the kind of strategy combination that can be sustained as an equilibrium given a (C_H, C_L) pair (i.e.,

the regions' shapes in Figure 1.3 will change as λ changes). Consider the region just beneath the dividing line of regions $N2$ and $C4$: as λ increases, it changes from belonging to region $N2$ to belonging to $C4$. This means that the firm will prefer the opaque algorithm when faced with a small λ but prefer a transparent algorithm when faced with a large λ . Although it appears counter-intuitive, it can be explained as follows. When λ is small, making the algorithm transparent is not effective enough to change the agents' behavior on the causal feature. However, when λ is large, the agents' investment on the causal feature will increase drastically, and the firm is able to hire agents who are on average more productive under the transparent scenario than under the opaque scenario.

The following proposition summarizes our findings about how λ affects the firm's decision on transparency:

Proposition 1 *An increase in λ has the following effects on the firm's decision on transparency:*

1. *The area of regions $C1$, $C2$, and $C4$ increases, which means that the transparent algorithm is preferred under more (C_H, C_L) value pairs.*
2. *Within regions $C1$ and $C2$, the conditions on β to make the transparent algorithm preferred to the opaque algorithm become stricter (i.e. a larger β is needed).*

We now discuss the impact of θ on algorithmic transparency. Taking the derivative of the expressions of β_1 and β_2 in Equations 1.10 and 1.11 with respect to θ yields:

$$\frac{\partial \beta_1}{\partial \theta} = -2\lambda R + \lambda \alpha + R - \alpha. \quad (1.15)$$

$$\frac{\partial \beta_2}{\partial \theta} = 0. \quad (1.16)$$

It can be shown that $\frac{\partial \beta_1}{\partial \theta}$ is smaller than 0 in the parameter ranges that we consider. This means that, in region $C1$, as the proportion of the H type agents increases, the conditions on β to make algorithmic transparency more desirable become milder. Since $\frac{\partial \beta_1}{\partial \theta} = 0$, in region $C2$, θ has no influence on the firm's decision on algorithmic transparency. The intuition

behind these results are as follows. There are two possible negative effects – the loss of the correlational feature and a reduced degree of separation on the causal feature – and one possible positive effect of algorithmic transparency – the increase in agents’ investment on the causal feature. In region $C1$, after the firm switches from the opaque to the transparent algorithm, neither the degree of separation on the causal feature (from pooling at 0 to pooling at 1) nor the increase in agents’ investment in the causal feature (all agents in the market, irrespective of their types, increase their causal feature from 0 to 1) changes with θ . The negative effect of the loss of the correlational feature is smaller when θ is high, because there are fewer L type agents in the market who can be mistakenly hired by the firm due to the loss of the correlational feature. Thus, a smaller β is needed to offset this negative effect. In region $C2$, a higher θ mitigates the two negative effects (again, there are fewer L type agents on the market, and the increase in the number of L type agents being mis-classified and hired by the firm due to the reduced degree of separation on the causal and the loss of the correlational feature will be smaller), but also reduces the positive effect (fewer agents increase their investment in the causal feature as the firm switches to the transparent algorithm). Overall, θ does not affect the value of β needed to make the transparent algorithm preferable.

The following proposition summarizes our findings about how θ affects the firm’s decision on transparency:

Proposition 2 *In region $C1$, a higher θ will increase the firm’s incentive to make the algorithm transparent. In other regions, θ has no impact on the firm’s decision on algorithmic transparency.*

1.4.5 Agents’ Welfare

In Section 1.4.3, we have specified conditions under which the transparent algorithm will yield a strictly higher payoff to the firm than the opaque algorithm. We will next investigate the impact of algorithmic transparency on the agents’ welfare.

The total payoff across all agents in the equilibrium is summarized in the following lemma.

Lemma 3 *For each equilibrium outcome shown in Figure 1.1 (for the opaque scenario) and Figure 1.2 (for the transparent scenario), the corresponding total payoff for all agents is given by*

$$\begin{aligned}
\Pi_{agents_{O1}} &= (\lambda\theta + (1 - \lambda)(1 - \theta))R \\
\Pi_{agents_{O2}} &= \theta(R - C_H) \\
\Pi_{agents_{O3}} &= (\lambda\theta + (1 - \lambda)(1 - \theta))R - C_H\theta - C_L(1 - \theta) \\
\Pi_{agents_{O4}} &= \frac{(\lambda\theta + (1 - \lambda)(1 - \theta))(R - C_H)}{\lambda} \\
\Pi_{agents_{O5}} &= \frac{(2R\lambda - R - C_H\lambda - C_L\lambda + C_L)\theta}{\lambda} \\
\Pi_{agents_{T1}} &= 0 \\
\Pi_{agents_{T2}} &= \theta(R - C_H) \\
\Pi_{agents_{T3}} &= R - C_H\theta - C_L(1 - \theta),
\end{aligned}$$

where $\Pi_{agents_{O_i}}$ denotes the agents' total payoff in case i of the opaque scenario and $\Pi_{agents_{T_i}}$ denotes the agents' total payoff in case i of the transparent scenario.

As previously discussed, Figure 1.3 shows how the agents' behavior changes on the causal feature when the algorithm is made transparent. First, we consider the three regions where the agents' behavior on the causal feature does not change (regions $N1$, $N2$, and $N3$). In directly comparing the agents' payoff in the equilibrium, we have the following observations:

- In region $N1$, the agents play the strategies in case 1 in both the opaque and transparent scenarios. $\Pi_{agents_{O1}} > \Pi_{agents_{T1}}$. Therefore, the agents will receive a higher total payoff under the opaque algorithm in this region.
- In region $N2$, the agents play the strategies in case 2 in both the opaque and transparent scenarios. Since $\Pi_{agents_{O2}} = \Pi_{agents_{T2}}$, in this region, the agents are indifferent to whether the algorithm is opaque or transparent.

- In region $N3$, the agents play the strategies in case 3 in both the opaque and transparent scenarios. $\Pi_{agents_{O3}} < \Pi_{agents_{T3}}$. Therefore, the agents will receive a higher payoff under the transparent algorithm in this region.

In regions $N1$, $N2$, and $N3$, the agents' behavior on the causal feature does not change. Since improving the correlational feature is assumed to be costless, the agents' total cost stays the same before and after the algorithm is made transparent. The only thing that varies is the benefit they can obtain under the firm's hiring strategy. In region $N1$, more agents will be hired under the opaque algorithm (a λ portion of the H type agents and a $1 - \lambda$ portion of the L type agents are hired under the opaque algorithm but no one will be hired under the transparent algorithm). In region $N3$, more agents will be hired under the transparent algorithm (a λ portion of the H type agents and a $1 - \lambda$ portion of the L type agents are hired under the opaque algorithm and everyone will be hired under the transparent algorithm). In region $N2$, the same number of agents will be hired regardless of whether the algorithm is opaque or transparent (only the H type agents will be hired).

Next, we consider the four regions where agents' behavior on the causal feature changes after the algorithm is made transparent (regions $C1$, $C2$, $C3$, and $C4$). By directly comparing the agents' payoffs in the equilibrium, we obtain the following observations:

- In region $C1$, the agents' strategies and their total payoff change from opaque scenario 1 to transparent scenario 3. $\Pi_{agents_{T3}} \leq \Pi_{agents_{O1}}$ iff

$$C_H\theta + C_L(1 - \theta) \geq (1 - \lambda\theta - (1 - \lambda)(1 - \theta))R. \quad (1.17)$$

The smallest possible value for the left-hand side (LHS) of the inequality is reached when a (C_H, C_L) pair lands at the lower left corner in region $C1$, or in other words, when $C_H = (1 - \lambda)R$ and $C_L = \lambda R$. It can further be shown that this smallest value equals the right-hand side (RHS). Thus, Equation 1.17 is satisfied for any (C_H, C_L) pair in region $C1$. In this region, the transparent algorithm will give the agents a lower total

payoff compared with the opaque algorithm.

- In region $C2$, the agents' strategies and their total payoff change from opaque scenario 4 to transparent scenario 3. $\Pi_{agents_{T3}} \geq \Pi_{agents_{O4}}$ iff

$$\frac{1-\lambda}{\lambda}C_H - C_L \geq \left(\frac{1}{\lambda} - 2\right)R. \quad (1.18)$$

The smallest possible value for the LHS of the inequality is reached when a (C_H, C_L) pair lands at the lower right corner in region $C2$, or in other words, when $C_H = (1-\lambda)R$ and $C_L = \lambda R$. It can further be shown that this smallest value equals the RHS. Thus, Equation 1.18 is satisfied for any (C_H, C_L) pair in region $C2$. In this region, the transparent algorithm will give agents a higher payoff compared with the opaque algorithm.

- In region $C3$, the agents' strategies and their total payoff change from opaque scenario 5 to transparent scenario 3. $\Pi_{agents_{T3}} \geq \Pi_{agents_{O5}}$ iff

$$\left(2\theta - \frac{\theta}{\lambda} - 1\right)C_L \geq \left(2\theta - \frac{\theta}{\lambda} - 1\right)R. \quad (1.19)$$

Since $2\theta - \frac{\theta}{\lambda} - 1 \leq 0$, and $C_L \leq R$ in region $C3$, Equation 1.19 is satisfied for any (C_H, C_L) pair in region $C3$. In this region, the transparent algorithm will give agents a higher payoff compared with the opaque algorithm.

- In region $C4$, the agents' strategies and their total payoff change from opaque scenario 1 to transparent scenario 2. $\Pi_{agents_{T2}} \leq \Pi_{agents_{O1}}$ iff

$$\theta(R - C_H) \leq (\lambda\theta + (1-\lambda)(1-\theta))R. \quad (1.20)$$

The largest possible value for the LHS is reached when a (C_H, C_L) pair lands at the lower bound of region $C4$, or in other words, when $C_H = (1-\lambda)R$. It can be shown

that this largest value is smaller than the RHS. Thus, Equation 1.20 is satisfied for any (C_H, C_L) pair in region $C4$. In this region, the transparent algorithm will give agents a lower payoff compared with the opaque algorithm.

The following theorem summarizes our findings regarding the agents' welfare under the opaque and transparent algorithms.

Theorem 2 *Whether the agents are better off under either the opaque or the transparent algorithm depends on the values of (C_H, C_L) , and this dependence is shown in Figure 1.4. In regions $N3$, $C2$ and $C3$, the agents' welfare is higher under the transparent algorithm. In regions $N1$, $C1$ and $C4$, the agents' welfare is higher under the opaque algorithm. In region $N2$, the agents' welfare is not affected by algorithmic transparency.*




Figure 1.4: Comparison of agents' welfare in the transparent and opaque scenarios

The intuition behind Theorem 2 can be explained as follows. The pattern here is that the agents as a whole will prefer the opaque algorithm when C_H and C_L are large. They will prefer the transparent algorithm when C_H and C_L are small, and will be indifferent when C_L is large but C_H is small. Making the algorithm transparent may force the agents to invest in the causal feature. Of course, they can also benefit from the increased investment in the causal feature (i.e., a higher chance of being hired). The cost of this investment increases with C_H and C_L , but the benefit does not vary with C_H and C_L . Consequently, the agents

will be worse (better) off under the transparent algorithm if C_H and C_L are large (small).

Comparing Theorem 1 with Theorem 2, we find that the firm's and agents' interests conflict in some regions. For example, in regions $N3$ and $C3$, the transparent algorithm will give the agents a higher payoff, but the firm prefers the opaque algorithm. In region $C4$, the opaque algorithm will give the agents a higher payoff but the firm prefers the transparent algorithm. In regions $C1$ and $C2$, whether there is a conflict of interest between the firm and the agents depends on the value of β .

In region $C2$, if the condition of $\beta > \beta_2$ is satisfied, both the firm and the agents would prefer the transparent algorithm over the opaque algorithm. As for the total welfare, it is not difficult to see that in region $N1$, the opaque algorithm will result in a higher total welfare than the transparent algorithm, because both the firm and the agents prefer the opaque algorithm in the region. In region $N2$, the total welfare is the same under the opaque and transparent algorithms. In region $N3$, the firm prefers the opaque algorithm while the agents prefer the transparent algorithm. It turns out that the transparent algorithm will result in a higher total payoff. In regions $C1$ - $C4$, the total welfare could be higher either under the opaque algorithm or under the transparent algorithm, depending on the values of β , C_H , and C_L .

1.5 The Stackelberg Model

Our analysis in Section 1.4 has focused on the setting where the firm does not have commitment power. As noted in Section 3.1, this is useful to highlight the insight that the firm could still prefer the transparent algorithm to the opaque algorithm even in the absence of the 'first mover advantage'. In this section, we consider the case of 'full transparency' where the firm publishes all details of the hiring algorithm, including the features being used and the hiring strategy, and the firm has commitment power on the published algorithm. In this case, the Stackelberg model would be a more appropriate model when analyzing the

transparent scenario. Our main objective in this section is to show that the key insights in the previous section will only be further strengthened when the firm switches from using partial transparency to using full transparency.

We start with discussing the firm's equilibrium payoffs. The exact payoffs for all cases are summarized in Lemma 4.

Lemma 4 *The equilibrium outcome depends on the values of (C_H, C_L) , and this dependence is shown in Figure 1.5. The corresponding total payoffs for the firm are given by*

$$\begin{aligned}\Pi_{firm_{FT1}} &= 0 \\ \Pi_{firm_{FT2}} &= \theta(\alpha + \beta - R) \\ \Pi_{firm_{FT3}} &= \theta(\alpha + \beta) + (1 - \theta)\beta - R \\ \Pi_{firm_{FTS}} &= \frac{C_L\theta(\alpha + \beta - R)}{R},\end{aligned}$$

where $\Pi_{firm_{FTi}}$ denotes the firm's total payoff in case i , $i \in \{1, 2, 3, S\}$.




Figure 1.5: Equilibrium outcome in the transparent scenario in the Stackelberg model

Compared with the payoffs discussed in Lemma 2, we can see that the firm always gets a *weakly* higher payoff in the full transparency scenario than in the partial transparency scenario. This is so because, in the full transparency scenario, the firm can always commit to

the equilibrium hiring strategy in the partial transparency scenario and achieve the same payoff as in the partial transparency scenario, which sets a lower bound on the firm's payoff.¹¹ Given that the firm's payoff is weakly higher in the full transparency scenario than in the partial transparency scenario, naturally, the conditions under which full transparency is preferred over opacity are less strict than the conditions under which partial transparency is preferred over opacity. In particular, it can be shown that there are conditions under which the firm prefers full transparency, but not partial transparency, over opacity. These conditions are summarized in Proposition 3.

Proposition 3 *The firm prefers full transparency but not partial transparency over opacity if any of the following conditions is satisfied:*

- Value pair $(\frac{C_L}{R}, \frac{C_H}{R})$ falls in region C1 and $\beta \leq \beta_1$;
- Value pair $(\frac{C_L}{R}, \frac{C_H}{R})$ falls in region C2 and $\beta \leq \beta_2$ and $C_L > C_L^{C2}$;
- Value pair $(\frac{C_L}{R}, \frac{C_H}{R})$ falls in region C3 and $C_L > C_L^{C3}$;
- Value pair $(\frac{C_L}{R}, \frac{C_H}{R})$ falls in region N3 and $C_L > C_L^{N3}$;

where $C_L^{C2} = R - \frac{R^2(1-\theta)(1-\lambda)}{(\alpha-R)\theta\lambda}$, $C_L^{C3} = \frac{(2\lambda-1)R\theta}{\lambda}$, $C_L^{N3} = R\lambda\theta + \frac{R(1-\lambda)(1-\theta)(\beta-R)}{\theta(\alpha+\beta-R)}$.

In Section 1.4.4, we discussed how the values of θ and λ affect the firm's preference on algorithmic transparency under partial transparency. Most of the results in Section 1.4.4 will not change qualitatively when full transparency is considered. The detailed discussions and analysis can be found in Appendix A.3.3. In what follows, we briefly discuss the key findings:

1. In general, when λ is large, the conditions under which the transparent algorithm is preferred over the opaque algorithm become stricter. However, there are also (C_H, C_L) pairs under which the firm switches from never preferring transparency to possibly preferring transparency (depending on the values of the other parameters) when λ gets

¹¹To be more specific, in the full transparency scenario, if the firm commits to the equilibrium hiring strategy in the partial transparency scenario, P^* , the agents' best response would be q_T^* . The resulting firm's payoff is the same as the firm's equilibrium payoff in the partial transparency scenario. If there exists a hiring strategy that can lead to a payoff greater than the firm's equilibrium payoff in the partial transparency scenario, the firm will choose that hiring strategy. Therefore, the firm's payoff is weakly higher in the full transparency scenario than in the partial transparency scenario.

larger. This finding is consistent with that under the partial transparency model.

2. In general, when θ increases, there will be more (C_H, C_L) pairs under which the firm will prefer the transparent algorithm over the opaque algorithm. This finding is unique under full transparency since, under partial transparency, θ does not affect the firm's decision on algorithmic transparency in most regions. This difference is due to the fact that in some regions, the firm can get a higher payoff under full transparency than under partial transparency, and the payoff difference is increasing in θ . In other words, when θ is large, the firm's payoff in the full transparency scenario will increase by a larger amount compared to the partial transparency scenario, which gives the firm more incentive to make the algorithm transparent.

Next, we discuss how agents' welfare is affected by full transparency. Lemma 5 summarizes the agents' equilibrium payoff in the full transparency scenario.

Lemma 5 *In the full transparency scenario, for each equilibrium outcome shown in Figure 1.5 the corresponding total payoff for agents is given by*

$$\begin{aligned}\Pi_{agents_{FT1}} &= 0 \\ \Pi_{agents_{FT2}} &= \theta(R - C_H) \\ \Pi_{agents_{FT3}} &= R - C_H\theta - C_L(1 - \theta) \\ \Pi_{agents_{FTS}} &= \theta(C_L - C_H),\end{aligned}$$

where $\Pi_{agents_{FTi}}$ denotes the agents' total payoff in case i of the full transparency scenario.

Comparing the agents' payoff in the full transparency scenario with that in the partial transparency scenario, we can see that agents' welfare becomes strictly lower in region S and stays the same in all other regions. The finding is interesting as it shows that, as the level of transparency in the algorithm increases (i.e., from partial to full transparency), the firm may become better off while the agents may become worse off. It strengthens our conclusion in Section 1.4.5 that, in most cases when the firm prefers transparency, the agents will be

worse off under the transparent algorithm. The results presented in this section are consistent with recent research in both economics and computer science, which shows that compared to the case where the decision maker does not have commitment power, under the Stackelberg competition where the decision maker has commitment power, the decision maker can receive a higher payoff but at a cost of decreasing agents' welfare (Frankel and Kartik, 2019, Milli et al., 2019).

1.6 Extensions

In this section, we analyze several model extensions that relax some of the assumptions in our main model described in Section 1.3. Due to space limitation, we focus our discussion on the impact of algorithmic transparency on the firm, and briefly report the key findings here. The detailed discussions and proofs can be found in Appendix A.2. Overall, the results show that relaxing those assumptions does not alter the main results and insights of the study.

First, we relax the assumption that the cost of improving the correlational feature is close to zero. We consider the case where the cost of improving the correlation feature is substantial: c_h for the H type agents and c_l for the L type agents where $c_l \geq c_h > 0$. We find that under this condition, the firm's equilibrium payoff in the transparent scenario weakly increases while the firm's equilibrium payoff in the opaque scenario is not affected. Thus the main result of the paper, under certain conditions, algorithmic transparency benefits the firm, is strengthened. Second, we address the fixed wage assumption by allowing the firm to strategically choose the wage. After solving this extended model, we find that the firm in the transparent scenario is able to set a lower wage than in the opaque scenario without needing to worry about worsening the degree of separation between the two types of agents on the causal feature. Since the agents' productivity (affected by α and β) is assumed to be unaffected by wage, the firm will benefit more from endogenizing the wage in the transparent scenario than in the opaque scenario; thus, the firm's preference for the transparent algorithm

will be strengthened. Lastly, we show that even when the agents severely underestimate or overestimate the prediction power of the correlational feature, λ , the regions in which the firm prefers the transparent algorithm will not vanish.

1.7 Conclusion

1.7.1 Summary of Results

In this paper, we studied how firm and agent welfare is affected by algorithmic transparency. We allowed the agents to be strategic such that they can invest in their causal and correlational features to increase their chances of being hired in response to the firm's algorithm. We also investigated how the predictive power of the correlational feature, the market composition in terms of the fraction of H type agents, and the impact of the causal feature on agent productivity affect the firm's decision to make their algorithm transparent or opaque.

As a first result, we identified a broad set of conditions under which the firm would be better off with algorithmic transparency than opacity. Our second result is that the agents may not always be better off under algorithmic transparency. Our third result is that, even when the correlational feature has high predictive power in the opaque scenario, the firm could still be better off making the algorithm transparent. Our final result is that, when the fraction of H type agents on the market is high, the firm would be better off by making its algorithm transparent. We also provided several extensions to our main model. After relaxing several assumptions and considering several model alternatives, we found that the main insights of the paper do not change.

1.7.2 Implications for Managers

Our paper shows that managers using machine learning models for decision making could be better off by making their algorithms transparent. Algorithmic transparency does not always

mean a loss of predictive power. In some cases, it can in fact lead to greater predictive power. In other cases, while it may reduce predictive power, it can still make managers better off by improving the desirability of the whole market. Our results are particularly promising for managers, as they are now facing growing calls to make their algorithms transparent.

We identified a set of conditions where managers should prefer algorithmic transparency. There are three factors that managers should consider: (a) access to a good set of causal features; (b) the predictive power of the correlational features; (c) the market composition in terms of the fraction of H type (desirable) agents.

We provide some guidance on what makes a good causal feature here. The causal feature serves two purposes: (a) a signaling purpose – H type agents have a cost advantage on this feature, and thus, it can help separate H type agents from L type agents; and (b) a human capital purpose – the feature itself contributes to the productivity of the agents. To serve the signaling purpose well, the identified causal feature should be neither too costly nor too cheap to improve. If it is too costly, no one will improve it. By contrast, if it is too cheap, everyone will improve it. In terms of the human capital purpose, the higher the feature’s impact on productivity, the better it is. Even when this causal feature is unable to completely separate the H type agents from the L type agents, if it is moderately costly and contributes to productivity, the firm could still be better off. Typically algorithm designers are not focused on causality or identifying causal features. Our results indicate that they should. The recent stream of research in computer science that examines causal inference in machine learning models bodes well for them in this regard.

The second factor that managers should consider is the predictive power of the correlated features. Intuitively, managers may think of keeping their algorithms opaque when the correlational features provide significant predictive power. Our results show that this thinking is incorrect. We show that firms are more likely to be better off by making their algorithms transparent when correlational features provide significant predictive power in the opaque counterpart. Though incorrectly, managers would be particularly concerned about making

the algorithms transparent when the marginal effect of the causal feature in separating the H type agents from the L type agents is small in the presence of correlational features in an opaque algorithm. However, they should realize that the effect of the causal feature is suppressed largely due to the strategic behavior of the agents when the algorithm is opaque. All agents, but more importantly, the H type agents, underinvest in the causal feature when they know the correlational feature can separate them from the L type agents: the higher the predictive power of the correlational feature, the lower the incentive of the agents to invest in the causal feature. When the algorithm is made transparent, in the new equilibrium, the agents have a higher incentive to invest in the causal feature, making the firm better off.

The third factor managers should consider is the market composition in terms of the fraction of the H type agents. If the causal feature is unable to separate the H type agents from the L type agents, algorithmic transparency could still benefit the firm if the cost of improving the causal feature is moderate and the market is composed of more H type agents.

Overall, our results suggest that managers should not view manipulation by agents as bad. Rather, they should embrace it and use algorithmic transparency as a lever for motivating agents to invest in more desirable actions.

Managers can infer what parameter range (especially the (C_H, C_L) combination) their specific context falls in, which affects whether algorithmic transparency improves the firm's payoff, from the nature of the job opening they try to fill and the causal features they choose to consider, as well as the historical data they have about the past hiring outcomes. For example, compared to holding a bachelor's degree, if the causal feature considered is holding a graduate degree, the cost of improving the causal feature would be higher for both types of agents, and the H type agents' cost advantage would likely be higher. In addition, agents' strategies observed in the historical data, which most likely corresponds to the opaque scenario, can help managers pin down which condition (region shown in Figure 1.1) they are facing and predict agents' strategies under the transparent algorithm. The value of θ can be estimated from the historical data, and the value of λ is given by the algorithm trained on

historical data under algorithmic opacity.

1.7.3 Implications for Public Policy

There are two key arguments typically put forth in support of algorithmic transparency. First, making algorithms transparent could highlight any hidden biases in algorithms and make them accountable (Citron and Pasquale, 2014a). Second, the users who are affected by an algorithm's decision making have the right to see which factors affect decisions made about them. Our paper has implications related to this second argument.

Recent legislation like the General Data Protection Regulation has afforded individuals a right to explanation under which firms have to provide an explanation regarding how a decision has been made by their algorithms (Goodman and Flaxman, 2017). Our results show that such a regulation may not improve consumer welfare. When agents know which features in an algorithm affect important decisions about them, they can improve those features. Consequently, algorithmic transparency is generally viewed as helping the agents at the cost of the firm. However, our results show that the agents may not benefit from algorithmic transparency as expected. When algorithms are opaque, they derive their accuracy from both causal and correlational features. Therefore, the H type agents do not need to invest in the costly causal feature to separate themselves from the L type agents. In the transparent scenarios, in many cases, the agents have to invest in the costly causal feature to achieve similar or even less separation with no change in their wage.

1.7.4 Generalizability of the Results

While our model focuses on the job hiring scenario which is an example of a screening problem, our model and results are generalizable to other screening problems. Specifically, problems comprising of two types of economic agents who have asymmetric information and who are attempting to engage in a transaction. The "screener" (i.e. agent with less information, e.g. the firm in our case) attempts to gain further insight or knowledge into the private

information of the other agents (i.e. hidden type of the agent in our case). ML/AI algorithms are proving useful in helping the firms “screen” the applicants better. Job hiring is only one example of a screening problem. Other examples include, insurance markets (e.g. car insurance, health insurance) or credit lending markets.

1.7.5 Future Research Directions

Firms have typically kept their algorithms opaque to protect them from gaming by agents. In this study, we show that this strategy may not be the best, and the firm could be better off by making its algorithm transparent in the presence of strategic users. However, there are three other reasons as to why firms may still not want to make their algorithm transparent – interpretability, privacy and competition. With access to big data and large computational power, machine learning models have become complicated to the extent that they are rendered uninterpretable. While recent research has made advances in developing interpretable machine learning models, Bertsimas et al. (2019) show that model interpretability comes at a cost of accuracy. As a result, when considering the issue of algorithmic transparency, it may be interesting to consider the tradeoff between interpretability and accuracy. Similarly, when a firm makes its algorithm transparent, this can lead to privacy concerns. Others may be able to infer information about agents when they are selected by a transparent algorithm. In these cases, algorithmic transparency may impose a privacy cost on agents. Future research can investigate algorithmic transparency in the presence of privacy concerns. Finally, it would be interesting to investigate whether and how algorithmic transparency may affect the intensity of competition among firms. We believe these are interesting avenues for future research.

Chapter 2

Algorithmic Lending, Competition, and Strategic Information Disclosure

2.1 Introduction

Financial lenders use proprietary machine learning (ML) algorithms to evaluate credit risk for unsecured financial products, but these algorithms are not disclosed to borrowers who are directly impacted by their outcomes (Citron and Pasquale, 2014b). Furthermore, the approval outcomes for a consumer may differ between lenders, with one lender approving the application while another may deny it (Shaffer, 1998, Carmichael, 2017). This disparity can result in significant uncertainty for consumers when determining which lender to approach in the unsecured financial product market. A recent survey conducted by Experian revealed that 40% of respondents reported being denied a credit card (Experian, 2020). Additionally, 69% of surveyed consumers “wished they could know in advance if their credit card application would be approved”.

Financial lenders can reduce borrower uncertainty by offering pre-approval tools to borrowers.¹ The utilization of algorithms for screening renders it a relatively low-cost measure

¹The use of pre-approval tools by lenders can offer potential borrowers a clearer understanding of their chances of approval, thereby reducing uncertainty. Nonetheless, the effectiveness of pre-approval tools is

to provide pre-approval outcomes to prospective borrowers. In this scenario, according to the unraveling theory, lenders should reveal pre-approval outcomes freely to borrowers as it would encourage the potential borrowers who would be eventually approved by the lender to apply. Nonetheless, there exists a significant discrepancy in the provision of pre-approval outcomes to borrowers by financial lenders. Our analysis of several lenders' websites shows a discrepancy in the availability of pre-approval tools for credit cards.² We found that lenders do not offer pre-approval tools for credit cards at all times, and only a few lenders provide such tools at any given time.

We investigate how competition among lenders affects their incentives to provide free pre-approval checking tools to borrowers. We answer why even when lenders are symmetric, asymmetric revealing of the pre-approval outcome is an equilibrium outcome. We further examine how the accuracy of the lenders' algorithms, the accuracy of the borrowers' beliefs on the approval outcomes, and the riskiness of the market affect the lenders' decisions to reveal the pre-approval outcomes. Finally, we investigate how a policy that mandates pre-approval provision affects borrower surplus.

Formally, we analyze a multi-stage game in which duopoly lenders use ML algorithms to approve or reject borrowers who apply for their unsecured financial products. We analyze the lenders' equilibrium behavior in revealing pre-approval outcomes. The borrowers are modeled as H type (non-defaulters) and L type (defaulters). The borrowers decide which lender to apply to or they can choose not to apply to either. The two lenders are symmetric in their algorithms' accuracy and offer financial products that are identical in terms of non-price features (e.g. loan amount, and credit limit, etc.). The lenders first simultaneously decide (1) whether to provide pre-approval checking tools to borrowers or not, and then simultaneously

subject to certain limitations. During the pre-approval process, borrowers grant the lender permission to conduct a soft inquiry on their credit report. While soft inquiries do not provide lenders with as comprehensive credit information as hard inquiries, they do not negatively impact a borrower's credit score. Although pre-approval cannot guarantee formal approval of a borrower's application, it greatly reduces borrowers' uncertainty. For simplicity, we assume that if the lender provides the pre-approval tool, the pre-approval outcome aligns with the final decision

²For example, on March 15, 2023, Citi Bank, Chase and PNC were not offering the pre-approval tools for their credit cards whereas American Express, Capital One and Discover Card were.

decide (2) the price/interest rate that they will charge for their financial product. We use the sub-game perfect Nash equilibrium (SPNE) as our solution concept. Our most important and non-intuitive finding is the asymmetric equilibrium in revelation even when the two lenders are symmetric. In particular, we find that when the lenders' algorithms are accurate, the SPNE is *asymmetric*. One lender chooses to reveal the pre-approval outcomes to borrowers, while the other lender chooses not to reveal them. Only when the algorithms' accuracy is sufficiently low, both lenders will reveal the pre-approval outcomes.

Non-revelation of pre-approval by lenders can cause borrowers to make inaccurate assumptions about their eligibility, leading to qualified borrowers choosing not to apply. If both lenders reveal their algorithms, they can correct the inaccurate assumptions of borrowers and benefit from market expansion. However, it may intensify competition on interest rates as borrowers can make informed decisions and choose the lender offering the most favorable terms. In contrast, if only one lender discloses its algorithm, that lender can benefit from obtaining borrowers who wrongly believe that the competing lender would reject them. In contrast to the symmetric revealing scenario, in the asymmetric revealing case, a considerable number of consumers who only expect the revealing lender to approve their application are of H type. This subset of consumers is particularly appealing to the revealing lender as it can extract a larger surplus from them without facing competition. The revealing lender must balance between extracting surplus from the *price-insensitive* borrower segment who believe that only the revealing lender will approve them and attracting the *price-sensitive* borrower segment that thinks both lenders will approve them when deciding on the interest rate. Consequently, the revealing lender charges a higher interest rate than it would have if it only focused on the borrower segment that thinks both lenders will approve them. This allows the revealing lender to extract greater surplus by raising interest rates. Consequently, the softening of competition allows the non-disclosing lender to also extract surplus from the latter borrower segment.

The findings of our study demonstrate a novel association between lenders' pre-approval

revealing strategies and the level of product differentiation that emerges endogenously as a result of lenders' revealing decisions. Specifically, our research shows that lenders can strategically use asymmetric revealing to differentiate themselves and mitigate competition. By asymmetrically revealing their algorithms, lenders create asymmetry in interest rates for products with identical non-price characteristics, which ultimately diversifies borrowers' preferences. In the asymmetric equilibrium, the lender who reveals always receives a higher equilibrium payoff than the non-revealing lender. Therefore, when conditions support asymmetric equilibrium, a lender should take the opportunity to reveal its algorithm first. We also provide discussions on how various modeling parameters, including accuracy of the lenders' algorithms, accuracy of borrowers' prior beliefs of approval, and the riskness of the market, affect the existence of the asymmetric revealing outcome.

From a societal welfare perspective, lenders' provision of accurate approval odds enhances credit market efficiency by helping borrowers avoid suboptimal application decisions. However, our additional analysis suggests that regulating lenders to provide pre-approval tools may have some drawbacks. Specifically, lenders have less incentive to invest in algorithmic accuracy under the "both reveal" case because the increased accuracy intensifies competition. Conversely, this is not the case in the asymmetric revealing equilibrium. Therefore, policymakers should be cautious when considering regulations on algorithmic transparency, as mandatory transparency may reduce lenders' incentive to invest in algorithmic screening technologies. In the long term, this could hinder the allocation of financial resources to more creditworthy borrowers. While borrower surplus is maximized when both lenders provide accurate approval odds given a fixed algorithm accuracy, this may not hold true when the accuracy of the algorithms is endogenously determined by lenders.

The rest of this paper is organized as follows: §2 provides relevant literature and explain our contributions to it, §3 introduces the general setup of our model, §4 contains the bulk of the analysis, §5 discusses an important extension to the main model, and §6 concludes.

2.2 Contributions to Literature

In this section, we provide a literature review and highlight how our study contributes to the existing literature. Our paper is closely related to previous research on strategic information provision under competition. The unraveling theory suggests that a seller with private information about their product will reveal all information when the cost of disclosure is negligible and credibility is high (Grossman, 1981, Milgrom, 1981). However, this theory contradicts what is observed in the real world. Previous studies have explored competition as a factor that may prevent full disclosure of information. These studies typically model firms as vertically and/or horizontally differentiated in quality, with borrowers uncertain about the quality of the products (Kuksov and Lin, 2010, Board, 2009, Levin et al., 2009, Gu and Xie, 2013, Hotz and Xiao, 2013, Guo and Zhao, 2009). Firms make a strategic decision on quality revelation or fit revelation based on whether such revelation would intensify or soften competition. Our study is related to this stream of research, but instead of product fit or product quality, it focuses on the strategic revelation of approval odds, which is important information for borrowers seeking a financial product. We show how competitive forces lead to an asymmetric equilibrium where one lender strategically reveals accurate approval odds while the other does not.

The finding of asymmetric information disclosure is not new in the literature. Nevertheless, our research distinguishes itself from existing studies through the novel context of our model, the underlying mechanisms we investigate, and the distinct implications we draw. Unlike prior papers, which mainly focus on the disclosure of product fit and quality in conventional goods, our analysis is tailored to financial products such as credit cards and unsecured personal loans. These products are unique due to an inherent approval phase, which introduces significant uncertainty for consumers. This aspect enables an exploration into the strategic disclosure of approval odds and its subsequent impact on market competition and consumer welfare. Moreover, the driving force behind our findings of asymmetric disclosure diverges significantly

from the mechanisms posited in earlier works. While prior literature finds that firms with asymmetric quality may choose different revealing strategies to signal themselves (Guo and Zhao, 2009), our study identifies new insights and a novel mechanism and demonstrates that the revelation of pre-approval outcomes creates endogenous differentiation even when firms are initially identical. Lastly, our research extends beyond the existing literature, which often examines how outcome disclosures are influenced by product quality (Gu and Xie, 2013), disclosure sequence (Guo and Zhao, 2009) search costs (Ghosh and Galbreth, 2013) and consumer risk/loss preferences (Zhang and Li, 2021). We study how such disclosure decisions are contingent upon critical factors such as the accuracy of lenders' predictive algorithms, accuracy of borrowers' prior belief on approval, and the inherent risk of the market. Our unique analytical framework thus offers fresh perspectives on the strategic pre-approval revelation in financial markets.

Our paper also contributes to the current literature on credit market competition by examining the potential of asymmetric information disclosure regarding pre-approval outcomes in reducing interest rate competition. While past research has predominantly focused on information asymmetry in lenders' evaluation of borrowers' creditworthiness (Hauswald and Marquez, 2003, 2006), our study emphasizes the opposite direction's asymmetry, where borrowers are uncertain about their approval odds. Previous studies have explored several information acquisition and exchange approaches that directly impact screening algorithm performance. In contrast, our investigation centers on lenders' disclosure of algorithmic decisions to borrowers, which may indirectly moderate competition intensity without altering the screening algorithm's effectiveness. Other related studies suggest that lenders can reduce competition through specialized screening (Hauswald and Marquez, 2006), coordination (Bouckaert and Degryse, 2004) or product differentiation (Villas-Boas and Schmidt-Mohr, 1999). In light of this, our research explores a novel approach to mitigating competition by examining asymmetric pre-approval disclosure.

Our research contributes to the existing literature on firms' product differentiation

strategies in the context of heterogeneous consumer preferences (Tirole and Jean, 1988, Hauser, 1988, Moorthy, 1988). Despite assuming identical products in our model, we observe that borrowers exhibit preference heterogeneity *ex ante* due to their belief that they have a greater chance of loan approval for one product over the other. Our findings demonstrate that asymmetric disclosure of the algorithmic decision-making process can strategically differentiate products and diversify borrower preferences. Consequently, asymmetric pre-approval revelation by symmetric lenders leads to differentiated products and attenuates competition.

2.3 Model

We consider a duopoly credit market with two competing lenders that sell financial products to borrowers. We start by describing borrowers and lenders, then we explain the sequence of decisions.

2.3.1 Borrowers

We model two types of borrowers, high-quality (non-defaulter) and low-quality (defaulter), denoted as H type and L type, respectively. We use θ to denote the portion of H type borrowers in the market.³ Borrowers are considering to apply for a financial product (e.g., an unsecured personal loan or a credit card). We assume that the amount of the loan is fixed, and we normalize it to 1. A borrower's utility from applying conditional on her type and the approval outcome is as follows:

$$U = \begin{cases} M_h - m - b & \text{if she is } H \text{ type and gets approved} \\ M_l - m - b & \text{if she is } L \text{ type and gets approved} \\ -m, & \text{if she gets rejected} \end{cases}$$

³We model borrowers as either defaulters or non-defaulters, but our model can be easily generalized to the

where $M_h(M_l)^4$ is the overall benefit that a $H(L)$ type borrower can get from the financial product (e.g., the monetary value and/or convenience brought by a credit card). b is the price of the financial product (e.g., the interest rate of the loan or the fee of the credit card)⁵, and $-m$ is the cost of applying (e.g., a hard inquiry on credit record and time/effort spent on the application process). The utility a borrower will receive if she does not apply is normalized to zero. We assume each borrower has a unit demand⁶, and to keep the model simple, we do not model borrowers' multi-stage application (i.e., get rejected by one lender then apply for the other lender)⁷. Here you need to explain more meaningfully, how you address the referee comment that the person may apply to multiple credit cards. Added an extension, problem solved. To facilitate discussion, hereafter, we will use unsecured personal loans as an example of the financial product.

2.3.2 Lenders

There are two symmetric lenders⁸ in the market, providing homogeneous loans to borrowers. That is, the loans have identical non-price features but may come with different interest rates. Each lender has its own screening algorithm. We assume that the algorithms of the two lenders are independent⁹ but have the same level of accuracy¹⁰, and the algorithms are equally accurate in predicting the positive and negative cases (i.e., the true positive rate equals the true negative rate). Mathematically, the accuracy of the lenders' algorithms is characterized by $P_b \in (\frac{1}{2}, 1]$ ¹¹: the algorithm predicts a H type applicant as H type with

case where the H type borrowers default with a probability of ϵ_1 while the L type default with a probability of $1 - \epsilon_2$, where $\epsilon_1 < 0.5$ and $\epsilon_2 < 0.5$.

⁴In many other models, M_h is assumed to be larger than M_l , since L type borrowers have to incorporate the dis-utility associated with the reputation loss once they default. However, this assumption is not necessary in our setting.

⁵Hereafter, we use the term 'price' and the term 'interest rate' interchangeably

⁶We can think of this as that the marginal benefit of getting a second financial product decreases sharply.

⁷In Appendix B.2.1, we allow for this multi-stage application in a model extension

⁸This symmetric lender assumption will be relaxed in the model extension, see Appendix B.2.2

⁹This independence assumption will be relaxed in the model extension, see Appendix B.2.4.

¹⁰This equal-accuracy assumption will be relaxed in the model extension, see Appendix B.2.3

¹¹Hereafter, we use subscript b or B to denote all the lender side parameters since we use 'Bank' as an example for the lender.

probability P_b , and a L type applicant as L type with probability P_b .

The value of a non-defaulter to a lender equals the price set by the lender: b , and the value of a defaulter to a lender equals the negative of the loan amount, which has been normalized to -1 . The utility that a lender receives if it approves n_H H type applicants and n_L L type applicants is given by

$$\Pi = n_H b - n_L \tag{2.1}$$

2.3.3 Borrowers' Belief

The borrowers' application behavior (whether to apply and which lender to apply to) depends on their beliefs about the approval chance. Such beliefs may not align with the truth: the lender's algorithms are opaque, instead of knowing whether they will be approved directly from the lenders' algorithm, borrowers can only infer the approval chance from various other channels such as online forums, past experiences, or communication with other borrowers. While positively correlated with actual outcomes, these beliefs are not entirely precise. We use a parameter P_c to capture this accuracy of borrowers' belief on approval outcomes¹²: specifically, among all the applicants who will be approved by a lender, only P_c portion of them believe they will be approved before applying, the rest of them ($1 - P_c$ portion) incorrectly believe they will be rejected. Symmetrically, among all the applicants who will be rejected by a lender, only P_c portion of them believe they will be rejected before applying, the rest of them ($1 - P_c$ portion) incorrectly believe they will be approved. You should explain here how this belief system aligns with people knowing their own type. E.g. A H type knows his type and that the bank is accurate with P_b . The H type should think that the probability of approval by the bank is P_b . However, you are saying that the H type will think that his probability of approval is P_c . There is something missing in logic here. Under this belief

¹²Hereafter, we use subscript c or C to denote all the borrower-side parameters

system, borrowers of both types are subject to the same probability of having incorrect beliefs about the approval outcomes. Because H type borrowers are more likely to be approved than L type borrowers, H type borrowers are more likely to believe they will be approved than L type borrowers. Providing the pre-approval tool will help reduce this incorrect belief: In our model, we assume if the lender provides the pre-approval tool, the borrowers will know the approval outcomes exactly before applying, in other words, borrowers' beliefs on approval will be consistent with the actual approval outcomes.

2.3.4 Game Sequence

The sequence of stages in the game between the two lenders and the borrowers is as follows.

- **STAGE 1: Pre-approval Revealing Stage** Each lender chooses whether or not to provide the pre-approval checking tool simultaneously. Each lender's decision is observed by the other lender and the borrowers.
- **STAGE 2: Pricing Stage** Each lender chooses an interest rate b ($b \in [0, \bar{b}]$)¹³ for its financial product.
- **STAGE 3: Applying Stage** Borrowers observe the interest rates set by both lenders and observe the pre-approval decision(s) from the lender(s) who provides the pre-approval tool. Borrowers choose which lender to apply to or choose not to apply to either.
- **STAGE 4: Approving Stage** The lenders receive applications and use their algorithms to approve or reject borrowers. The payoffs to the lenders and the borrowers are then determined.

It is worth highlighting that lenders lack the ability to monitor the outcomes of competing lenders. Consequently, the lender is unable to differentiate between borrowers who will receive approvals from the competing lender and those who will be declined by the competing lender.

¹³ \bar{b} is the maximum interest rate that can be set by the lenders, for example, required by the law.

We use sub-game perfect Nash equilibrium (SPNE) as our solution concept. We solve for the SPNE using backward induction. We proceed with the analysis in the next section. A summary of notations can be found in Appendix B.1.

2.3.5 Illustration

Before the detailed analysis of the equilibrium of the game. We show the key intuition through an illustration.

In the absence of the pre-approval tool, the mismatch between borrowers' beliefs and the actual approval outcome creates inefficiency in the application process – some qualified borrowers choose not to apply and some non-qualified borrowers apply but are rejected. Providing the pre-approval tool will reduce applicants' sub-optimal applying behavior. In the competitive scenario, it will also affect the competitive level of the market: If a borrower believes she will be approved by both lenders, the two lenders have to compete for this borrower, we say this borrower is in the lenders' *common segment*. By contrast, if a borrower believes she will only be approved by lender 1 but not lender 2, we say this borrower is in lender 1's *captive segment*. Intuitively, if borrowers in the common segment are attractive (larger size and higher portion of H type borrowers), competition will be intensified. If borrowers in the captive segment are of high quality, competition will be softened. The lenders' pre-approval revelation decision will determine the composition of the size and quality of lenders' common and captive segments, consequently affecting the intensity of the competition.

Figure 2.1 shows a visual representation of the composition of different segments under different revealing scenarios. We use circles of different shades of gray to represent the borrowers' quality loosely, the black dots represent borrowers who are qualified under both lenders' algorithms, they are very likely to be H types given that both lenders' algorithms predict them as H types independently. The gray dots represent borrowers who are qualified under only one of the two lenders' algorithms, these borrowers are less likely to be H types.




Figure 2.1: Illustration

The four regions in each sub-figure are divided based on borrowers' beliefs on the approval outcome, for example, the upper left segment represents borrowers who believe will be approved by both lenders, this is also the common segment for both lenders. The upper right segment represents borrowers who believe will only be approved by lender 1, this is lender 1's captive segment. The black dots represent borrowers who are qualified under both lenders' algorithms, the gray dots represent borrowers who are qualified under only one of the two lenders' algorithms, the white dots represent borrowers who are not qualified under either lenders' algorithms.

The white dots represent borrowers who are not qualified under either of the two lenders' algorithms. When both lenders reveal the pre-approval outcomes (as shown in figure 2.1a), the common segment becomes very attractive since it only contains borrowers who will be approved by both lenders (black dots). The competition is very intense in this case. When neither lender reveals the pre-approval outcomes (as shown in figure 2.1b), the common segment becomes less attractive since it mixes in many borrowers who will only be approved by one lender (gray dots) but incorrectly believe will be approved by both lenders. However, the captive segment becomes more attractive, since it mixes in some high-quality borrowers who will be approved by both lenders (black dots) but believe will only be approved by lender 1 or 2. Consequently, lenders have less incentive to compete for the common segment but instead focus on capturing the captive segment, the competition is softened. In the asymmetric revealing case (as shown in figure 2.1c), the captive segment is attractive for the revealing lender since it mixes in some high-quality borrowers who will be approved by both lenders (black dots) but believe will only be approved by the revealing lender. Differently,

the non-revealing lender's captive segment is less attractive: all borrowers who believe will be rejected by the revealing lender will indeed be rejected, given that the revealing lender reveals the pre-approval outcomes. The revealing lender has incentives to focus on captive segments while the non-revealing lender can focus on the common segment, thus the competition is softened.

2.4 Analysis

Now we move to the detailed analysis.

2.4.1 Additional Assumptions

To avoid invalid and uninteresting equilibrium, we focus on the following parameter ranges:

Assumption 1

$$\bar{b}\theta - (1 - \theta) \geq 0.$$

Assumption 1 says that granting loans to borrowers is *ex ante* efficient under interest rate \bar{b} . This assumption ensures that lenders can get non-negative equilibrium payoffs under mixed strategy settings in all sub-games.¹⁴

Assumption 2

$$M_h - m - \bar{b} > 0$$

$$M_l - m - \bar{b} > 0$$

The two conditions in Assumption 2 ensure that a borrower will not choose the outside option (i.e., not to apply for either lender) if she believes that she will be approved by at least one lender. This is to avoid uninteresting cases where \bar{b} is so high that no one ends up applying

¹⁴Similar assumptions have been made in several papers, for example, in (Hauswald and Marquez, 2003) and (Hauswald and Marquez, 2006), to ensure that lenders will not have the incentive to engage in ruinous competition to obtain negative profits in equilibrium.

regardless of the belief about the approval outcomes.

2.4.2 Analysis: Approving Stage

We begin our analysis from the last stage, the Approving Stage. The lenders' decisions in this stage are trivial: They make the approving decisions based on their specific screening algorithm's prediction. Specifically, they will approve (reject) borrowers who are predicted by their own algorithm as H (L) type.

2.4.3 Analysis: Applying Stage

After seeing the interest rates set by the two lenders, as well as the pre-approval decisions (if applicable), borrowers in this stage decide whether to apply and if so which lender to apply to. It is straightforward to see that borrowers will only consider the lender that they think will approve them. Specifically, (1) if a borrower believes that she will be rejected by both lenders, she will choose not to apply to either lender; (2) if a borrower believes that she will be approved by only one of the lenders, she will apply to that lender; (3) if a borrower believes that she will be approved by both lenders, she will apply to the lender who charges a lower interest rate.

2.4.4 Analysis: Pricing Stage

In this stage, the two lenders anticipate borrowers' applying strategies in Stage 3, and choose pricing strategies that depend on the revealing outcome of Stage 1. We begin the analysis by solving each sub-game that results from lenders' pre-approval revelation choices in Stage 1.

There are three possible outcomes (sub-games) in Stage 1:

- **R-R:** Both lenders reveal the pre-approval outcomes.
- **N-N:** Neither lender reveals the pre-approval outcomes.
- **R-N:** One lender reveals the pre-approval outcomes, while the other does not.

We analyze the cases (sub-games) R-R, N-N, and R-N in succession.

Common Segment vs Captive Segment

Borrowers make decisions to apply to a lender based on the expected predictions they anticipate from each lender. These expectations result in four distinct borrower segments: HH , HL , LH , and LL . The HH segment represents borrowers who believe they will be approved (predicted as H type) by both lenders, while the HL segment represents borrowers who believe they will be approved by Lender 1 but rejected by Lender 2. Similarly, the LH segment represents borrowers who believe they will be rejected by Lender 1 but approved by Lender 2, and the LL segment represents borrowers who believe they will be rejected by both lenders. One thing to keep in mind is that these segments are not created based on the actual approval decision, but based on the borrowers' belief of approval.

The distribution of H and L type borrowers across the four segments varies across the three sub-games: R-R, N-N, and R-N. Table 2.1 provides an illustration of the distribution of H and L type borrowers across the four segments in the R-R case. In this table, the top and bottom expressions in each cell indicate the number of H and L type borrowers, respectively, in the corresponding segment. To calculate the number of H type borrowers in a segment, we use the probability that each lender independently predicts the H type borrower correctly, denoted as P_b , and the fraction of H type borrowers in the population, denoted as θ . Specifically, $P_b^2\theta$ is the number of H type borrowers that fall in the HH segment. The number of H and L type borrowers in the other segments can be calculated similarly as reported in Table 2.1.

Given our assumption that $P_b > 1/2$, we observe that the HH segment mostly consists of H type borrowers while the LL segment mostly consists of L type borrowers. Consequently, lenders view borrowers in the HH segment as the least risky and hence most valuable, whereas those in the LL segment are perceived as the most risky and least valuable.

It follows that borrowers who believe to receive a classification of HH would be indifferent

		Prediction from Lender 1	
		H	L
Prediction from Lender 2	H	$P_b^2\theta$ $(1 - P_b)^2(1 - \theta)$	$P_b(1 - P_b)\theta$ $P_b(1 - P_b)(1 - \theta)$
	L	$P_b(1 - P_b)\theta$ $P_b(1 - P_b)(1 - \theta)$	$(1 - P_b)^2\theta$ $P_b^2(1 - \theta)$

Table 2.1: Common Segment vs Captive Segment for the R-R case

Note: From Lender 1's perspective, the upper left cell is the common segment while the lower left is its captive segment; From Lender 2's perspective, the upper left cell is the common segment while the upper right cell is its captive segment. The top line in each cell denotes the number of H type borrowers and the bottom line in each cell denotes the number of L type borrowers in the corresponding segment.

between two lenders charging the same interest rate. As a result, lenders would compete with each other in a Bertrand fashion to attract borrowers from this segment, which we refer to as the common segment, as the probability of such borrowers applying to either lender is non-zero. On the other hand, borrowers in the HL (LH) segment will exclusively apply to Lender 1 (2) under the belief that only Lender 1 (2) will approve their application. We refer to these segments as the captive segment for Lender 1 (2). Therefore, while lenders would compete to lure borrowers from the common segment, no competition would arise for borrowers in the captive segment.

R-R: Both reveal

In the R-R case, both lenders provide the pre-approval tools. A borrower knows exactly whether she will be approved by each lender or not. To make the discussion easier, we use the following notations: We refer to the predictions of a borrower(k)'s type by Lender 1's and Lender 2's algorithms as $Y_{k,b}^1$ and $Y_{k,b}^2$, respectively. In the R-R case, lenders' predictions are directly passed on to borrowers through the pre-approval tool. A is the set of all borrowers in the market. We define the following sets which correspond to borrowers in the segments

defined in Section 2.4.4:

$$A_{hh}^{rr} = \{k \in A | Y_{k,b}^1 = H, Y_{k,b}^2 = H\}$$

$$A_{hl}^{rr} = \{k \in A | Y_{k,b}^1 = H, Y_{k,b}^2 = L\}$$

$$A_{lh}^{rr} = \{k \in A | Y_{k,b}^1 = L, Y_{k,b}^2 = H\}$$

$$A_{ll}^{rr} = \{k \in A | Y_{k,b}^1 = L, Y_{k,b}^2 = L\}$$

That is, A_{hh}^{rr} is the set of borrowers who are predicted as H type by both lenders, A_{hl}^{rr} is the set of borrowers who are predicted as H type by Lender 1 and predicted as L type by Lender 2, etc. Let N_{xy}^{rr} denote the number of borrowers and V_{xy}^{rr} denote the fraction of H type borrowers in the set A_{xy}^{rr} where $x, y \in \{h, l\}$. The values of N_{xy}^{rr} and V_{xy}^{rr} can be calculated from Table 2.1. Since neither lender will approve borrowers in A_{ll}^{rr} , we omit the calculations of N_{ll}^{rr} and V_{ll}^{rr} below.

$$N_{hl}^{rr} = P_b(1 - P_b)\theta + P_b(1 - P_b)(1 - \theta)$$

$$V_{hl}^{rr} = \frac{P_b(1 - P_b)\theta}{N_{hl}^{rr}} = \theta$$

$$N_{hh}^{rr} = P_b^2\theta + (1 - P_b)^2(1 - \theta)$$

$$V_{hh}^{rr} = \frac{P_b^2\theta}{N_{hh}^{rr}} > \theta$$

The most profitable segment for two lenders is A_{hh}^{rr} , as their algorithms predict that the borrowers in this segment are of H type. This prediction increases the posterior probability of these borrowers being H type, and thus the lenders compete intensely for this segment. On the other hand, the profitability of A_{hl}^{rr} and A_{lh}^{rr} , the captive segments of the respective lenders, is lower, resulting in less competition. However, the profitability of A_{hh}^{rr} intensifies competition, while the profitability of A_{hl}^{rr} moderates it. It is worth noting that lenders do

not have access to their competitors' algorithms or predictions, and cannot differentiate between borrowers in the common and their own captive segment. This implies that lenders cannot treat borrowers in the common and their own captive segment differently (e.g., charge different interest rates).

The optimal interest rate, b , for the lenders is derived in this context. Lenders are indifferent to the borrowers in A_{hh}^{rr} if they offer the same interest rates. Borrowers in A_{hl}^{rr} and A_{lh}^{rr} only apply to their respective lender. Lenders set their interest rates simultaneously, and a pure strategy equilibrium does not exist due to the incentive to undercut the competitor's interest rate to attract borrowers in A_{hh}^{rr} . Therefore, a mixed strategy equilibrium is studied, where both lenders use the same mixed strategy of randomizing their b with a probability distribution characterized by a cumulative density function $F^{rr}(b)$.¹⁵ The equilibrium strategy is summarized in Lemma 6, and the proof can be found in Appendix B.3.1.

Lemma 6 *The CDF of the distribution of lender's equilibrium pricing (interest rate setting) strategy is shown as follows:*

$$F^{rr}(b) = \begin{cases} 0 & \text{if } b < \underline{b}^{rr} \\ \frac{bP_b\theta - P_b\theta - P_b(-\bar{b}P_b\theta + \bar{b}\theta - P_b\theta + P_b + \theta - 1) + P_b + \theta - 1}{bP_b^2\theta + P_b^2\theta - P_b^2 - 2P_b\theta + 2P_b + \theta - 1} & \text{if } \underline{b}^{rr} \leq b < \bar{b} \\ 1 & \text{if } b \geq \bar{b} \end{cases}$$

Where $\underline{b}^{rr} = \frac{P_b\theta(-\bar{b}P_b + \bar{b} - P_b + 2) + P_b(P_b - 2) - \theta + 1}{P_b\theta}$

Each lender's equilibrium payoff is:

$$\Pi_{rr} = P_b(-\bar{b}P_b\theta + \bar{b}\theta - P_b\theta + P_b + \theta - 1)$$

An interesting observation from Lemma 6 is that lenders' equilibrium profit will decrease

¹⁵In the real world, such a strategy can be interpreted as lenders offering different interest rates for periods of a varying length, promotional interest rates or cash back rewards of a varying amount, which effectively changes b (Varian, 1980).

as the algorithms' accuracy increases. Taking derivative of lenders' equilibrium payoff with respect to P_b , we get

$$\frac{\partial \Pi_{rr}}{\partial P_b} = (\bar{b}\theta + \theta - 1)(1 - 2P_b) \leq 0 \quad (2.2)$$

As P_b increases, the common segment becomes more profitable, which intensifies competition. As competition increases, the equilibrium interest rate gets lower, which decreases the profit from both the common and the captive segments.

N-N: Neither Reveal

In the N-N case, borrowers do not have access to the pre-approval tools. Thus not every borrower has the correct belief of the approval outcome. In Table 2.2, we summarize the distribution of H and L type borrowers in different segments based on the lenders' predictions and the borrowers' beliefs of the predictions. As shown in Table 2.2, there are 16 segments. Each segment is defined by the combination of the lenders' predictions and the borrowers' beliefs of the predictions in the following sequence – Lender 1's prediction, borrowers' beliefs of Lender 1's prediction, Lender 2's prediction and borrowers' beliefs of Lender 2's prediction. For example, the segment $HLHL$ includes all borrowers whom Lender 1 predicts as H type, but believe Lender 1 predicts them as L type, Lender 2 predicts as H type, but believe Lender 2 predicts them as L type. In Table 2.2, the number of H (L) type borrowers that belong to a segment are shown in the top (bottom) row in the corresponding cell. The calculations are straightforward. For example, to calculate the number of H type borrowers in $HLHL$, we multiply the following: (i) P_b , the probability of Lender 1 predicting the H type correctly, (ii) $1 - P_c$, the portion of the borrowers have incorrect belief of the Lender 1's prediction, (iii) P_b , the probability of Lender 2 predicting the H type correctly, (iv) P_c , the portion of the borrowers have incorrect belief of the Lender 2's prediction, and (v) θ , the fraction of borrowers who are H type in the population.

Without loss of generality, we discuss the results from Lender 1's perspective. Since the

				Prediction from Lender 1			
				H		L	
				Borrowers' Belief of Prediction		Borrowers' Belief of Prediction	
				H	L	H	L
Prediction from Lender 2	H	Borrowers' Belief of Prediction	H	$P_c^2 P_b^2 \theta$ $P_c^2 (1 - P_b)^2 (1 - \theta)$	$P_c (1 - P_c) P_b^2 \theta$ $P_c (1 - P_c) (1 - P_b)^2 (1 - \theta)$	$P_c (1 - P_c) P_b (1 - P_b) \theta$ $P_c (1 - P_c) P_b (1 - P_b) (1 - \theta)$	$P_c^2 P_b (1 - P_b) \theta$ $P_c^2 P_b (1 - P_b) (1 - \theta)$
			L	$P_c (1 - P_c) P_b^2 \theta$ $P_c (1 - P_c) (1 - P_b)^2 (1 - \theta)$	$(1 - P_c)^2 P_b^2 \theta$ $(1 - P_c)^2 (1 - P_b)^2 (1 - \theta)$	$(1 - P_c)^2 P_b (1 - P_b) \theta$ $(1 - P_c)^2 P_b (1 - P_b) (1 - \theta)$	$P_c (1 - P_c) P_b (1 - P_b) \theta$ $P_c (1 - P_c) P_b (1 - P_b) (1 - \theta)$
	L	Borrowers' Belief of Prediction	H	$P_c (1 - P_c) P_b (1 - P_b) \theta$ $P_c (1 - P_c) P_b (1 - P_b) (1 - \theta)$	$(1 - P_c)^2 P_b (1 - P_b) \theta$ $(1 - P_c)^2 P_b (1 - P_b) (1 - \theta)$	$(1 - P_c)^2 (1 - P_b)^2 \theta$ $(1 - P_c)^2 P_b^2 (1 - \theta)$	$P_c (1 - P_c) (1 - P_b)^2 \theta$ $P_c (1 - P_c) P_b^2 (1 - \theta)$
			L	$P_c^2 P_b (1 - P_b) \theta$ $P_c^2 P_b (1 - P_b) (1 - \theta)$	$P_c (1 - P_c) P_b (1 - P_b) \theta$ $P_c (1 - P_c) P_b (1 - P_b) (1 - \theta)$	$P_c (1 - P_c) (1 - P_b)^2 \theta$ $P_c (1 - P_c) P_b^2 (1 - \theta)$	$P_c^2 (1 - P_b)^2 \theta$ $P_c^2 P_b^2 (1 - \theta)$

Table 2.2: Common Segment vs Captive Segment for the N-N case

Note: From Lender 1's perspective, the cells (Row 1, Column 1) and (Row 3, Column 1) combined are the common segment while the cell (Row 2, Column 1) and (Row 4, Column 1) combined are its captive segment. From Lender 2's perspective, cell (Row 1, Column 1) and (Row 1, Column 3) combined are the common segment while cells (Row 1, Column 2) and (Row 1, Column 4) combined are its captive segment. The top line in each cell denotes the number of H type borrowers while the bottom line in each cell denotes the number of L type borrowers.

lender would not approve a borrower whom its own algorithm classifies as L type, Lender 1 will not approve any borrower who falls under columns 3 and 4 in Table 2.2. Further, borrowers who believe they will not be approved by a lender would not apply to that lender. Hence, the only relevant borrowers for Lender 1 are the ones who fall under Column 1. We define the following sets to capture the borrowers in Column 1 of Table 2.2 for ease of explanation. Note that $Y_{k,c}^1$ ($Y_{k,c}^2$) denotes an borrower(k)'s belief of predictions for Lender 1 (2).

$$A_{hh}^{nn} = \{k \in A | Y_{k,c}^1 = H \wedge Y_{k,c}^2 = H \wedge Y_{k,b}^1 = H\}$$

$$A_{hl}^{nn} = \{k \in A | Y_{k,c}^1 = H \wedge Y_{k,c}^2 = L \wedge Y_{k,b}^1 = H\}$$

Specifically, A_{hh}^{nn} denotes the set of borrowers who are predicted as H type by Lender 1

and who believe will be predicted as H type by Lender 1. This is the common segment since borrowers in this segment do not have a specific preference for either lender. In Table 2.2, the cell (Row 1, Column 1) and (Row 3, Column 1) together constitute A_{hh}^{nn} . A_{hl}^{nn} denotes the set of borrowers who are predicted as H type by Lender 1, and who believe will be approved by Lender 1 but rejected by Lender 2. This is the captive segment for Lender 1 since the borrowers in this segment will always apply to Lender 1. In Table 2.2, the cells (Row 2, Column 1) and (Row 4, Column 1) together constitute A_{hl}^{nn} . Further, we use N_{xy}^{nn} to denote the number of borrowers and V_{xy}^{nn} to denote the fraction of H type borrower in the set A_{xy}^{nn} where $x, y \in \{h, l\}$. The values of these quantities can be easily calculated from Table 2.2 and are reported in the Appendix B.3.2.

In comparison to the R-R case, the level of competition among lenders in the N-N case is relatively lower. The degree of competition is affected by the relative profitability of the common and captive segments. Mathematically, this can be represented as $N_{hh}^{nn} - N_{hl}^{nn} < N_{hh}^{rr} - N_{hl}^{rr}$ and $V_{hh}^{nn} - V_{hl}^{nn} < V_{hh}^{rr} - V_{hl}^{rr}$. Consequently, in contrast to the R-R case where the A_{hh}^{rr} segment is highly profitable relative to the A_{hl}^{rr} segment, the difference between the profitability of the A_{hh}^{nn} and A_{hl}^{nn} segments is diminished in the N-N case. This is due to a higher probability of H type borrowers erroneously assuming that they will only be approved by one of the lenders as a result of mis-specified beliefs. As a result, they become part of the captive segment for lenders in the N-N case, thus making the captive segment more attractive.

We next solve for the optimal strategies in setting b for the two lenders in the N-N sub-game. As in the R-R case, a pure strategy equilibrium in interest rates does not exist in the N-N case either. Therefore, we study the mixed strategy equilibrium, and focus on the symmetric Nash equilibrium, i.e., the two lenders use the same mixed strategy. We use a probability distribution with a CDF $F^{nn}(b)$ to characterize lenders' pricing strategy in the N-N case.

Lemma 7 *The CDF of the distribution of lender's equilibrium pricing strategy in b is as*

follows:

$$F^{nn}(b) = \begin{cases} 0 & \text{if } b < \underline{b}^{nn} \\ \frac{bN_{hh}^{nn}V_{hh}^{nn} + bN_{hl}^{nn}V_{hl}^{nn} + N_{hh}^{nn}V_{hh}^{nn} - N_{hh}^{nn} + N_{hl}^{nn}V_{hl}^{nn} - N_{hl}^{nn}(\bar{b}V_{hl}^{nn} + V_{hl}^{nn} - 1) - N_{hl}^{nn}}{N_{hh}^{nn}(bV_{hh}^{nn} + V_{hh}^{nn} - 1)} & \text{if } \underline{b}^{nn} \leq b < \bar{b} \\ 1 & \text{if } b \geq \bar{b} \end{cases}$$

Where $\underline{b}^{nn} = \frac{\bar{b}N_{hl}V_{hl} - N_{hh}V_{hh} + N_{hh}}{N_{hh}V_{hh} + N_{hl}V_{hl}}$

Each lender's equilibrium payoff is:

$$\Pi_{nn} = N_{hl}^{nn}(\bar{b}V_{hl}^{nn} + V_{hl}^{nn} - 1)$$

The proof of Lemma 7 can be found in Appendix B.3.2.

R-N: Asymmetric Revealing

In the context of the R-N case, one of the lenders provides pre-approval tool to borrowers, enabling borrowers to receive predictions that accurately reflect the true approval outcomes for that particular lender. Conversely, the other lender chooses not to provide the pre-approval tools, which results in borrowers relying on their mis-specified beliefs. Since the two lenders are symmetric, without loss of generality, we will assume that Lender 1 is the lender that provides the pre-approval tool, while Lender 2 is the lender that does not.

Similar to what we have done for the R-R and N-N case, in Table 2.3, we show the distribution of H and L type borrowers across different segments based on the lenders' predictions and borrowers' beliefs of the predictions. As can be seen in Table 2.3, there are 8 segments. Each segment is defined by the combination of the lenders' predictions and borrowers' beliefs of lenders' predictions in the following sequence – Lender 1's prediction, Lender 2's prediction and the borrowers' beliefs of Lender 2's prediction. For example, the segment HLH includes all borrowers whom Lender 1 predicts as H type, Lender 2 predicts

as L type, and who believe Lender 2 predicts them as H type. In Table 2.3, the number of $H(L)$ type borrowers that belong to a segment is shown in the top (bottom) line in the corresponding cell.

				Prediction from Lender 1	
				H	L
Prediction from Lender 2	H	Borrowers' Belief of Prediction	H	$P_c P_b^2 \theta$ $P_c (1 - P_b)^2 (1 - \theta)$	$P_c P_b (1 - P_b) \theta$ $P_c P_b (1 - P_b) (1 - \theta)$
			L	$(1 - P_c) P_b^2 \theta$ $(1 - P_c) (1 - P_b)^2 (1 - \theta)$	$(1 - P_c) P_b (1 - P_b) \theta$ $(1 - P_c) P_b (1 - P_b) (1 - \theta)$
	L	Borrowers' Belief of Prediction	H	$(1 - P_c) P_b (1 - P_b) \theta$ $(1 - P_c) P_b (1 - P_b) (1 - \theta)$	$(1 - P_c) (1 - P_b)^2 \theta$ $(1 - P_c) P_b^2 (1 - \theta)$
			L	$P_c P_b (1 - P_b) \theta$ $P_c P_b (1 - P_b) (1 - \theta)$	$P_c (1 - P_b)^2 \theta$ $P_c P_b^2 (1 - \theta)$

Table 2.3: Common Segment vs Captive Segment for the R-N case

Note: From Lender 1's perspective, the cells (Row 1, Column 1) and (Row 3, Column 1) together are the common segment, while the cells (Row 2, Column 1) and (Row 4, Column 1) together are its captive segment. From Lender 2's perspective, the cell (Row 1, Column 1) is the common segment, while the cell (Row 1, Column 2) is its captive segment. The top line in each cell denotes the number of H type borrowers while the bottom line in each cell denotes the number of L type borrowers.

The only relevant borrower segments for Lender 1 are the ones that fall under Column 1; the only relevant borrower segments for Lender 2 are those that fall in Row 1. We define the following sets to capture the borrowers in Column 1 and/or Row 1 of Table 2.3 for ease of explanation.

$$A_{hh}^1 = \{k \in A | Y_{k,b}^1 = H \wedge Y_{k,c}^2 = H\}$$

$$A_{hh}^2 = \{k \in A | Y_{k,b}^1 = H \wedge Y_{k,c}^2 = H \wedge Y_{k,b}^2 = H\}$$

$$A_{hl}^1 = \{k \in A | Y_{k,b}^1 = H \wedge Y_{k,c}^2 = L\}$$

$$A_{lh}^2 = \{k \in A | Y_{k,b}^1 = L \wedge Y_{k,c}^2 = H \wedge Y_{k,b}^2 = H\}$$

Specifically, A_{hh}^1 denotes the set of borrowers who will be approved by Lender 1 and believe will be approved by both Lender 1 and Lender 2. A_{hh}^2 denotes the set of borrowers who will be approved by Lender 2 and believe will be approved by both Lender 1 and Lender 2. These are the common segment for Lender 1 and Lender 2 respectively since borrowers in these sets do not have a preference between the two lenders. In Table 2.3, the cells (Row 1, Column 1) and (Row 3, Column 1) together constitute A_{hh}^1 , the cell (Row 1, Column 1) corresponds to A_{hh}^2 . A_{hl}^1 denotes the set of borrowers who will be approved by Lender 1 and believe will only be approved by Lender 1. This is the captive segment for Lender 1. A_{lh}^2 denotes the set of borrowers will be approved by Lender 2 and believe will only be approved by Lender 2. In Table 2.3, the cells (Row 2, Column 1) and (Row 4, Column 1) together constitute A_{hl}^1 , and the cell (Row 1, Column 2) corresponds to A_{lh}^2 .

Again, N_{xy}^k denotes the number of borrowers, and V_{xy}^k denotes the fraction of H type borrowers in set A_{xy}^k where $x, y \in \{h, l\}$ and $k \in \{1, 2\}$. Lenders set interest rates, b_1 and b_2 simultaneously to compete. As before, a pure strategy equilibrium in interest rates does not exist. We study the mixed strategy equilibrium. We use the CDFs $F_1(b)$ and $F_2(b)$ to characterize Lender 1's (revealing lender) and Lender 2's (non-revealing lender) pricing strategies respectively. The lenders' equilibrium strategies are summarized in Lemma 8; the proof as well as the values for N_{xy}^k and V_{xy}^k can be found in Appendix B.3.3.

Lemma 8 *The CDF for the revealing lender's equilibrium pricing strategy in b is*

$$F_1(b) = \begin{cases} 0 & \text{if } b < \underline{b}^{rn} \\ \frac{bN_{hh}^2 V_{hh}^2 + bN_{lh}^2 V_{lh}^2 - k_2 + N_{hh}^2 V_{hh}^2 - N_{hh}^2 + N_{lh}^2 V_{lh}^2 - N_{lh}^2}{N_{hh}^2 (bV_{hh}^2 + V_{hh}^2 - 1)} & \text{if } \underline{b}^{rn} \leq b < \bar{b} \\ 1 & \text{if } b \geq \bar{b} \end{cases}$$

The CDF for the non-revealing lender's equilibrium pricing strategy in b is

$$F_2(b) = \begin{cases} 0 & \text{if } b < \underline{b}^{rn} \\ \frac{bN_{hh}^1 V_{hh}^1 + bN_{hl}^1 V_{hl}^1 - k_1 + N_{hh}^1 V_{hh}^1 - N_{hh}^1 + N_{hl}^1 V_{hl}^1 - N_{hl}^1}{N_{hh}^1 (bV_{hh}^1 + V_{hh}^1 - 1)} & \text{if } \underline{b}^{rn} \leq b < \bar{b} \\ 1 & \text{if } b \geq \bar{b} \end{cases}$$

$$\text{Where } \underline{b}^{rn} = \frac{\bar{b}N_{hl}^1 V_{hl}^1 - N_{hh}^1 V_{hh}^1 + N_{hh}^1}{N_{hh}^1 V_{hh}^1 + N_{hl}^1 V_{hl}^1}$$

The equilibrium payoffs to the revealing lender and the non-revealing lender are

$$\begin{aligned} \Pi_{rn}^1 &= N_{hl}^1 (\bar{b}V_{hl}^1 + V_{hl}^1 - 1) \\ \Pi_{rn}^2 &= ((1 + \bar{b})(N_{hh}^2 N_{hl}^1 V_{hh}^2 V_{hl}^1 + N_{hl}^1 N_{lh}^2 V_{hl}^1 V_{lh}^2) + N_{hh}^1 N_{hh}^2 (V_{hh}^2 - V_{hh}^1) \\ &\quad - N_{hh}^1 N_{lh}^2 (-V_{hh}^1 + V_{lh}^2) - N_{hh}^2 N_{hl}^1 V_{hl}^1 - N_{hl}^1 N_{lh}^2 V_{hl}^1) / (N_{hh}^1 V_{hh}^1 + N_{hl}^1 V_{hl}^1) \end{aligned}$$

Mathematically, we have $F_1(b) \leq F_2(b)$, $\forall b \in [0, \bar{b}]$. This implies that the revealing lender sets a higher interest rate than the non-revealing lender in general. Figure 2.2 illustrates the CDFs of the two lenders' strategies in the R-N case, together with the lenders' equilibrium strategies in the N-N and R-R cases. As is shown in the figure, the interest rate set in the R-R case is the lowest, indicating intense competition between the two lenders in that case.

Lender 1 derives benefits from two factors. Firstly, since it provides the pre-approval tool, any borrower of type H who will be approved by its algorithm will not mistakenly believe that Lender 1 will reject it. Secondly, since Lender 2 does not provide the approval tool, many borrowers of type H who will be approved by both lenders may incorrectly believe that only Lender 1 will approve them. In contrast, for Lender 2, the captive segment consists only of borrowers whom Lender 1's algorithm will reject. Therefore, the posterior probability of Lender 2's captive segment borrowers belonging to type H is lower. This makes the captive segment significantly more attractive for Lender 1 as compared to Lender 2. The reduced intensity of competition faced by Lender 1 as a result of this makes it possible for them




Figure 2.2: Lenders' equilibrium pricing strategies in each sub-game; $P_b = 0.8$, $P_c = 0.8$, $\theta = 0.7$, $\bar{b} = 0.5$

to charge a higher interest rate. The high interest rate charged by Lender 1 also leads to softened competition for Lender 2, which can potentially benefit from this scenario.

In Figure 2.3, we present the equilibrium payoff breakdown for each lender in the R-N case across different borrower segments. Since the lenders are playing mixed strategies, the profit breakdown is plotted as a function of b , assuming that the competing lender is using the equilibrium mixed strategy. The solid lines in the figure decrease with increasing b , as setting a higher value of b makes the focal lender's product less competitive in the common segment (A_{hh}^1 or A_{hh}^2). In contrast, the dashed lines increase with increasing b , as borrowers in segments A_{hl}^1 or A_{lh}^2 are captive, and the profit generated from them is proportional to b . It is noteworthy that compared to the non-revealing lender, a larger proportion of the revealing lender's profit is generated from its captive segment.

It is not difficult to check that $\frac{\Pi_{rn}^2}{\Pi_{rn}^1} = P_c \leq 1$. This means the revealing lender gets a higher profit in equilibrium than the non-revealing lender, and the difference in the two lenders' profits decreases in P_c . As P_c increases, borrowers' beliefs of the approval outcome without the pre-approval tool become more accurate. Hence, the R-N case starts resembling more like the R-R case where the competition on interest rates is high driving profit down.




Figure 2.3: Lenders' equilibrium profit break down in the R-N case

Note: The solid lines represent the profit generated from common segment borrowers; the dashed lines represent the profit generated from captive segment borrowers; $P_b = 0.8$, $P_c = 0.8$, $\theta = 0.7$, $\bar{b} = 0.5$

2.4.5 Analysis: Pre-approval Revealing Stage

The Payoff Matrix and the Pre-approval Revealing Equilibrium

So far, we have solved all possible sub-games after Stage 2. We now solve for the equilibrium in Stage 1. The “pre-approval revealing game” that the two lenders are facing in Stage 1 is shown in the following table,

		Lender 2	
		Not Reveal	Reveal
Lender 1	Not Reveal	(Π_{nn}, Π_{nn})	(Π_{rn}^1, Π_{rn}^2)
	Reveal	(Π_{rn}^2, Π_{rn}^1)	(Π_{rr}, Π_{rr})

where Π_{rr} , Π_{nn} , Π_{rn}^1 and Π_{rn}^2 denote lenders' equilibrium payoffs under different sub-games, as defined in Lemmas 6, 7, and 8.

We prove that $\Pi_{rn}^1 > \Pi_{nn}$ always holds (see Appendix B.3.4), which rules out N-N as an equilibrium outcome. One of the lenders in the N-N case always has the incentive to deviate and reveal the pre-approval to benefit from the many more H type borrowers who would apply to it both in the common and the captive segments when the competing lender chooses

not to reveal.

The relationship between Π_{rn}^2 and Π_{rr} is contingent on the value of P_b . Specifically, when P_b is relatively high, Π_{rn}^2 is greater than Π_{rr} . In the R-R case, this implies that the common segment is more attractive than the captive segment, leading to intense competition and reduced profits for both lenders. However, in the R-N case, the non-revelation by lender 2's softens the competition, allowing lender 1 to focus on the attractive captive segment and charge higher interest rates to extract surplus. This is particularly beneficial when P_b is high. Hence, lender 2 would choose to adhere to the R-N case when P_b is high, as it has more incentive to avoid intensified competition.

Conversely, when P_b is low, the captive segment in the R-R case is composed of many H type borrowers, making the common segment only marginally attractive. Consequently, the competition is less intense in the R-R case. Hence, the benefits of the R-N case in softening competition are minimal when P_b is low. However, when P_b is low, the benefit of correcting borrowers' inaccurate beliefs about their approval odds is higher, making deviation from the R-N case more attractive for lender 2.

We summarize the pre-approval revealing equilibrium in Stage 1 in the following proposition.

Proposition 4 *The pre-approval revealing equilibrium is determined by lenders' algorithm's accuracy, P_b , as follows:*

1. *If lenders' algorithm's accuracy P_b is sufficiently low, i.e., $P_b < P_b^0$, then $\Pi_{rn}^2 < \Pi_{rr}$, and the revealing equilibrium is symmetric such that both lenders choose to reveal the pre-approval outcomes.*
2. *If lenders' algorithm's accuracy P_b is sufficiently high, i.e., $P_b \geq P_b^0$, then $\Pi_{rn}^2 \geq \Pi_{rr}$, and the revealing equilibrium is asymmetric such that only one lender chooses to reveal the pre-approval outcomes.*

The threshold P_b^0 above is a function of P_c and θ :

$$P_b^0 = \frac{\bar{b}P_c\theta + \bar{b}\theta + 3P_c\theta - 3P_c + \theta - 1 + \sqrt{\Delta}}{4\bar{b}P_c\theta + 2\bar{b}\theta + 4P_c\theta - 4P_c + 2\theta - 2}$$

$$\begin{aligned} \Delta = & \bar{b}^2 P_c^2 \theta^2 + 2\bar{b}^2 P_c \theta^2 + \bar{b}^2 \theta^2 - 2\bar{b} P_c^2 \theta^2 + 2\bar{b} P_c^2 \theta + 4\bar{b} P_c \theta^2 - 4\bar{b} P_c \theta + 2\bar{b} \theta^2 \\ & - 2\bar{b}\theta + P_c^2 \theta^2 - 2P_c^2 \theta + P_c^2 + 2P_c \theta^2 - 4P_c \theta + 2P_c + \theta^2 - 2\theta + 1 \end{aligned}$$

The proof can be found in Appendix B.3.4. The dependence of revealing equilibrium on P_b and P_c is shown in figure 2.4, which graphically illustrates Proposition 4. It can be seen that an asymmetric equilibrium will occur unless the accuracy of lenders' algorithm is very low.




Figure 2.4: Algorithm revealing equilibrium depending on P_b and P_c when $\theta = 0.8$, $\bar{b} = 0.5$

Figure 2.5 further illustrates the equilibrium profits under the pricing equilibrium in the three possible sub-games resulting from the two lenders' revealing decisions (i.e., R-R, N-N, and R-N). SPNE profits are denoted in black, and off-equilibrium profits are denoted in gray. By deviating from the symmetric N-N case to the R-N case, the revealing lender is able to get a higher profit. On the other hand, the non-revealing lender does not become worse off because of the softened competition on b .




Figure 2.5: Equilibrium profits in each pre-approval revealing sub-games when $P_c = 0.7$, $\theta = 0.8$, $\bar{b} = 0.5$

Note. SPNE profits are denoted in black, and off-equilibrium profits are denoted in gray. When $P_b < P_b^0$ (the vertical line in the figure), the SPNE is both lenders revealing the pre-approval tools in Stage 1 and setting interest rates according to Lemma 6 in Stage 2, and the corresponding equilibrium profits for the two lenders are represented by the black portion of the “R-R case” line. When $P_b \geq P_b^0$, the SPNE is one lender revealing the pre-approval tool and the other not revealing in Stage 1, and the revealing and non-revealing lenders setting interest rates according to Lemma 8. The corresponding equilibrium profit for the revealing lender is represented by the black portion of the “R-N case Lender 1” line and that for the non-revealing lender is represented by the black portion of the “R-N case Lender 2” line.

Effects of P_c and θ on revealing SPNE

The effects of P_c and θ on SPNE are summarized in Proposition 5.

Proposition 5 P_b^0 decreases with P_c and increases with θ .

The proof can be found in Appendix B.3.5. Proposition 5 states that the threshold that P_b needs to exceed for asymmetric revealing equilibrium to be sustained decreases with P_c . When P_c is high, the borrowers’ beliefs of the approval outcome are pretty accurate, only a small fraction of borrowers will hold incorrect beliefs. Consequently, there is little benefit from correcting borrowers’ incorrect beliefs about their approval outcomes. Hence, even at smaller values of P_b where there is little difference between the competition intensity in the R-N and the R-R cases, lender 2 would not deviate from R-N to R-R case.

Moreover, when θ , the portion of H type borrowers in the population, increases, the

threshold on P_b that is required to sustain asymmetric equilibrium increases. In comparison to the intensity of competition, the benefit of correcting borrowers' incorrect beliefs about approval outcomes is affected at a greater rate by θ . This is because, a greater fraction of borrowers who incorrectly think they will be denied by the non-revealing lender are likely to be H type when θ is high. As a result, a higher value of θ encourages the non-revealing lender in the R-N case to reveal its algorithm and switch to the R-R case.

2.4.6 Borrower Surplus and Social Welfare

We next examine borrowers' surplus under each sub-game resulting from lenders' decisions on pre-approval revelation in Stage 1. We proceed by first calculating total surplus of lenders and borrowers, and then subtracting the lenders' equilibrium payoff. Since the interest rates set by the lender do not directly affect total surplus,¹⁶ our approach to calculating borrower surplus avoids integration over b . Previously, we have divided borrowers into several segments according to the lenders' predictions and borrowers' beliefs on the predictions. Some segments of borrowers never apply to either of the lenders, and those borrowers generate zero total surplus to the system. Some segments of borrowers will apply but will not be approved, and those borrowers generate a negative surplus because there is a cost associated with being rejected. Other segments of borrowers will apply and will be approved, and whether they generate a positive total surplus to the system or not depends on their types.

The comparison of total surplus and borrower surplus in the three pre-approval revealing cases is summarized in Proposition 6.

Proposition 6 *Both total surplus and borrower surplus are the highest in the R-R case, followed by the R-N case, and the lowest in the N-N case. That is, $TS^{rr} > TS^{rn} > TS^{nn}$,*

¹⁶The reason is as follows. Borrowers who do not apply to either lender generate no surplus to the society; those who apply and are rejected each generate a negative surplus of $-m$; those who apply and are approved also generates a fixed amount of surplus, which depends on their type but not on b , because the interest payments are surplus transferred from borrowers to lenders. Also, the number of borrowers in each of the segments is not affected by b either, because as long as a borrower believes she will be approved by at least one lender, she will apply regardless of the interest rate.

and $CS^{rr} > CS^{rn} > CS^{nn}$.

The full derivation and expressions of total surplus and borrower surplus can be found in Appendix B.3.6. Here we focus on discussing the comparison between the R-N case and the R-R case because these are the two cases that can be sustained in the equilibrium. We first look at why $TS^{rr} > TS^{rn} > TS^{nn}$. In the R-R case, there are no mis-specified beliefs by borrowers regarding their approval outcomes. In contrast, borrowers' beliefs about their approval odds in the N-N case are the least accurate. The incorrect beliefs regarding approval outcomes lead borrowers to make sub-optimal application decisions, e.g. borrowers who will be approved choose not to apply whereas borrowers who will be rejected choose to apply. The number of such non-optimal decisions is the lowest in the R-R case and highest in the N-N case. These non-optimal decisions create a negative social surplus. A borrower who applies but gets rejected will create $-m$ social surplus, and a borrower who would be approved but does not apply will result in an opportunity cost to social surplus with a size of $\theta(M_h - m) + (1 - \theta)(M_l - m - 1)$.¹⁷ Hence, the social surplus is the highest in the R-R case, followed by the R-N case and then the N-N case.

As for borrower surplus, when comparing CS^{rn} with CS^{rr} , notice that (1) the number of borrowers who are approved in the R-R case is larger than in the R-N case; (2) the number of borrowers who are rejected in the R-R case is smaller than in the R-N case; (3) the interest rate set in the R-R case is on average lower than in the R-N case. (This is the case because $F^{rr}(b) > F_1(b)$, $F^{rr}(b) > F_2(b)$, $\forall b \in [0, \bar{b}]$.) In other words, in the R-R case, more borrowers are approved, fewer borrowers are rejected, and the interest rates set by the lenders are lower compared with the R-N case. Consequently, borrowers' welfare is higher in the R-R case than in the R-N case. The borrower surplus in the N-N case, CS^{nn} , is the smallest among the three. The reason is that there are a large number of borrowers who make sub-optimal applying decisions because of the large market friction in terms of borrowers' beliefs of the

¹⁷This value is always positive. Rewrite the expression as $[\theta\bar{b} + (1 - \theta)(\bar{b} - 1)] + [\theta(M_h - m - \bar{b}) + (1 - \theta)(M_l - m - \bar{b})]$. The first part is greater than 0 under Assumption 1, and the second part is greater than 0 following the first inequality in Assumption 2.

approval. Even when one lender sets a higher interest rate in the R-N case than it would do in the N-N case, moving from N-N to R-N will still benefit borrowers.

2.5 Mandating Algorithm Revelation

As discussed previously, both borrower surplus and social surplus are the highest in the R-R case. However, when P_b is high, only one lender would want to reveal the pre-approval outcomes. Does it imply that mandatory disclosure could be socially desirable from the policy makers' perspective? We tackle this question by showing the other side of the story – the effect of mandatory revealing on lenders' incentive to improve their algorithms. While in the previous analysis, we assume that the lenders' algorithm accuracy P_b is exogenous, in this extension we consider it a decision that lenders have to make prior to Stage 1 in the current model.

In this extension, we study two scenarios: a mandatory scenario where lenders are required to reveal the pre-approval outcomes, and a voluntary scenario where such requirements are absent. We will focus on the parameter range where the two lenders make asymmetry revealing decisions voluntarily, because outside of this range, mandatory revealing is redundant. In other words, we assume that $P_b > P_b^0$, where P_b^0 is defined in Proposition 4. We assume there are advances in technology that could help lenders increase their algorithms' accuracy from P_b to P_b^* without any cost, and study lenders' incentives to take on this opportunity in both scenarios. Intuitively, it may appear that a lender should always improve its algorithm's accuracy when such improvement is costless. However, we will demonstrate that this is not the case. Below, we start with the mandatory scenario. The lenders' equilibrium strategies in upgrading their algorithms in this scenario are summarized in Proposition 7.

Proposition 7 *In the mandatory revealing scenario, there is a unique symmetric mix-strategy equilibrium in upgrading screening algorithms, where each lender chooses to upgrade its algorithm with probability $P^M = \frac{1-P_b}{P_b^*}$.*

The proof of Proposition 7 can be found in Appendix B.3.7. An interesting observation is that P^M decreases in P_b and P_b^* , which means that when lenders' algorithms' accuracy is relatively high, lenders will have little incentive to further increase it. The reason is that when lenders' algorithms' accuracy is high, competition is especially intense in the R-R case. Any further increase in algorithms' accuracy would make the common segment even more profitable compared to the captive segment thus further intensify the competition on the interest rate, which would drive down lenders' equilibrium profits. Since the mandatory revealing policy is only relevant in the case where P_b is relatively large (i.e., $P_b > P_b^0$), lenders will have a relatively small incentive to improve their screening algorithms if revealing is made mandatory.

We next consider the voluntary scenario. The lenders' equilibrium strategies in upgrading their algorithm in the voluntary scenario are shown in Proposition 8

Proposition 8 *In the voluntary revealing scenario, the revealing lender always chooses to upgrade its algorithm, while the non-revealing lender chooses to upgrade the algorithm if $\theta \leq \theta^0$ or $P_c \leq P_c^0$, where*

$$\begin{cases} \theta^0 = \frac{2P_bP_b^*P_c - P_bP_b^* - 2P_bP_c + P_b + 2P_b^{*2}P_c - P_b^{*2} - 3P_b^*P_c + 2P_b^* + P_c}{-2P_bP_c + P_b + P_b^{*2}(2P_c - 1)(\bar{b} + 1) + P_b^*((\bar{b} + 1)(2P_bP_c - P_b - P_c) + 2 - 2P_c) + P_c} \\ P_c^0 = \frac{P_b + P_b^*}{2P_b + 2P_b^* - 1} \end{cases}$$

The proof of Proposition 8 can be found in Appendix B.3.8. Proposition 8 says that in the voluntary scenario, both lenders will adopt the higher accuracy algorithm unless both θ and P_c are high, in which case only the revealing lender chooses to upgrade the algorithm. We show in the appendix that upgrading the algorithm is a dominant strategy for the revealing lender. The intuition is as follows: In Section 2.4, we have shown that in the R-N case, the non-revealing lender chooses to focus on capturing the common segment because borrowers in the common segment are the ones who will be approved by the revealing lender, so these borrowers are of high quality on average. At the same time, the revealing lender chooses to

focus on its captive segment because although these borrowers believe will be rejected by the non-revealing lender, many of them are actually of high quality and will be approved by the non-revealing lender. Consequently, the competition is softened. If the revealing lender increases its algorithm's accuracy, the non-revealing lender's common segment will become even better because of the fact that the revealing lender's approval says more when the screening algorithm is accurate. At the same time, the revealing lender can rely more on its own prediction, thus it will have more incentive to exploit its captive segment, and therefore, competition will be further softened. By contrast, the increase in the non-revealing lender's algorithm will intensify the competition, because it will increase the quality of the revealing lender's common segment and the non-revealing lender's captive segment. Thus the non-revealing lender has to consider the relative effect sizes of improved screening ability and increased competition due to an increase in its algorithm's accuracy. When $\theta \leq \theta^0$, the market is risky in the sense that there is large fraction of L type borrowers, and there are significant benefits from a better screening ability. Similarly, when $P_c \leq P_c^0$, the direct impact of an increase in the accuracy of the non-revealing lender's algorithm on intensifying the competition is fairly weak. As a result, the non-revealing lender would choose to increase its algorithm's accuracy only when $\theta \leq \theta^0$ or $P_c \leq P_c^0$.

Our analysis above suggests that the implications of policy makers' decisions on regulating algorithm revelation are not as straightforward as they appear intuitively. Since increasing the accuracy of screening algorithms can help allocate fund to more creditworthy borrowers, failing to do so may potentially hurt borrower surplus. The following proposition specifies a sufficient condition under which mandatory algorithm revealing will hurt consumer surplus.

Proposition 9 *Mandatory algorithm revealing will hurt borrower surplus if $\theta \leq \theta^0$ and $P_c > \underline{P}_c$, Where θ^0 is defined in Proposition 8, and*

$$\begin{aligned}
P_c &= \frac{-B + \sqrt{B^2 - 4AC}}{2A} \\
A &= 2\bar{b}P_b^{*2}\theta - \bar{b}P_b^*\theta + 2P_b^{*2}\theta - 2P_b^{*2} - 3P_b^*\theta + 3P_b^* + \theta - 1 \\
B &= m\theta + (P_b^{*2} - P_b)(\bar{b}\theta + 3m - 2M_h\theta + 2M_l\theta - 2M_l - \theta + 1) + P_b^*(m - 2m\theta) \\
C &= 1 - M_h\theta + 2P_b^2(-\bar{b}\theta - \theta + 1) + 2P_b(\bar{b}\theta + \theta - 1) - \theta \\
&\quad + P_b^{*2}(-\bar{b}\theta - 2m + M_h\theta - M_l\theta + M_l) + 2P_b^*(m + M_l\theta - M_l) \\
&\quad + (-m + M_h\theta - M_l\theta + M_l + \theta - 1)(P_b^4 - 2P_b^3 + P_b^2)/P_b^{*2}
\end{aligned}$$

Proposition 9 suggests that mandating revealing of screening algorithms could hurt consumer surplus when P_c is high but θ is low. First, the condition on θ is to ensure that the lenders have an incentive to increase accuracy in the R-N case, according to Proposition 8. The intuition behind the P_c related condition will become clear once we layout the upside and downside of such a mandatory revealing policy. On the upside, mandatory revealing will reduce market friction by eliminating borrowers' misspecified beliefs on approval. On the downside, it will impede lenders from adopting more accurate algorithms. The benefit from the upside is decreasing in P_c since the benefit of revealing the algorithm is low if borrowers' assumptions of approval are already accurate. Moreover, the loss due to the downside is increasing in P_c : when borrowers' assumptions of approval are more accurate (i.e., P_c is higher), borrowers sub-optimal applying behavior is reduced, and the benefit of lenders adopting more advanced algorithms becomes larger since any improvement in screening accuracy will directly improve the chances that credit is allocated to the more deserving H type borrowers.

2.6 Conclusions

2.6.1 Summary of Results

The use of ML algorithms and data driven decision making by financial lenders seems win-win for both the lenders and the borrowers – lenders are able to better screen borrowers for their credit worthiness and deserving borrowers are more likely to get approved for credit (Fu et al., 2021b, Wei et al., 2016, Chan et al., 2022, Netzer et al., 2019). However, as lenders get better at screening the borrowers, the later continue to face considerable uncertainty in their chances of getting approved for credit by a lender (Citron and Pasquale, 2014b, Fu et al., 2020).

This paper explores the factors influencing lenders’ decisions to provide approval odds to borrowers through the pre-approval tools. Despite the advantages of using machine learning algorithms and data-driven decision-making, borrowers face significant uncertainty in their chances of credit approval as lenders improve their screening processes. Our analysis reveals that lenders strategically use asymmetric revealing of pre-approval outcomes to soften competition and focus on different borrower segments, particularly when their algorithms’ accuracy is high, the borrowers’ belief of approval is also accurate, or the market is risky with a significant portion of low-creditworthiness borrowers.

The asymmetric revealing equilibrium creates product differentiation, where the revealing lender focuses more on its captive segment and charges higher interest rates, while the non-revealing lender focuses on the common segment and charges lower rates. The borrower surplus and total surplus are the highest when both lenders reveal their algorithms and the lowest when neither does. However, mandating lenders to provide accurate approval odds may not necessarily improve borrower surplus as it reduces lenders’ incentive to improve algorithm accuracy compared to voluntary revelation.

2.6.2 Managerial Implications

Our analysis yields several key findings that offer insights for future research and practice on algorithm revelation in financial lending. Firstly, the decision to provide personalized pre-approval outcomes is complex, as it can intensify or soften competition depending on the circumstances. Our model suggests that the accuracy of lenders' algorithms, accuracy of borrowers' belief of approval, as well as the market composition, mediate these effects. Additionally, our results indicate that at least one lender should always provide pre-approval tools, while both lenders should reveal their pre-approval outcomes in certain conditions. This highlights the promise of market forces leading to partial algorithmic transparency.

Secondly, the competitive implications of algorithm/pre-approval revelation are more severe in markets with highly accurate lenders' algorithms. This can lead to self-debilitating competition on interest rates, and lenders should avoid providing personalized pre-approval tools if their competitor is already doing so. Thirdly, the prior beliefs of borrowers regarding their approval odds can moderate the effects of algorithm/pre-approval revelation. When the borrowers' belief is less accurate, the market expansion effect from providing approval tools is stronger, and both lenders can benefit from revealing their pre-approval outcomes.

Fourthly, asymmetric pre-approval revelation can generate product differentiation, softening competition among financial lenders. When conditions for an asymmetric equilibrium are met, the revealing lender is better off than the non-revealing lender, and the non-revealing lender need not worry about being hurt by the asymmetric action of the competing lender. Finally, lenders' algorithms' accuracy allows for differentiation in financial products, even those with similar non-price features. While high accuracy reduces loss due to misclassification, it also intensifies competition, and lenders should carefully consider its impact.

2.6.3 Implications for Public Policy

Lenders' provision of pre-approval tools can enhance the efficiency of credit markets from a social welfare perspective by helping borrowers avoid non-optimal application decisions that waste resources. However, this does not imply that policy makers should always enforce lenders to disclose pre-approval outcomes to borrowers. It is suggested that policy makers understand the strategic reasons behind lenders' reluctance to reveal such information and consider the potential impact of mandatory pre-approval revealing on lender competition.

When lenders are compelled to reveal the pre-approval outcomes, they may have less incentive to invest in screening technologies, as shown in our analysis. Thus, when considering regulations on mandatory information disclosure, policy makers should consider how such regulations may impact subsequent competition and the incentive for lenders to invest in algorithmic screening technologies.

In general, policy makers should be especially cautious about mandatory pre-approval revealing in relatively risky markets and when the borrowers' beliefs of the approval are accurate. In those cases, the benefit of reduced market friction may be outweighed by the loss due to lack of adoption of better screening technologies by lenders.

2.6.4 Limitations

There are several factors that our model does not account for, providing opportunities for future research. For instance, we only consider lenders' competition on interest rate, we do not model their competition on other dimensions such as financial product features. Furthermore, we do not consider algorithmic gaming by borrowers or other channels through which borrowers can apply for loans, such as intermediary platforms. Investigating how these factors impact pre-approval revelation strategies would be interesting for future research. Despite these limitations, our paper provides valuable insights into algorithmic lending and the strategies lenders employ to reveal approval odds to borrowers. Our model explains the

observed asymmetric information revealing behavior in credit markets and introduces a new dimension of strategy to leverage competition, which previous models have not captured. Overall, our paper contributes to a better understanding of the interplay between competition and information revelation strategies.

2.6.5 Funding and Competing Interests

All authors certify that they have no affiliations with or involvement in any organization or entity with any financial interest or non-financial interest in the subject matter or materials discussed in this manuscript. The authors have no funding to report.

Chapter 3

Algorithms, Artificial Intelligence and Simple Rule-Based Pricing

3.1 Introduction

In recent years, the rapid advancement of information technology has revolutionized the landscape of pricing methods, with the emergence of software solutions that enable the automatic setting and adjustment of prices. These automated pricing algorithms have gained significant traction and widespread adoption, particularly in the context of online shopping platforms.

Simple rule-based algorithms, an important type of automated pricing algorithms, operate according to predetermined protocols defined by human agents and have gained significant popularity on e-commerce websites. Most major e-commerce marketplaces, including Amazon, eBay, Shopify, Walmart, and Google Shopping, provide simple rule-based pricing algorithms as free pricing tools to individual sellers. These algorithms typically allow sellers to set prices that are below, equal to, or above the lowest price in the predefined market, taking into account user-specified listing conditions and fulfillment methods.¹ For sellers requiring

¹See https://sellercentral.amazon.com/gp/help/external/200836360?ref=efph_200836360_cont_201186860&language=en_US.

greater flexibility, such as the ability to match prices of specific competitors or alternate between different strategies based on market conditions, more complicated rule-based pricing solutions are available through various third-party repricer services, including ChannelEngine, RepricerExpress, Aura, and Informed.co.

Despite their convenience and ease of use, simple rule-based algorithms can sometimes lead to problematic situations. A notable example occurred several years ago when algorithmic pricing first started to gain traction. A book priced at over \$23 million attracted significant attention, and upon investigation, it was discovered that two sellers were using automated pricing algorithms that pegged to each other, causing the prices to spiral out of control rapidly. Indeed, simple rule-based algorithms may prove suboptimal in many cases, potentially triggering "a race to the bottom" and harming sellers' long-term profitability.

In recent years, more sophisticated AI-based pricing algorithms have appeared, such as those based on reinforcement learning (RL). Different from the simple rule-based algorithms, these AI algorithms feature self-learning capabilities that enable them to autonomously identify optimal price points that maximize long-term profits through iterative pricing experiments. RL algorithms have recently been employed to solve various business problems, including display advertisement measurement (Tunuguntla, 2021), sequential coupon targeting (Liu, 2022, Wang et al., 2021), ad sequencing (Rafieian, 2022), and ad recommendations (Aramayo et al., 2022). Unsurprisingly, these algorithms have demonstrated significantly better performance compared to traditional methods. In the pricing domain specifically, researchers from Alibaba, the Chinese e-commerce giant, found that employing reinforcement learning-based pricing algorithms led to substantially improved profits for thousands of products, outperforming expert manual pricing (Liu et al., 2019). While none of the major platforms currently provide AI-based pricing algorithms as a built-in feature, some third-party repricer services have begun offering these sophisticated pricing methods to businesses.

Despite the AI-based algorithms being in the early stages of adoption in practice, they have garnered significant attention from researchers in various fields, including Economics,

Law, and Marketing. Among the research questions explored, one of the most crucial issues is whether these intelligent algorithms can lead to collusion. The findings from these studies suggest that in repeated price competitions within simulated market environments, RL algorithms have demonstrated a remarkable ability to learn sophisticated pricing strategies that soften competition and enable the maintenance of supra-competitive prices in competitive markets (Calvano et al., 2020, Hansen et al., 2021, Klein, 2021, Johnson et al., 2020, Asker et al., 2022). This has raised concerns about the potential for algorithmic collusion and its implications for market efficiency, consumer welfare, and regulatory frameworks. As AI-based pricing algorithms become more prevalent, understanding their impact on market dynamics and developing appropriate governance mechanisms will be critical to ensuring fair competition and protecting consumer interests.

While almost all the research work focus on the competition between AI pricing algorithms, disproportional less attention has been paid to the simple rule-based algorithm. However, simple rule-based algorithms are actually very prevalent on online marketplaces. Analyzing a sample of 1,600 best-selling items on Amazon, Chen et al. (2016) find that more than one-third of sellers were using simple rule-based pricing strategies in 2015. A 2017 E-Commerce sector inquiry found that 50% of retailers in the European Union track their online competitors' prices, and 70% of them automatically adjust their prices according to simple rules. With such a prevalence of simple rule-based algorithms, it is natural to wonder how a simple rule-based algorithm would perform when competing with AI-based algorithms. Surprisingly, this topic remains understudied in the literature, despite the widespread use of simple rule-based algorithms. Our work is among the first to systematically examine this problem, focusing on the equilibrium behavior and dynamics when algorithms of different levels of sophistication interact. This problem is of great importance in understanding sellers' strategic choices when faced with competitors adopting newly emerged, sophisticated pricing technologies

Specifically, the central research question of this study is whether a seller should adopt a simple rule-based pricing algorithm or an RL pricing algorithm when competing against

rivals using RL pricing algorithms. This question is internally important, as it determines the final equilibrium of algorithm choice and the prevalence of each type of pricing algorithm in the long run.

Our study adopts a three-pronged strategy to address the research questions at hand. The first approach, follows the established economics literature that examines competition among pricing algorithms in simulated market environments, with a specific focus on the canonical Logit demand model as demonstrated in previous works such as Calvano et al. (2020), Klein (2021), and Asker et al. (2022). The second approach leverages a more sophisticated structural demand model that takes into account realistic model parameters and consumer behavior during search and purchase processes. In our third approach, we employ theoretical analysis with simplifying assumptions to identify boundary conditions that determine when Q-learning algorithms and rule-based algorithms may dominate each other.

In the baseline case, we use the *Undercut Lowest Price (UL)* algorithm as an example of a simple rule-based pricing algorithm, as it is one of the most commonly applied in practice. This algorithm operates by setting the price slightly below the lowest price offered by competitors in the market. For the RL pricing algorithm, we employ the *Q-learning* algorithm, which is a basic yet popular choice in the field of reinforcement learning.

To compare the behavior and performance of these algorithms, we consider two distinct scenarios: the Q-Q scenario and the Q-R scenario. In the Q-Q scenario, both the focal seller and the competing seller adopt Q-learning algorithms. In the Q-R scenario, the focal seller uses a simple rule-based algorithm (UL), while the competing seller employs a Q-learning algorithm. In each scenario, we allow the algorithms to interact repeatedly, simulating the ongoing competition in a market environment. We record and compare the market outcomes across three key aspects: 1), Equilibrium price and profit: 2), Learning dynamics. 3), Deviation and punishment strategies:

Our research findings reveal a compelling result: the focal seller achieves higher profits in the Q-R scenario compared to the Q-Q scenario. We attribute this outcome to the distinct

learning dynamics present in each scenario and the specific properties of the simple rule-based strategy. In the Q-Q scenario, both Q-learning algorithms engage in simultaneous learning through exploration, which leads to non-stationarity in the environment. This non-stationarity makes it more challenging for either algorithm to learn the most profitable collusive strategy. Conversely, in the Q-R scenario, the environment becomes stable because the strategy of the simple rule-based algorithm is fixed. This stability enables the competitor's Q-learning algorithm to learn the optimal policy more easily. Interestingly, the Q-learning algorithm's optimal learning also benefits the simple rule-based algorithm, which is pegged to it. The simple rule-based algorithm, in this case, represents a variant of the well-known "Tit for Tat" strategy, which has proven to be a formidable strategy in repeated games. Our study also sheds light on the off-equilibrium strategies learned by the Q-learning algorithm in different scenarios. In the Q-Q scenario, we observe that the algorithms learn a multi-period punishment strategy. In contrast, the Q-learning algorithm in the Q-R scenario, learns an "always cooperate" strategy.

Importantly, our findings remain consistent and robust when considering various other factors. We have examined the impact of different algorithm parameters, and different market structures, including scenarios with product differentiation, and demand uncertainty. We have also tested the robustness of our findings by considering alternative simple rule-based pricing algorithms beyond UL.

The findings from our structural modeling and theoretical modeling approaches consistently support our main conclusions. The analysis based on the structural demand model suggests that a focal seller can achieve a significant profit increase of approximately 9% by adopting a simple rule-based algorithm. Our theoretical analysis reveals that the focal seller stands to gain more by adopting a simple rule-based algorithm if the discount factor is sufficiently high.

Our findings have significant managerial implications. In markets where competitors are adopting advanced pricing technologies, it may be more practical and profitable for a seller to employ a simpler rule-based pricing strategy instead of attempting to match the

sophistication of their competitors. Furthermore, our research has important implications for policymakers concerned with the potential for tacit collusion facilitated by pricing algorithms. While much attention has been focused on the collusive risks associated with the widespread adoption of sophisticated pricing algorithms, our findings suggest that policymakers should also pay close attention to markets where the majority of pricing algorithms employed are simple and rule-based.

The rest of the paper is organized as follows: In Section 3.2, we discuss relevant literature. Section 3.3 expounds on the model setup and Section 3.4 presents the results of the experiments conducted in the oligopoly logit market environment. Section 3.5 outlines the structural demand model and the corresponding results. Section 3.6 proposes a theoretical framework. Finally, in Section 3.7, we provide a discussion and conclusion of our findings.

3.2 Relevant Literature

This study contributes to the growing literature on algorithmic pricing and collusion, which investigates how and why pricing algorithms might increase market prices. Several potential channels have been discussed, including the ability of algorithms to predict future demand levels, affecting incentives to deviate to lower prices during periods of high predicted demand and discouraging collusion (Miklós-Thal and Tucker, 2019, O’Connor and Wilson, 2021). Studies focusing on learning algorithms have demonstrated the ability of self-learning intelligent algorithms (i.e., RL) to learn to charge supra-competitive prices without coordination or communication, a phenomenon known as tacit collusion. Explanations for this phenomenon include the generation of prices correlated with other algorithms, leading to overestimation of price sensitivity (Hansen et al., 2021), and the learning of punishment strategies by intelligent algorithms with memory to thwart deviation and sustain supra-competitive prices (Calvano et al., 2020).

Despite their widespread use, simple rule-based algorithms have not received much

attention in the literature. Our paper is among the first to systematically examine the pricing dynamics and equilibrium in competitive markets that include both RL and simple rule-based algorithms, particularly focusing on the strategic selection of algorithms by sellers based on the level of sophistication of the pricing algorithms. Our study makes at least two notable contributions to the algorithmic pricing literature. First, we highlight a factor that has been largely overlooked but could lead to increased market prices: the use of simple rule-based algorithms may result in higher equilibrium prices compared to more sophisticated RL pricing algorithms, especially in competitive scenarios involving other RL algorithms. Secondly, our analysis extends beyond merely examining the outcomes and equilibrium prices, we delve deeper into the learning dynamics and identify the underlying mechanisms that lead to different pricing results. We offer new perspectives on algorithmic collusion by demonstrating how the presence of multiple self-learning algorithms can introduce non-stationarity into the market environment in the learning process, making it more difficult to search for optimal strategies. Conversely, we explore how simple rule-based pricing algorithms can contribute to a more stationary market environment, allowing for the emergence of near-monopoly pricing strategies among intelligent algorithms.

While empirical studies on algorithmic pricing are relatively rare, some notable exceptions exist. For example, Assad et al. (2020) found that the adoption of algorithmic pricing by German gasoline retailers could have increased margins by up to 38%. Similarly, Brown and MacKay (2021) showed that the adoption of pricing algorithms significantly alters pricing patterns, and Chen et al. (2016) provided empirical evidence of the impact of pricing algorithms on market pricing behaviors.

The empirical examination of RL pricing algorithms is particularly challenging, with most of the literature relying on simulations using the canonical logit demand model (Calvano et al., 2020, Hansen et al., 2021, Johnson et al., 2020, Asker et al., 2022). Methodologically, our study extends the analysis by estimating a more realistic structural demand model and conducting counterfactuals based on this model, in addition to the typical logit demand

model simulations.

Our analytical framework extends the existing, yet limited, theoretical models of RL algorithms in the literature, which often assume 'stateless' algorithms for simplification (Hansen et al., 2021, Banchio and Mantegazza, 2023). We consider the multiplicity of states in our model, providing a richer representation of pricing strategies.

3.3 Competition With Pricing Algorithms

3.3.1 The Market

We study the oligopoly price competition in a repeated game setting. The market consists of I consumers ($i = 1, 2, \dots, I$), where each consumer demands at most one product. There are J differentiated products in the market provided by J sellers ($j = 1, 2, \dots, J$). Each seller sells their own product (with a constant marginal cost mc_j) in each and every period without capacity constraints. At the beginning of each time period ($t = 0, 1, \dots$), the sellers set prices for their products (p_j^t) simultaneously, and the price history from previous periods is common knowledge. At the end of each period, a per-period profit is realized for each seller $r_j^t = (p_j^t - mc_j)D_j^t$, where D_j^t is the per-period demand for seller j determined by a demand function $D_j^t(p_j^t, p_{-j}^t)$.

For model tractability, we use discretized prices throughout the paper. We use p_j^{Nash} and p_j^{Mono} to denote the Bertrand-Nash price (competitive price) and monopoly price of the stage game for seller j , respectively. We then construct a set of equally spaced grid points $A = \{p_{j,k}\}, k \in \{1, \dots, K\}$ where $p_{j,1} = p_j^{Nash}$, $p_{j,K} = p_j^{Mono}$, and $p_{j,2}$ to $p_{j,K-1}$ are evenly spaced in between. The space between two consecutive price grids $\Delta p_j = (p_j^{Mono} - p_j^{Nash})/(K - 1)$. We use $K = 10$ throughout the paper unless specified otherwise.

3.3.2 The Algorithm

Rule-Based Algorithm

There are three main types of rule-based pricing, depending on how price changes are triggered. The first type is *Competitor-based pricing*, where price changes are triggered by competitors' actions. This is the most common strategy used in practice – most large e-commerce platforms and almost all large third-party repricing platforms provide such tools. For example, the 'Target Low Price'² tool provided by Amazon allows sellers to adjust prices to match the lowest price in the market. On *informed.co*, sellers can specify the competitors they want to compete with, by either star rating, fulfillment methods, or item condition (whether new or used). Sellers can then choose to always set prices equal to these competitors' prices or lower than them by a fixed amount or percentage. The seller can also choose a minimum price to make sure the price never goes below that price.³ These repricers can react to competitors' price changes almost instantaneously, for example, according to *AlphaRepricer*, it can reprice a product every two minutes⁴, which means that any change in competitors' prices will trigger a reaction within two minutes. The second type of rule-based algorithm is *Sales-based pricing*, where price changes are triggered by the change in sales volume in the past. For example, Amazon provides sellers with tools to automatically decrease the price if its sales volume drops below a certain threshold in a certain period of time.⁵ The third type of rule-based algorithm is *Time-based pricing*, where price changes are triggered by the time of the day or day of the week. For example, using *RepricerExpress* – a third party repricing solution – sellers can stop repricing or raise their price at a pre-specified time of the day.⁶ Table 3.1

²This 'Target Low Price' mechanism should not be confused with the 'Match Price Policy' or 'Match Price Guarantee' provided by retailers like Walmart and Target. The latter guarantees that exactly the same product cannot be found elsewhere at a lower price. The former is a tool that individual sellers can use to adjust prices according to competitors' prices, where the matches are not necessarily defined by multiple sellers selling identical products, and prices are not necessarily matched exactly.

³<https://help.informed.co/en/articles/2319357-build-your-own>

⁴<https://alpharepricer.com/blogs/why-is-alpha-repricer-better-than-other-repricers/>

⁵<https://sellercentral.amazon.com/gp/help/external/FRJDFLPWZSAG67>

⁶<https://support.repricer.com/sleep-mode>

summarizes the commonly used rule-based pricing algorithms.

Table 3.1: Rule-based Algorithm Type

Algorithm Type	Provider	Example
Competitor-based	Amazon, RepricerExpress, Informed.co, BQool, ChannelAdvisor, SellerEngine, SellerActive, Aura, AlphaRepricer	Match Low Price (Always set price equal to the lowest price in the market)
Sales-based	Amazon, SellerEngine, SellerActive	Sales Velocity Pricing
Time-based	RepricerExpress	Sleep-mode Pricing

In this paper, we focus on the most commonly applied rule-based algorithm in the competitive market—the competitor-based pricing algorithm. Its popularity comes from the fact that it can react to opponent sellers’ price changes almost instantaneously to prevent loss. By contrast, the sales-based pricing algorithm adjusts price only after the effect of opponent sellers’ price change has been reflected in one’s own sales. Most third-party repricing software gives sellers options to customize the competitor-based repricing rule, including choosing which seller they wish to compete with, choosing to match price exactly or undercut, and whether or not to set a minimum price to prevent loss. In the benchmark case, we focus on a widely used competitor-based pricing algorithm: *Undercut Lowest Price (UL)*, where prices are adjusted to undercut the lowest price among competitors by a fixed amount or fixed percentage. In Section 3.4.2, we study three variants of *UL*: *Undercut Lowest Price with a Minimum Price*, *Target Lowest Price*, *Target Lowest Price with a Minimum Price*.

Specifically in our setting, the simple rule-based algorithm will always undercut the *last period* lowest price among competitors by one price grid. The policy function of the simple rule-based algorithm Φ_R is specified in Equation 3.1

$$\Phi_R : p_j^t = \min(p_{-j}^{t-1}) - \Delta p_j, \forall t \quad (3.1)$$

Q-learning Algorithm

Following the emerging literature on algorithmic pricing in economics (e.g. Calvano et al. (2020), Asker et al. (2022), Johnson et al. (2020), Waltman and Kaymak (2008)), we use

Q-learning as a representative example of AI-powered learning algorithm. There are several reasons for this choice. First, Q-learning is a *model-free* reinforcement learning algorithm, which means that it does not need any knowledge of the environment. It is well suited for pricing applications since market conditions and demand functions are often unknown to sellers. Second, Q-learning is one of the most popular RL algorithms used by computer scientists: the implementation is simple and the number of hyper-parameters to tune is significantly smaller compared with neural network based algorithms. Third, the logic behind Q-learning is clear, which makes it possible for us to study its learning dynamics.

A Q-learning algorithm can learn the optimal strategy in a stationary Markov decision process: in period t , given the state ($s^t \in S$), a Q-learning seller j takes an action ($p_j^t \in A$). A reward (r_j^t) is realized and the current state transits to a new state in the next period ($s^{t+1} \in S$) according to some time-invariant state transition rule. The goal of the Q-learning algorithm is to maximize the long-term discounted cumulative reward ($E[\sum_t \delta^t r_j^t]$), where $\delta < 1$ is the discount factor. The algorithm is forward-looking since it does not simply take the action that maximizes the current period reward. Instead, it also takes into account the effect of the current period action on future period payoffs through state transitions. The value to seller j given a state can be represented by the Bellman's value function

$$V_j(s^t) = \max_{p \in A} \{E[r_j^t | s^t, p] + \delta E[V(s^{t+1}) | s^t, p]\} \quad (3.2)$$

For the Q-learning agent, the long-term value of taking action p_j^t at state s^t is represented by a 'state-action value function' – Q-function

$$\begin{aligned} Q_j(s^t, p_j^t) &= E(r_j^t | s^t, p_j^t) + \delta V_j(s^{t+1}) \\ &= E(r_j^t | s^t, p_j^t) + \delta E[\max_{p \in A} Q(s^{t+1}, p) | s^t, p_j^t] \end{aligned} \quad (3.3)$$

If the state space (S) and action space (A) are finite, the Q-values can be stored in a $|S| \times |A|$ table. Given this Q-table, after observing state s^t , the algorithm will take the action in set A

that maximizes the Q-value. The core of the Q-learning algorithm is to learn the Q-values, which can be achieved by interacting with the environment repeatedly and updating the Q-table constantly. Specifically, at time t an action p_j^t is taken, the Q-value $Q_j(s^t, p_j^t)$ will be updated as follows, while the Q-values for other state-action combinations remain unchanged:

$$Q_j(s^t, p_j^t) = (1 - \alpha)Q(s^t, p_j^t) + \alpha(r_j^t + \delta \max_{p \in A} (Q_j(s^{t+1}, p))) \quad (3.4)$$

where the additional parameter α is the learning rate. Intuitively, the update rule moves the current Q-value α “steps” towards the current-period reward plus the discounted value of the next period state.

Since Q-learning uses a table to store Q-values. Thus it can only work in finite action space. We use the price grid constructed earlier as the action space for seller j : $A_j = \{p_{j,k_j}\}, k_j \in \{1, \dots, 10\}$. Throughout the paper, we allow the Q-learning algorithm to have a single-period memory. That is, the state consists of the last-period prices of all sellers. Specifically, $S = \{(p_{1,k_1}, p_{2,k_2}, \dots, p_{J,k_J}), k_1, k_2, \dots, k_J \in \{1, \dots, 10\}\}$. The size of the state space $|S| = 10^J$. Although seemingly restricted, the number of potential strategies with one-period memory is actually extremely large. Specifically, there are in total 10^{10^J} potential strategies for the Q-learning algorithm to choose from.

Like other learning algorithms, Q-learning faces the exploration-exploitation trade-off. On the one hand, it tries to pick the optimal action based on the current knowledge; on the other hand, it needs to explore other options that have not been visited to learn about their potential. A commonly used exploration strategy is called the ϵ -greedy strategy. It takes the greedy action with probability $1 - \epsilon$ and randomly picks an action from the action space with probability ϵ . In practice, ϵ is often set to decrease over time, which allows the algorithm to explore intensively at first, then take the greedy action once the learning process is completed. Unless otherwise mentioned, in this paper we set ϵ to decay exponentially over time: $\epsilon(t) = e^{-\omega t}$, where ω is a parameter that controls how fast ϵ decays.

The Q-learning algorithm needs a stopping rule. In the paper, given the size of the state

space and the action space, we set each session length to 500,000 iterations. Our ex-post check shows more than 99% of the sessions are stabilized when the session ends (i.e., the optimal actions do not change for all states). With regards to initialization, given that the learning algorithm possesses no prior knowledge of the competitor’s strategy, we follow Calvano et al. (2020) and initialize the Q-value of taking action p in state s as the long-term profit by charging price p assuming the competitor selects the price randomly.

Once the learning is finished, the Q-learning algorithm’s pricing policy Φ_Q can be specified as in Equation 3.5:

$$\Phi_Q : p_j^t = \arg \max_p Q(s^t, p), \forall t \tag{3.5}$$

3.3.3 Competition

We begin with the case of 2 sellers in the market ($J = 2$). We study two competitive scenarios: 1) when both sellers use a Q-learning algorithm (Q-Q scenario hereafter) and 2) when a seller uses a simple rule-based algorithm to compete with a Q-learning algorithm (Q-R scenario hereafter). Without loss of generality, we assume seller 1 always uses a Q-learning algorithm, and seller 2, the focal seller, uses Q-learning in the Q-Q scenario and the simple rule-based algorithms in the Q-R scenario. Algorithm 1 and Algorithm 2 provide the details of the competition in the Q-Q and Q-R scenarios in pseudo-code form, respectively.

3.4 Market Environment 1: Logit Demand Model

Our initial focus is on the market setting in which the demand model adheres to the canonical logit demand model. We use a_j to denote the product quality of the product j , We use μ to denote the level of horizontal differentiation among the products. The demand function is as follows:⁷

⁷We normalize the number of consumers to 1.

Algorithm 1 Q-learning seller v.s. Q-learning seller

Parameters: $\alpha \in (0, 1]$, $\delta \in (0, 1)$, $\omega \in (0, \infty)$

Initialize $Q_1(s, p)$, $Q_2(s, p)$, for all $s \in S$, $p \in A$

Initialize $t = 1$

Initialize s^1 **while not converge do**

end

$\epsilon \leftarrow e^{-\omega t}$

$p_1^t \leftarrow \begin{cases} \arg \max_p Q_1(s^t, p) & \text{with probability } 1 - \epsilon \\ \text{a random action} & \text{with probability } \epsilon \end{cases}$

$p_2^t \leftarrow \begin{cases} \arg \max_p Q_2(s^t, p) & \text{with probability } 1 - \epsilon \\ \text{a random action} & \text{with probability } \epsilon \end{cases}$

Calculate r_1^t, r_2^t

$s^{t+1} \leftarrow (p_1^t, p_2^t)$

$Q_1(s^t, p_1^t) \leftarrow (1 - \alpha)Q_1(s^t, p_1^t) + \alpha(r_1^t + \delta \max_{p \in A}(Q_1(s^{t+1}, p)))$

$Q_2(s^t, p_2^t) \leftarrow (1 - \alpha)Q_2(s^t, p_2^t) + \alpha(r_2^t + \delta \max_{p \in A}(Q_2(s^{t+1}, p)))$

$t \leftarrow t + 1$

Algorithm 2 Q-learning seller v.s. Simple rule-based seller

Parameters: $\alpha \in (0, 1]$, $\delta \in (0, 1)$, $\omega \in (0, \infty)$, Δp

Initialize $Q_1(s, p)$, for all $s \in S$, $p \in A$

Initialize $t = 1$

Initialize s^1 **while not converge do**

end

$\epsilon \leftarrow e^{-\omega t}$

$p_1^t \leftarrow \begin{cases} \arg \max_p Q_1(s^t, p_1^t) & \text{with probability } 1 - \epsilon \\ \text{a random action} & \text{with probability } \epsilon \end{cases}$

$p_2^t \leftarrow s^t(1) - \Delta p$

Calculate r_1^t, r_2^t

$s^{t+1} \leftarrow (p_1^t, p_2^t)$

$Q_1(s^t, p_1^t) \leftarrow (1 - \alpha)Q_1(s^t, p_1^t) + \alpha(r_1^t + \delta \max_{p \in A}(Q_1(s^{t+1}, p)))$

$t \leftarrow t + 1$

$$D_j^t = \frac{\exp(\frac{a_j - p_j^t}{\mu})}{\sum_j \exp(\frac{a_j - p_j^t}{\mu}) + \exp(\frac{a_0}{\mu})} \quad (3.6)$$

where a_0 is a parameter that captures the demand for the outside good. As mentioned before, the marginal cost of product j is c_j , the profit for the seller j in period t is $r_j^t = (p_j^t - mc_j)D_j^t$.

In the baseline model, we focus on the case where $J = 2$ and the two sellers are symmetric

(i.e., $mc_1 = mc_2$, $a_1 = a_2$). We assume that the demand function is unknown to either party, this assumption highlights an advantage of using a sophisticated learning algorithm over a simple rule-based algorithm: The optimal price cannot be calculated *ex ante*, but can be learned through repeated interactions.

Algorithm Parameters

For Q-learning algorithm hyper-parameters we choose $\alpha = 0.15$, $\omega = 1.5 \times 10^{-5}$. As for parameters for the demand function, in the baseline model, we set $a_1 = a_2 = 2$, $a_0 = 0$, $mc_1 = mc_2 = 1$, and $\mu = 0.25$. We use these configurations for the experiment unless otherwise mentioned⁸.

3.4.1 Experiment Outcomes: Baseline Setting

In this section, we present the results obtained from the experiment and compare the effectiveness of the Q-learning and the simple rule-based algorithms in the Q-Q and Q-R scenarios respectively on three aspects: 1) equilibrium prices and profits, 2) learning process and dynamics, and 3) the punishment strategies learned by the algorithms.

Price and Profit Comparison

For both the Q-Q scenario and Q-R scenario, we run the experiment 100 times using the baseline parameters. The prices charged by seller 2 (averaging over the last 1000 periods) and the corresponding single period profit for seller 2 (averaging over the last 1000 periods) are reported in Table 3.2. The values within parentheses represent the standard deviations across the 100 runs.

The results for the Q-Q scenario are similar to those shown in Calvano et al. (2020). Specifically, the two Q-learning algorithms learn to charge supra-competitive price (Bertrand-Nash price is 1.47), however, the price they converge to (around 1.79) is significantly lower

⁸In Section 3.4.2, we will extend the baseline model by varying the value of α , ω , a_0 , and μ

	Q-learning Algorithm	Rule-based Algorithm
Price	1.787(0.070)	1.870(0.000)
Profit	0.320(0.021)	0.361(0.000)

Table 3.2: Price and Profit Comparison for Focal Seller (Baseline)

Note: The column corresponding to Q-learning algorithm (Rule-based algorithm) represents the results for the focal seller in the Q-Q (Q-R) scenario.




Figure 3.1: Learning dynamics

Note: In the figure, the points on the solid lines represent the moving average (window size = 1000) of the average price across the 100 experiments at any t (number of iterations). The upper and lower bound of the shaded area represent the maximum and the minimum value within the moving window.

than the monopoly price (1.92). By contrast, in the Q-R scenario, the Q-learning algorithms quickly learn to charge a price close to the monopoly price, which benefits the simple rule-based algorithm adopter compared with the Q-Q scenario.

Learning Dynamic Comparison

After observing the distinctions in equilibrium price and profits in the two scenarios, we next compare their learning dynamics. From the 100 experiments in each scenario, we keep track of the prices charged by the algorithms in each period and report in Figure 3.1 the average price charged by the two sellers as a function of the number of iterations.

Learning Speed and Experiment Loss

We first compare the number of iterations needed for convergence between the Q-Q and Q-R scenarios. In the Q-Q scenario, roughly 350,000 iterations are needed for convergence. In the Q-R scenario, by contrast, only 40,000 iterations are needed. We next compare the loss incurred during the experimentation process. We take the equilibrium profit after prices converge as benchmark. In each period during experimentation, we compute the difference between the profit seller 2 gets and the benchmark profit, and we sum them up to measure the total loss in each scenario. It turns out that seller 2's experiment loss is on average 2.74 times higher in the Q-Q scenario compared with the Q-R scenario.

Learning Dynamics

Next, we examine the differences in learning processes between the two scenarios by utilizing the phase plot to illustrate the strategies employed by each seller. Specifically, we analyze the evolution of the strategy phase plots over time. The figure consists of four subfigures, each representing the strategy phase graph at a distinct learning stage, as indicated at the top of each sub-figure. The x-axis and y-axis in each phase graph represent the 10 different prices that can be chosen by seller 1 and seller 2, respectively. The 100 nodes in the graph represent the 100 states, which correspond to the price pair selected in the last period. For instance, the lower left node denotes the state in which both sellers chose Bertrand-Nash price in the previous period, while the upper right node represents the state where both sellers opted for the Monopoly price. The node size reflects the frequency at which a particular price pair is selected. Initially, all nodes are of equal size, as both algorithms choose prices randomly, resulting in all price pairs being selected equally frequently. The arrows in the graph denote the greedy strategies employed by the algorithms. That is, given a state, the arrow points to the price pair selected based on the Q values of the two algorithms. We observe that initially, all nodes point to the third-lowest node, which is the best response price when the competitor selects a price randomly.

When the learning starts, the algorithms transit from selecting prices randomly to favoring lower prices as the exploration rate diminishes. As the learning progresses to 39%, nodes




Figure 3.2: Learning Dynamics in the Q-Q scenario

located in the lower left corner are larger on average, indicating a higher frequency of selecting low-price pairs, which aligns with our expectations. Concurrently, as the competitor shifts from random price selection to favoring lower prices, the true Q values for states corresponding low prices decline. It is noteworthy that the Q values associated with low prices are updated more frequently and drop at a quicker rate than those for medium and high prices (since a Q-value is updated only when the corresponding action is taken, which happens more frequently for low prices). Eventually, the Q values for low prices dip below the Q values for

medium prices, prompting the algorithms' greedy strategies to switch from selecting low to medium prices. At the 51% mark, as shown in the third sub-figure, the large nodes move upward, indicating that the algorithms begin to coordinate on medium prices. At this stage, Q values for medium prices are updated more frequently, although they do not plummet as steeply as the Q values for low prices because the competitor adopts a less competitive strategy at this stage. Consequently, the Q values for medium prices never drop below those for high prices. As depicted in the final sub-figure, the upward trend eventually plateaus, resulting in both sellers repeatedly charging medium prices in the equilibrium.

The phase plots for the Q-R scenario are presented in Figure 3.3. In the Q-R scenario, the simple rule-based algorithm's strategy is fixed and the environment for the Q-learning algorithm is stable. The competitor faced by the Q-learning seller is not required to engage in random exploration, as observed in the Q-Q scenario. This advantageous feature enables the Q-learning algorithm to learn to coordinate on the highest price with remarkable speed.

Deviation and Punishment Strategy Comparison

To understand the mechanism driving the competition outcome discussed in the previous section, we study the strategies learned by the algorithms. For each experiment, we retrieve the two sellers' converged Q-tables to derive the sellers' strategies (i.e., what action to take in each state (s), which is $\arg \max_{a \in A} Q(s, a)$). We then force seller 1 to deviate from the learned strategy in one period. Seller 2 continues to play according to its learned strategy in this period, and in the subsequent period, both sellers play their learned strategies. We find that in most of the sessions, the two sellers will come back to the pre-deviation price within 10 moves (periods). We then plot the impulse-response functions derived from this exercise for all 100 experiment sessions ⁹. Specifically, shown in Figure 3.4a are the prices charged by the two sellers after the forced deviation of seller 1 in the Q-Q scenario. A similar plot is shown in Figures 3.4b and 3.4c for the Q-R scenario.

⁹In this experiment we focus on the sessions that do not lead to price cycles






Figure 3.3: Learning Dynamics in the Q-R scenario


We then take a closer look at the punishment strategies learned by the Q-learning algorithm in the Q-Q and Q-R scenarios. We use a directed graph to represent the limit strategies. Specifically, each node represents a state (a price pair), and each edge represents the state transition given the two firms' strategies. For example, an arrow pointing from node A to node B means that at the state specified in node A, the two algorithms' strategies will generate a price pair as specified in node B. The square is the absorbing state where there are no arrows pointing out. The shade of the nodes denotes the profit gain, darker



(a) Q-Q scenario



(b) Q-R scenario



(c) Q-R scenario

Figure 3.4: Deviation-Punishment scheme

The plot shows how prices evolve according to both sellers' limit strategies after an exogenous price cut for one of the sellers in 3 different scenarios. Figure (a): In the Q-Q scenario, seller 1 cuts the price by 5 grids in period 1. Figure (b): In the Q-R scenario, seller 1 (Q-learning algorithm) cuts the price by 5 grids in period 1. Figure (c): In the Q-R scenario, seller 2 (Simple rule-based algorithm) cuts the price by 5 grids in period 1.

nodes mean more collusive price pairs. The nodes' size represents the node's importance in the graph(as measured by betweenness centrality).




Figure 3.5: Graph of Strategies in Q-Q scenario

This is a visual representation of the strategies depicted through a directed graph. The square represents an absorbing node. All other nodes are represented by circles. The shades reflect the profit gain, with darker nodes indicating higher profit gains. The size of each node corresponds to its importance in the graph, as measured by betweenness centrality. In each graph, a blue line is used to highlight a specific deviation-punishment-reconcile path

In the Q-Q scenario, a deviation from the equilibrium state by seller 1 results in a return to the absorbing node rather than remaining stuck in another node. Thus, the punishment strategy in this scenario is not a grim trigger strategy. The graph illustrates a deviation-punishment-reconciliation path, represented by a blue line, which requires the two sellers to pass through several progressively darker nodes to return to the absorbing node. This




Figure 3.6: Graph of Strategies in Q-R scenario

strategy effectively punishes the deviating seller for multiple periods and eases the punishment over time representing a stick and carrot strategy.

The strategy graph in the Q-R scenario exhibits a distinct pattern: no matter which node the seller deviates to, it comes back to the equilibrium state almost immediately. The Q-learning algorithm employed in this scenario learns an ‘always cooperate’ strategy. The rationale for this is that the simple rule-based algorithm does not facilitate reconciliation unless the Q-learning algorithm brings the price back to the original level. The Q-learning algorithm picks this up from the past experience of deviation. Consequently, the algorithm increases the price immediately after a deviation to reduce the length of the punishment period and mitigate losses.

3.4.2 Experiment Outcomes: Alternative Settings

While the main intuition is shown in the baseline setting, we also work on several alternative settings to show the robustness of our results. In this section, we show how changes in 1), algorithm parameters, 2), demand structure and 3), different types of rule-based algorithms affect our main result. We put the study on multiple sellers and more advanced RL pricing algorithms in Appendix C.1.

Alternative Algorithm Settings

In the main model, algorithm parameters were initially set to their default values. This section extends the analysis by examining a broader spectrum of parameter values to assess the robustness of our findings. Specifically, we adjust three critical parameters within the Q-learning algorithm that are pivotal to its performance: the learning rate (α), the exploration decay rate (β), and the discount factor (δ).

For the learning rate α , although a typical value used in the Q-learning algorithm is 0.15, we investigate a range of values from 0.01 to 0.3 while keeping β and δ fixed at 1.5×10^{-5} and 0.95, respectively.

Regarding the exploration decay rate β , we examine values ranging from 0.1×10^{-5} to 2.8×10^{-5} while maintaining α at 0.15. To provide a better understanding of this parameter range, consider the following examples: with $\beta = 0.1 \times 10^{-5}$, after 250,000 iterations (half of the total number of iterations), the algorithm continues to explore (i.e., choose random actions) with a probability of 78%. Conversely, at the other extreme, with $\beta = 2.8 \times 10^{-5}$, the exploration probability has diminished to nearly 0 ($\epsilon = 0.1\%$).

As illustrated in Figure ??, the equilibrium profit for seller 2 in the QR scenario consistently surpasses the equilibrium profit for seller 2 in the QQ scenario across the explored parameter ranges. Figure ?? elaborates on the various permutations of α and β within the explored range, reinforcing our conclusion that the superior performance of the simple rule-based




Figure 3.7: Seller 2’s profit by using Q-learning algorithm and rule-based algorithm under different values of (a) α and (b) ω .

Note: The boxes in the figure represent the distributions of the seller 2’s profits from 100 experiments. The box extends from the first quartile to the third quartile, with a line at the median. The whiskers extend from the box by $0.5 \times$ the inter-quartile range (IQR).

algorithm over the Q-learning algorithm is not contingent upon specific parameter settings.




Figure 3.8: Percentage profit increase in seller 2’s profit when choosing Simple Rule-based algorithm over Q-learning algorithm.

Note: Each cell represents a permutation of α and ω . The values, represented by the shades of the color, show the percentage difference in profit percentage profit increase in seller 2’s profit when choosing the Simple Rule-based algorithm over the Q-learning algorithm.

We also explore how the discount factor δ alters the result. δ reflects the sellers’ or the

algorithms' time preference. We study how the patience of the seller affect the decision of choice of algorithms. In particular, while in the main model, we fix $\delta = 0.95$, here we explore a wider range of values. Figure ?? illustrates the result of comparison.




Figure 3.9: Seller 2's profit by using Q-learning algorithm and rule-based algorithm under different values of δ .

Note: The boxes in the figure represent the distributions of the seller 2's profits from 100 experiments. The box extends from the first quartile to the third quartile, with a line at the median. The whiskers extend from the box by $0.5 \times$ the inter-quartile range (IQR).

While the simple rule-based algorithm consistently outperforms Q-learning algorithms when the discount factor δ is large, this strict dominance vanishes when δ is small. In such cases, both the QQ and QR scenarios result in competitive equilibrium. This observation is logical, as supra-competitive prices are sustained by punishment strategies.

Without a sufficient level of patience, as represented by a lower δ value, the punishment strategies (whether deployed by the simple rule-based algorithm or the Q-learning algorithm) fail to effectively encourage the Q-learning competitor to learn cooperative behavior. Consequently, the Q-learning algorithm is less likely to converge to a collusive outcome when the discount factor is small, leading to competitive pricing in both the QQ and QR scenarios.

Different Rule Based Pricing Algorithms

We now examine the performance of other rule-based algorithms when they compete with the Q-learning algorithm. Specifically, we consider the following 3 types of competitor-based

algorithms and 2 fixed price strategies: 1) *Undercut Lowest Price with a Minimum Price (ULM)*: The agent using ULM always charges one price grid below the competitor’s last-period price but will never go below a user-defined minimum price p_{min} ¹⁰. 2) *Target Lowest Price (TL)*: This algorithm will match the competitor’s last-period price. 3) *Target Lowest Price with a Minimum Price (TLM)*: The agent using TLM always match the competitor’s last-period price but will never go below a user-defined minimum price p_{min} . 4) *Fixed Low Price (FL)*: The agent naively sets a fixed low price p_{low} ¹¹. 5) *Fixed High Price (FH)*: The agent naively sets a fixed high price p_{high} ¹².




Figure 3.10: (a) Price and (b) Profit comparison between QQ and QR scenario when seller 2 uses different Simple Rule-based algorithms.

Figure ?? presents a comparison of different rule-based algorithms versus a Q-learning algorithm in terms of both equilibrium price and equilibrium profit when the competitor employs a Q-learning algorithm. In the price comparison (Figure 3.10a), the dotted blue line represents the mean value of seller 2’s equilibrium price when choosing a Q-learning algorithm. The bars indicate the equilibrium prices when selecting different rule-based algorithms, with red bars representing the focal seller (seller 2) and gray bars representing the competing seller (seller 1).

¹⁰In this section, we set p_{min} to the 4th price grid

¹¹In this section, we set p_{low} to the 3rd price grid

¹²In this section, we set p_{high} to the 8th price grid

By comparing the red bars, we observe that all the rule-based algorithms considered here, except for the fixed low price (FL) strategy, result in a higher equilibrium price for the focal seller compared to the Q-learning algorithm. Furthermore, when examining the gray bars, we find that the competitor's prices in the three competitor-based algorithms are higher than those in the Q-learning algorithm, while the two fixed price strategies lead to lower competitor prices compared to the Q-learning algorithm. Consequently, as shown in Figure 3.10b, the three competitor-based algorithms achieve higher equilibrium profits than the Q-learning algorithm, whereas the two non-competitor-based algorithms yield lower equilibrium profits compared to the Q-learning algorithm.

The effectiveness of the simple rule-based algorithm lies not in its simplicity but in its ability to peg prices to the competitor's prices. When the competitor is a learning algorithm, it can effectively learn this pegging pattern and raise prices to reduce the intensity of competition. This simple pegging strategy resembles the tit-for-tat strategy developed in evolutionary game theory and exhibits a significant level of sophistication.

Alternative Market Structure

Horizontal Differentiation (Changes in μ):

According to the demand function specified in Equation 3.6, the degree of horizontal differentiation between the products is captured by the parameter μ . As μ approaches 0, the products become perfect substitutes, and as μ increases, the *cross price elasticity of demand* decreases, resulting in products that are more independent of each other. In the baseline scenario, we set $\mu = 0.25$, and now we consider five different values of μ ($\mu \in 0.1, 0.2, 0.3, 0.4, 0.5$). It is important to note that when μ changes, both the Bertrand-Nash price and the monopoly price will change accordingly.

Figure 3.11a presents the results of seller 2's profit under different levels of horizontal differentiation. As μ decreases (i.e., the degree of horizontal differentiation decreases), seller 2's equilibrium profit in the Q-Q scenario drops. This is because when the two products

become more similar, being undercut by competitors will result in lower short-term profits. Consequently, the algorithm will initially assign lower Q-values to high prices and higher Q-values to low prices, making it more challenging to coordinate on medium or high prices later on. However, this pattern is reversed in the Q-R scenario. In this case, the Q-learning algorithm is always able to charge the highest price, and the lower the value of μ , the higher the profit that a simple rule-based algorithm can achieve by consistently undercutting its competitor. Overall, the results demonstrate that the rule-based algorithm outperforms the Q-learning algorithm consistently across different levels of horizontal differentiation.




Figure 3.11: Seller 2's profit by using Q-learning algorithm and rule-based algorithm under different values of (a) μ and (b) Δa_0 .

Note: The boxes in the figure represent the distributions of the seller 2's profits from 100 experiments. The box extends from the first quartile to the third quartile, with a line at the median. The whiskers extend from the box by $0.5 \times$ the inter-quartile range (IQR).

Uncertainty in Demand (Changes in a_0):

To introduce additional demand uncertainty, we allow the quality of the 'outside good' to fluctuate from period to period. Specifically, in each period, the quality of the outside good (a_0) can take on either the value $-\Delta a_0$ or Δa_0 , each with a 50% probability. We conduct experiments with five different values of Δa_0 ($\Delta a_0 \in 0, 0.125, 0.25, 0.375, 0.5$) in this study. Similar to the previous case, when Δa_0 changes, both the Bertrand-Nash price and the monopoly price adjust accordingly.

Figure 3.11b presents the results of seller 2’s profit under different levels of demand uncertainty. The findings consistently show that the rule-based algorithm provides seller 2 with a higher equilibrium profit across all cases, regardless of the level of demand uncertainty introduced by the fluctuating quality of the outside good.

3.5 Market Environment 2: A Structural Demand Model

In addition to the stylized logit demand model, we also study more realistic demand models and see how algorithms behave under these alternative demand structures. In particular, in this section, we will work on a more sophisticated consumer non-sequential search demand model where model parameters are estimated from real-world data.

The non-sequential search model is constructed and estimated using transaction-level data from JD.com, a prominent Chinese e-commerce platform. The model consists of five vertically differentiated sellers and a unit mass of heterogeneous consumers. These consumers engage in a two-stage decision-making process: first, they decide which product to click on, and then they decide which product among the clicked ones to purchase.

Appendix C.2 provides details on the structure of the model. The incorporation of this non-sequential search model allows for a more realistic representation of consumer behavior in an online marketplace setting, capturing the inherent complexity of the decision-making process and the presence of multiple differentiated sellers.

Counterfactual: Algorithmic Pricing

Using the estimated model parameters, we study the competition of pricing algorithmic via counterfactuals. Specifically, based on the choice of the algorithms of the 5 sellers, we study two counterfactual scenarios: In counterfactual scenario 1, seller 1 uses a Q-learning pricing algorithm. Seller 3, 4, and 5 use a simple rule-based pricing algorithm¹³. In counterfactual


¹³Here again we use the ‘Undercut Lowest Price’ as a representative for simple rule-based algorithm

scenario 2, sellers 1 and 3 use Q-learning pricing algorithms. sellers 4 and 5 use simple rule-based algorithms. In both counterfactual scenarios, we compare seller 2’s equilibrium price/profit when the seller chooses between a Q-learning algorithm and a simple rule-based algorithm.


For both the simple rule-based and Q-learning pricing algorithms, we set the size of the price grids to 10 (i.e., there are 10 equally-spaced possible price levels), and these price grids are seller specific. Each seller’s price grid ranges from its Nash equilibrium prices (competitive price) to its monopoly prices. When a seller uses a simple rule-based algorithm, the algorithm undercuts the lowest price of all Q-learning sellers by one price level (e.g., if Q-learning seller 1 uses its 7th price level, and Q-learning seller 2 uses its 5th price level, the simple rule-based algorithm will set price to its 4th price level). We set other hyper-parameters, including learning rate α , exploration decay rate β , and discount factor δ , to those used in the the baseline case (in Section 3.4). We simulate the market dynamics in each of the two counterfactual scenarios under each of Seller 2’s possible algorithm choices (rule-based and Q-learning) 100 times.

The learning dynamics of the algorithms in the two scenarios are shown in Figures 3.12 and 3.13, and the mean and variance of prices and profits in the two scenarios are reported in Tables 3.3 and 3.4. In counterfactual scenario 1 (2), the focal seller’s profit will increase by 4.9% (3.9%) by using a simple rule-based algorithm instead of a Q-learning algorithm. In the last column of the two tables, we also report the price and profit of our focal seller if it uses the ‘optimal’ simple rule-based pricing algorithm, in which the algorithm undercut competitors’ prices by the number of price levels which results in the highest equilibrium profit. When using the ‘optimal’ rule-based strategy, the focal sellers’ profit will increase by 9.02% (6.8%) in comparison to that from using a Q-learning algorithm.

In counterfactual scenario 1, seller 2, along with other sellers who employ simple rule-based algorithms, is able to drive the price of seller 1’s product to near-monopoly levels. Seller 2 benefits from the relatively high average market price and consistently prices lower than




(a) Seller 2: Q-learning




(b) Seller 2: Simple rule-based

Figure 3.12: Learning dynamics: Counterfactual Scenario 1



(a) Seller 2: Q-learning



(b) Seller 2: Simple rule-based

Figure 3.13: Learning dynamics: Counterfactual Scenario 1

Seller 2:	Q-learning	Simple rule-based	Rule-based(optimal)
Price	57.6(1.2)	59.6(0.0)	56.8(0.0)
Profit	58052.9(1141.7)	60917.6(3.6)	63291.6(4.7)

Table 3.3: Price and Profit Comparison for Seller 2 (Counterfactual scenario 1)

Seller 2:	Q-learning	Simple Rule-based	Rule-based(optimal)
Price	56.9(0.6)	57.2(1.0)	55.7(1.0)
Profit	57994.8(1018.8)	60248.3(690.3)	61920.2(1212.6)

Table 3.4: Price and Profit Comparison for Seller 2 (Counterfactual scenario 2)

its main competitor, seller 1. However, if seller 2 were to adopt the Q-learning algorithm concurrently with seller 1, the non-stationarity introduced by simultaneous exploration would hinder both sellers from raising prices to monopoly levels, resulting in lower profits for seller 2.

In counterfactual scenario 2, the utilization of the Q-learning algorithm by both seller 1 and seller 3 creates a non-stationary environment. If seller 2 also implements the Q-learning algorithm, it would further exacerbate the non-stationarity, making it more challenging to search for a joint optimal price. This is because facing two competitors that explore randomly, as opposed to one, lowers the best response price. Consequently, the Q-learning algorithms assign higher Q-values to lower prices and lower values to high prices initially, making it difficult for the algorithm to move from low to middle or high prices. The graphical representation in Figure 3.13a demonstrates this issue, where initially, the prices sharply decline and struggle to rebound to the middle or high level. However, utilizing a simple rule-based algorithm by seller 2 can assist seller 1 and seller 3 in realizing the disadvantage of being stuck at a low price quickly. If they continue to charge low prices and, in addition, seller 2 charges an even lower price as a significant competitor, the two sellers would be incentivized to explore the high price region. This effect is evident in Figure 3.13b, where the initial price drop is much smaller than in Figure 3.13a, and the sellers increase the price to a much higher level in the later learning stage.

3.6 Theoretical Analysis

In the previous sections, we showed, through extensive simulations using the logit demand model and counterfactual analysis using an estimated demand model from real-world data, that simple rule-based algorithms provide higher profit compared to Q-learning algorithms to a seller when competing against others who use Q-learning algorithms. However, do these results hold under all conditions? To address this issue, we approach the problem theoretically and identify boundary conditions when simple rule-based algorithms would dominate Q-learning algorithms.

Modeling the RL pricing algorithm theoretically is challenging. The main source of complication is the stochastic and dynamic nature of RL algorithms. The updating process depends on the actions taken and the actions themselves may be determined by random draws. We develop a theoretical framework that captures the RL pricing dynamics to help understand how different equilibrium outcomes are reached in different competitive scenarios. In essence, the framework transforms the discrete-time learning problem into a continuous-time version, which allows us to characterize the Q-learning dynamics using a system of differential equations. This framework enables us to study theoretically how the ‘simultaneous exploration issue’ affects algorithms’ ability to coordinate on high prices and how stationarity brought by the rule-based algorithm solves this issue.

We put the details of the general theoretical framework in Appendix C.4. Here we present the result under a simple configuration. First, we limit the size of the price grid (m) to 2: $A = p_l, p_h$. Second, we assume both sellers fully explore ($\epsilon = 1$ when $t \leq t_0$) then fully exploit ($\epsilon = 0$ when $t > t_0$). Third, we use a simple linear demand function $D_i(p_i, p_{-i}) = 1 - p_i + bp_{-i}$, where $b \in (0, 1)$ capture the cross-price elasticity of the two goods.

Considering that the size of the price grid is 2, we set p_l to the Bertrand-Nash price and

p_h to the monopoly price of the stage game:

$$p_l = \frac{1}{2-b}, \quad p_h = \frac{1}{2-2b} \quad (3.7)$$

The stage game payoff matrix to seller 1 is shown in Table 3.5, where seller 1 is the row player and seller 2 is the column player.

	p_l	p_h
p_l	$r_{ll} = \frac{1}{(b-2)^2}$	$r_{lh} = \frac{-b^2 + 2b - 2}{2(b^3 - 5b^2 + 8b - 4)}$
p_h	$r_{hl} = \frac{3b-2}{4(b^3 - 4b^2 + 5b - 2)}$	$r_{hh} = \frac{1}{4(1-b)}$

Table 3.5: The Payoff Matrix of the Stage Game

The pricing strategies of sellers are determined by learned Q-values, and once learning has concluded, Q-values are examined to identify the equilibrium outcome. By examining the mapping of sellers' states to actions, we can identify the absorbing state(s) based on sellers' strategies. In the event that the Q-values for setting p_h are higher than those for setting p_l in any given state, the state (p_h, p_h) becomes the absorbing state, and both sellers will charge p_h in equilibrium. We put the details on the learning process (how Q-values evolve) in Appendix C.4 and present the main result here. In the remainder of this section, we use $Q_{(p_{k_1}, p_{k_2})-p_{k_0}}(t), \forall k_1, k_2, k_3 \in \{l, h\}$ denote seller's learned Q-values at time t that corresponding to charging price p_{k_0} in state (p_{k_1}, p_{k_2}) .

Lemma 9 *The equilibrium outcome in the Q-Q scenario is determined by δ and b , as follows:*

1. *If $\delta < \frac{b}{2-b}$, $\lim_{t \rightarrow \infty} Q_{(p_l, p_l)-p_l}(t) > \lim_{t \rightarrow \infty} Q_{(p_{k_1}, p_{k_2})-p_h}(t), \forall k_1, k_2 \in \{l, h\}$, both sellers charge p_l repeatedly as $t \rightarrow \infty$.*
2. *If $\frac{b}{2-b} \leq \delta < \sqrt{\frac{b}{2-b}}$, $\lim_{t \rightarrow \infty} Q_{(p_l, p_l)-p_l}(t) \leq \lim_{t \rightarrow \infty} Q_{(p_l, p_l)-p_h}(t)$, $\lim_{t \rightarrow \infty} Q_{(p_h, p_h)-p_l}(t) > \lim_{t \rightarrow \infty} Q_{(p_h, p_h)-p_h}(t)$, both sellers alternate between p_l and p_h as $t \rightarrow \infty$.*
3. *If $\delta \geq \sqrt{\frac{b}{2-b}}$, $\lim_{t \rightarrow \infty} Q_{(p_l, p_l)-p_l}(t) < \lim_{t \rightarrow \infty} Q_{(p_l, p_l)-p_h}(t)$, $\lim_{t \rightarrow \infty} Q_{(p_h, p_h)-p_l}(t) \leq \lim_{t \rightarrow \infty} Q_{(p_h, p_h)-p_h}(t)$, both sellers charge p_h repeatedly as $t \rightarrow \infty$.*

The proof of Lemma 9 is available in Appendix C.4.1. Lemma 9 highlights that both algorithms will cooperate only when they attach great importance to the future (δ is high) and when the two products are not too similar (b is low). The reasoning behind this is that the focal seller will benefit more from charging p_l than p_h since, in all states, the competitor randomly chooses between p_l and p_h . As exploration decreases, both sellers frequently charge p_l , causing the Q-values assigned to p_l in state (p_l, p_l) to drop. If the Q-value for p_l never falls below that of p_h , neither seller has the incentive to increase the price, and both sellers become stuck in state (p_l, p_l) . This situation arises when the initial value for p_h is extremely low, or equivalently when b is high and δ is low, making pricing above the competitor unprofitable in the short term. As soon as both sellers switch to charging p_h in state (p_l, p_l) , they alternate between state (p_l, p_l) and state (p_h, p_h) , and the Q-value assigned to charging p_l in state (p_h, p_h) decreases. If the Q-value for p_l drops below that of p_h in state (p_h, p_h) , both sellers will learn to charge p_h repeatedly. This happens when the initial value for p_h is high, or when b is low and δ is high, making pricing above the competitor profitable in the short term, or making the seller less concerned about short-term profits.

As for the Q-R scenario, we assume seller 1 uses a Q-learning algorithm, while seller 2 uses a simple rule-based algorithm and always set a price that equals seller 1's last period price. The Q-dynamics can be derived similarly (see Appendix C.4 for details). Using the same simple configurations as in Q-Q scenario, the final learning outcome is summarized in Lemma 10:

Lemma 10 *The equilibrium outcome in the Q-R scenario is determined by δ and b , as follows:*

1. If $\delta < \frac{1}{2}$, $\lim_{t \rightarrow \infty} Q_{(p_i, p_j) - p_l}(t) > \lim_{t \rightarrow \infty} Q_{(p_i, p_j) - p_h}(t)$, $\forall i, j \in \{l, h\}$, both sellers charge p_l repeatedly as $t \rightarrow \infty$.
2. If $\frac{1}{2} \leq \delta$, $\lim_{t \rightarrow \infty} Q_{(p_h, p_h) - p_l}(t) \leq \lim_{t \rightarrow \infty} Q_{(p_h, p_h) - p_h}(t)$, both sellers charge p_h repeatedly as $t \rightarrow \infty$.

The proof for Lemma 10 can be found in Appendix C.4.2.

By comparing the learning outcomes in the Q-Q and the Q-R scenario, Theorem 3 summarizes conditions under which sellers get higher equilibrium payoff in the Q-R scenario than in the Q-Q scenario.

Theorem 3 *Both sellers get weakly higher equilibrium profit in the Q-R scenario than in the Q-Q scenario if $\delta \geq \frac{1}{2}$.*

The proof can be found in Appendix C.4.3. Theorem 3 tells us that when facing a competitor who adopts a Q-learning pricing algorithm, the focal seller should adopt a simple rule-based algorithm unless the discount factor is smaller than a threshold 0.5. Although this threshold may not seem too small, it should be noted that the length between time periods is also small¹⁴. In practice, it is highly unlikely that any sellers will encounter a discount factor smaller than this value, which is why the previous experimental results consistently show that a simple rule-based algorithm outperforms a Q-learning algorithm.

3.7 Discussion and Conclusion

This study investigates the issue of algorithmic pricing competition, which has gained significant attention in recent times. Specifically, this paper aims to answer a crucial question: what kind of pricing algorithm should a seller use when competing with opponents who employ sophisticated learning algorithms? Our findings indicate that a simple rule-based algorithm can remarkably increase the seller's profit, even in the presence of competitors using sophisticated learning algorithms. Notably, implementing complex learning algorithms on online marketplaces can be both costly and risky, as it entails a prolonged exploration period, during which losses are likely to occur. Such a scenario may pose significant challenges for small sellers, who may not be well-equipped to bear the risks, thus raising concerns that they may be at a disadvantage compared to the pre-algorithmic pricing era. Nevertheless, our results suggest that such concerns are largely unfounded, as the use of a simple rule-

¹⁴For example, if the length between time periods is 4 hours, it transfers to a daily discount factor of 0.015.

based algorithm can yield better outcomes for a seller than utilizing a sophisticated learning algorithm when competing with competitors using advanced algorithms.

In the legal and policymaking spheres, algorithmic tacit collusion has been a topic of intense discussion since 2015, with policymakers expressing concern that independent learning algorithms have the potential to charge supra-competitive prices even in the absence of communication and coordination among sellers. Our analysis contradicts the conventional wisdom that more advanced algorithms are better equipped to collude and sustain higher prices. Instead, we find that the use of a sophisticated learning algorithm by one seller, in conjunction with simple rule-based algorithms used by others, can result in the sustainment of higher prices. Policymakers concerned with algorithmic collusion should, therefore, direct their attention toward this specific scenario, as it presents a unique challenge that demands tailored interventions.

Appendix A

Appendix for Chapter 1

A.1 Notation Summary

Notation	Meaning
α	The marginal effect of type on performance
β	The marginal effect of the causal feature on performance
λ	The portion of H type agents who are high on the correlational feature
θ	The portion of H type agents in the population
C_H	H type agents' cost of improving on the causal feature
C_L	L type agents' cost of improving on the causal feature
R	The job compensation paid by the firm once an agent is hired
State A	Low on the causal and high on the correlational feature in the opaque scenario
State B	High on the causal and high on the correlational feature in the opaque scenario
State C	Low on the causal and low on the correlational feature in the opaque scenario
State D	High on the causal and low on the correlational feature in the opaque scenario
State E	Low on the causal feature in the transparent scenario
State F	High on the causal feature in the transparent scenario

Table A.1: Notation summary

A.2 Discussions of Extensions of the Main Model

In this section, we analyze several extensions by relaxing some of the assumptions in our main model in Section 1.3. Due to space limitation, we will focus our discussions on the impact of

algorithmic transparency on the firm, and the main goal is to show that relaxing some of the assumptions does not alter the main results and insights of the study. Specifically, in Appendix A.2.1, we relax the assumption that the cost of improving the correlational feature is zero; in Appendix A.2.2, we relax the assumption that the wage is fixed and exogenously given; in Appendix A.2.3, we consider the case where agents have incorrect belief of λ . Unless otherwise noted, all the proofs of the results in this section can be found in Appendix A.4.

A.2.1 Gaming is Costly

In the main model, we have assumed that the cost of improving the correlation feature (gaming) is minimal regardless of the true type of the agent (we refer to this as the ‘costless’ setting). We now consider the case where the cost of improving the correlational feature is c_h for H type agents and c_l for L type agents with $c_l \geq c_h > 0$ (we refer to this as the ‘costly’ setting). This captures the situation where the correlational feature is not easy to game. For example, a Ph.D. program admission office might use the number of international conferences an applicant has attended in the past as a correlational feature, or a firm may use AI to analyze applicants’ behavior (e.g., pace or pitch of speaking) in the interview and use these verbal clues as correlational features. All these features require a certain amount of effort to game.

In the following, we first discuss the case where $c_h = c_l = c$ (we call this the ‘no-advantage’ case) and then we discuss the case where $c_h < c_l$ (we call this the ‘advantage’ case). Note that the change in the cost of improving the correlational feature has no impact on our analysis for the opaque scenario since, in the opaque scenario, the agents cannot manipulate their values on the correlational feature. As for the transparent scenario, we will show that the firm will get a weakly higher payoff in the ‘costly’ setting than in the ‘costless’ setting.

Recall that, in the ‘costless’ setting, all agents will eventually have value 1 on the correlational feature under the transparent algorithm; thus, States A and C collapse into a single State E whereas States B and D collapse into a single State F according to the

discussions in Section 1.3.2. In the ‘costly’ setting, some agents may decide not to improve their correlational feature even when the algorithm is transparent; therefore, agents’ possible ending states are A, B, C, and D as in the opaque scenario.

The firm’s equilibrium payoffs for the ‘no-advantage’ case are given below. Note that we divide the regions in the same way for Lemmas 6 and 7 below. However, the firm’s payoff in each region is different depending on the relative size of c with respect to R .

Lemma 11 *In the ‘costly’ setting, when $c \geq R$, the equilibrium outcome depends on the values of (C_H, C_L) , and this dependence is shown in Figure A.1. The corresponding total payoffs for the firm are given by*

$$\begin{aligned}\Pi_{1a}^{ex1} &= \lambda\theta(\alpha - R) - (1 - \lambda)(1 - \theta)R \\ \Pi_{2a}^{ex1} &= \theta(\alpha + \beta - R) \\ \Pi_{3a}^{ex1} &= \Pi_{3b}^{ex1} = \Pi_{3c}^{ex1} = \lambda\theta(\alpha + \beta - R) + (1 - \lambda)(1 - \theta)(\beta - R),\end{aligned}$$

where Π_i^{ex1} denotes the firm’s total payoff in region i . All payoffs are non-negative.

Lemma 12 *In the ‘costly’ setting, when $c < R$, the equilibrium outcome depends on the values of (C_H, C_L) , and this dependence is shown in Figure A.1. The corresponding total payoffs for the firm are given by*

$$\begin{aligned}\Pi_{1a}^{ex2} &= 0 \\ \Pi_{2a}^{ex2} &= \theta(\alpha + \beta - R) \\ \Pi_{3a}^{ex2} &= \theta(\alpha + \beta - R) + (1 - \theta)(\beta - R) \\ \Pi_{3b}^{ex2} &= \theta(\alpha + \beta - R) + (1 - \lambda)(1 - \theta)(\beta - R) \\ \Pi_{3c}^{ex2} &= \lambda\theta(\alpha + \beta - R) + (1 - \lambda)(1 - \theta)(\beta - R),\end{aligned}$$

where Π_i^{ex2} denotes the firm’s total payoff in region i . All payoffs are non-negative.

By directly comparing the firm’s total payoffs in Lemmas 11 and 12 with those in Lemma




Figure A.1: Equilibrium outcome in the ‘costly’ setting

2, it is not difficult to see that the firm is weakly better off in the ‘costly’ setting than in ‘costless’ setting. We formally summarize this observation in Proposition 10.

Proposition 10 *For the ‘no-advantage’ case, the firm gets weakly better payoffs in the ‘costly’ setting than in the ‘costless’ setting regardless of C_H , C_L , and c .*

In general, as c increases, the conditions under which algorithmic transparency is preferred become less strict (note that, for each region, the payoff in the ‘costly’ setting is at least as large as the payoff in the ‘costless’ setting). The intuition is as follows: If c is substantial, the firm might be able to exploit the information on the correlational feature without worrying about agents gaming the feature even when the algorithm is transparent. To be more specific, the firm can treat the agents in states $\{A, B\}$ and states $\{C, D\}$ differently without worrying that the agents in states $\{C, D\}$ will move to states $\{A, B\}$, reducing the degree of separation. Since the average productivity of agents in states $\{A, B\}$ is higher than the average productivity of the whole population, the firm can always get a higher payoff in the ‘costly’ setting than in the ‘costless’ setting by hiring a larger number of agents in states $\{A, B\}$ than in states $\{C, D\}$.

We next study the ‘advantage’ case, where the cost of improving on correlational feature for H type agents stays at $c_h = c$ while the cost for L type agents increases to $c_l > c$.

Proposition 11 summarizes the main findings:

Proposition 11 *The firm gets weakly better payoffs in the ‘advantage’ setting than in the ‘no advantage’ setting regardless of C_H , C_L , c_h and c_l .*

The intuition of the above result is straightforward: Increasing c_l could discourage L type agents from gaming when the algorithm is transparent, and this is always beneficial to the firm since agents in states A and B are on average more productive than agents in states C and D, and the firm always hire more agents from the former states (i.e. $P_A \geq P_C, P_B \geq P_D$). Discouraging L type agents from gaming will further increase the difference in agents’ quality (i.e., the proportion of H type agents) between states $\{A, B\}$ and states $\{C, D\}$ which, in turn, helps the firm to hire more H type agents while avoiding more L type agents.

To summarize, when the cost of improving the correlational feature is not zero, it can only strengthen the firm’s incentive to choose a transparent algorithm for any value of C_H and C_L .

A.2.2 Endogenous Wage

In the main model, we have shown that, when the wage R is fixed, the firm will be better off in making the algorithm transparent under certain conditions. A natural question is whether this result is primarily driven by the fixed wage assumption, that is, whether the result is driven by the fact that agents’ overall education and productivity increase in the transparent scenario and the firm can take full advantage of this increase without incurring any additional cost. If, in the opaque scenario, the firm can lower the wage while still keeping the agents’ strategies unchanged, it can offset some of the firm’s loss due to agents’ lower education level. In such a case, the equilibrium profit in the opaque scenario will become higher because the cost of hiring gets lower. Consequently, the opaque algorithm may become more attractive than the transparent algorithm and our result might be altered. Our objective in this subsection is to show that the above argument is only partly true: If we allow the firm to adjust the wage, while it is true that the equilibrium payoff in the opaque scenario can

be increased by lowering the wage, the firm can further lower the wage in the transparent scenario while making sure that it does not worsen the agents' degree of separation on the causal feature. As a result, making the algorithm transparent would still be preferable under certain conditions.

In the main model, the firm is restricted to pay the same wage R regardless of the decision on algorithmic transparency. In this extension we allow the firm to use R_O in the opaque scenario and R_T in the transparent scenario, where R_O and R_T are smaller than R ,¹ and we denote the optimal R_O and R_T as R_{O*} and R_{T*} respectively. We further assume that $R_O \geq R$ where $R = \max(\theta\alpha, \beta + \frac{\theta(1-\lambda)\alpha}{\theta(1-\lambda)+(1-\theta)\lambda})$, consistent with Assumptions 2 and 1.3.4. *The purpose of this extension is not to fully solve for the equilibrium in the endogenous wage setting, but to show that the major insights from our main model still hold even if the firm is allowed to strategically lower the wage R .*

In this discussion, we will only restrict our attention to the regions that have been previously identified to favor transparent scenario (i.e., regions $C1$, $C2$, and $C4$ defined in Section 1.4.3). Specifically, we assume the firm will prefer the transparent algorithm if the wage is fixed at R_0 , where $(\frac{C_L}{R_0}, \frac{C_H}{R_0})$ lands in regions $C1$, $C2$, and $C4$ defined in Section 1.4.3. We then let the firm lower the wage from R_0 to R_{O*} and R_{T*} in the opaque and transparent scenarios, respectively. We denote the firm's payoff under R_{O*} (R_{T*}) where R_0 is in region Ci in the opaque (transparent) scenario as Π_{Ci}^O (Π_{Ci}^T), $i \in \{1, 2, 4\}$.

Proposition 12 $\Pi_{Ci}^O \leq \Pi_{Ci}^T, \forall i \in \{1, 2, 4\}$.

Proposition 12 tells us that, for the cases that we have previously identified to favor the transparent scenario, the transparent scenario will still be preferred when the firm is allowed to strategically lower the wage. The intuition behind this proposition is as follows: In the opaque scenario, H type agents hold an advantage on the correlational feature and this will impede them from improving the causal feature. Only when the wage is high will they be incentivized to improve the causal feature and separate themselves from L type

¹Here we only consider the firm's decision on lowering R . The case where the firm strategically raises R is

agents. Such an advantage does not exist after the algorithm is made transparent so a lower wage is sufficient to separate different types of agents. Consequently, the firm is able to set a lower wage in the transparent scenario than in the opaque scenario without worrying of decreasing agents' *Dos* on the causal feature or decreasing the average value on the causal feature. Since agents' productivity (captured by α and β) is assumed to be unaffected by wage, the firm will benefit more from endogenizing the wage (due to lower a wage expense) in the transparent scenario than in the opaque scenario; thus, the firm's incentive to choose a transparent algorithm will be strengthened.

A.2.3 Agents Have Incorrect Belief of λ

In the main model we have assumed that agents have a correct belief of λ . We now relax this assumption and consider two possibilities: (1) Agents overestimate λ as λ_O and (2) Agents underestimate λ as λ_U . We assume that

$$0.5 < \lambda_U < \lambda < \lambda_O < 1.$$

Furthermore, similar to Assumptions 2 and 1.3.4, we assume

$$\frac{(\theta\lambda_U + (1 - \theta)(1 - \lambda_U))R}{\theta\lambda_U} < \alpha.$$

and

$$\frac{\theta(1 - \lambda_U)\alpha}{\theta(1 - \lambda_U) + (1 - \theta)\lambda_U} < R - \beta.$$

The above conditions ensure that, in agents' (incorrect) belief, the firm will have the incentive to include the correlational feature in the algorithm.

Note that the agents' belief about λ is only relevant in the opaque scenario. In the transparent scenario, the agents' equilibrium strategies are the same as discussed in Section

 similar but less straightforward.

1.4.2. In the opaque scenario, for both possibilities (1) and (2), it is the agents' belief of λ , not the true λ , that will determine their strategies on the causal feature. However, the true λ will determine the firm's payoff given the agents' strategies. There are still five cases as defined in Section 1.4.1. However, as shown in Figure A.2, the boundaries of these cases are now determined by λ_O or λ_U . Within each case, the firm's payoff function stays the same as in Section 1.4.1. The comparison between the opaque and transparent scenarios can be performed in a similar way as in Section 1.4: We can divide the area below the 45-degree line into seven regions, $N1$ to $N3$ and $C1$ to $C4$, as shown in Figure A.3.




Figure A.2: Equilibrium outcome when agents have an incorrect belief of λ

Similar to our results in Section 1.4, in regions $N1$ to $N3$, the firm will prefer the opaque algorithm. In region $C4$, the firm will prefer the transparent algorithm. In region $C1$ to $C3$, whether the firm prefers the transparent or opaque algorithm depends on the value of β . Generally speaking, when agents overestimate λ , under a wider range of conditions, the firm will prefer the transparent algorithm. This is reasonable: Overestimating the predictive power of the unknown correlational feature further prevents agents from competing on the causal feature in the opaque scenario. As previously discussed, making the algorithm transparent intensifies agents' competition on the causal feature, which benefits the firm. This benefit is larger when agents overestimate λ . Our result also shows that even if agents underestimate




Figure A.3: The comparison between transparent and opaque scenario when agents have incorrect belief on λ

λ , there are still cases in which the transparent algorithm is preferred by the firm.

A.3 Omitted Discussions

A.3.1 Derivation of the Lower and Upper Bound of α and β

In this paper we assume α , the marginal effect of agents' type being H on their productivity, to be in a certain range to eliminate uninteresting scenarios:

$$\frac{(\theta\lambda + (1 - \theta)(1 - \lambda))R}{\theta\lambda} < \alpha < \frac{R}{\theta}.$$

In the opaque scenario, when there is no agent who improves the causal feature, we want the firm to hire some agents based on the information in the correlational feature instead of not hiring anyone. (Not hiring anyone in this case is uninteresting because it will trivially drive everyone to improve the causal feature.) Thus we want α to be large enough to incentivize the firm to hire agents who have value 1 on the correlational feature. In the transparent scenario, when there is no agent who improves the causal feature, all

the agents are mixed together in the feature space: they all have the same values on both the causal and correlational features. The firm will either hire everyone or not hire anyone, depending on whether the average productivity of all the agents exceeds the salary or not. We want α to be small enough that the firm will not hire anyone in this case. (Hiring everyone in this case is uninteresting because no one will have the incentive to improve on the causal feature regardless of the cost of improving.) Specifically, rewrite the left inequality as $\theta\lambda\alpha + (1 - \theta)(1 - \lambda) \times 0 > (\theta\lambda + (1 - \theta)(1 - \lambda))R$. In the initial distribution of the opaque scenario (i.e., where everyone has a value of 0 on the causal feature), there are $\theta\lambda$ H type agents and $(1 - \theta)(1 - \lambda)$ L type agents who have value 0 on the causal feature and value 1 on the correlational feature. The inequality means that their total productivity (left hand side) should be larger than the total salary paid to them (right hand side). In other words, the firm has an incentive to hire all of these agents. If this is not the case, then the firm will not hire any agents with value 0 on the causal feature even if they have value 1 on the correlational feature, which will trivially incentivize the agents to improve the causal features. The right inequality means that in the transparent scenario where the correlational feature is gamed, if everyone has value 0 on the causal feature, the firm will not hire anyone.

In this paper we also assume β , the marginal effect of education on the agent's productivity, to be in a certain range to eliminate uninteresting scenarios:

$$R - \theta\alpha < \beta < R - \frac{\theta(1 - \lambda)\alpha}{\theta(1 - \lambda) + (1 - \theta)\lambda}.$$

Rewrite the left inequality as $\theta(\alpha + \beta) + (1 - \theta)\beta > (\theta + (1 - \theta))R$. In the transparent scenario, when everyone games on the correlational feature and everyone improves the causal feature, there are θ H type agents and $1 - \theta$ L type agents who have value 1 on both features. The inequality means that their total productivity (left-hand side) is larger than the total salary paid to them (right-hand side). In other words, the firm will have an incentive to hire all of them. If this is not the case, then in the transparent scenario, no one will improve on the causal feature and the firm will end up hiring nobody. As for the right part inequality,

rewrite it as $\theta(1 - \lambda)(\alpha + \beta) + (1 - \theta)\lambda\beta < (\theta(1 - \lambda) + (1 - \theta)\lambda)R$. In the opaque scenario, when no one games on the correlational feature but everyone improves the causal feature, there are $\theta(1 - \lambda)$ H type agents and $(1 - \theta)\lambda$ L type agents who have value 1 on the causal feature but value 0 on the correlational feature. This inequality means that these agents' total productivity (left-hand side) is smaller than the total wage paid to them (right-hand side). In other words, the firm will have no incentive to hire anyone of them. If this is not the case, then in the opaque scenario, improving the causal feature will guarantee an agent to be hired regardless of her value on the correlational feature, which will again, lead to an uninteresting equilibrium.

A.3.2 Discussion on the case where β does not satisfy Assumption

1.3.4

In Assumption 1.3.4, we restrict β in a certain range to avoid uninteresting cases. We now discuss how the relaxation of this assumption will affect our result. Specifically, we consider two cases (1) the case where $\beta \geq R - \frac{\theta(1-\lambda)\alpha}{\theta(1-\lambda)+(1-\theta)\lambda}$ (i.e., β exceeds the upper bound ($\bar{\beta}$) defined in Assumption 1.3.4) and (2) the case where $\beta \leq R - \theta\alpha$ (i.e., β is below the lower bound (β) defined in Assumption 1.3.4), this includes the case $\beta = 0$.

We first consider a case where $\beta \geq R - \frac{\theta(1-\lambda)\alpha}{\theta(1-\lambda)+(1-\theta)\lambda}$. The dependence of equilibrium outcomes on parameters in the transparent scenario is the same as shown in the original model. However, the dependence of equilibrium outcomes on parameters in the opaque scenario has changed, as shown in Figure A.4. The meaning of different cases are defined in Section 1.4.1.

Compared with the opaque scenario in our original model, relaxing this assumption will eliminate the two partial separating equilibria. As a result, the firm will be indifferent between transparent and opaque algorithm in the region denoted as case 3 in Figure A.4 (i.e., the case where all agents get education and the firm hires everyone). In other regions, the preference of transparency is exactly the same as our original model. This is a degenerate case of our




Figure A.4: Equilibrium outcome in the opaque scenario when $\beta \geq R - \frac{\theta(1-\lambda)\alpha}{\theta(1-\lambda)+(1-\theta)\lambda}$

original model.

We now consider a case where $\beta \leq R - \theta\alpha$. The dependence of equilibrium outcomes on parameters in the opaque scenario is the same as shown in the original model. However, the dependence of equilibrium outcomes on parameters in the transparent scenario has changed, as shown in Figure A.5.




Figure A.5: Equilibrium outcome in the transparent scenario when $\beta \leq R - \theta\alpha$

Compared with the transparent scenario in our original model, the previous pure strategy equilibrium case 3 will be replaced by a partial separating equilibrium case 5. As a result, the firm will get weakly higher payoff in the region denoted as case 5 in Figure A.5, since in

case 5 some L type agents do not get education and thus are separated apart from H type agents. The firm's incentive to choose transparent algorithm weakly increases in this region. In other regions, the preference of transparency is exactly the same as our original model.

A.3.3 Discussion on the Effect of θ and λ on Decision on Algorithmic Transparency Under the Stackelberg Model

The effect of θ and λ on the firm's choice on opacity and full transparency is summarized in Proposition 13.

Proposition 13 *An increase in λ has the following effects on the firm's decision on transparency:*

1. *The area of regions C1, C2, and C4 increases, which means that the transparent algorithm is preferred under more (C_H, C_L) value pairs.*
2. *Within regions C2, the conditions on β to make the transparent algorithm preferred to the opaque algorithm become stricter (i.e. a larger β is needed).*
3. *Within regions C2, C3, and N3 the conditions on C_L to make the transparent algorithm preferred to the opaque algorithm become stricter (i.e. a larger C_L is needed).*

An increase in θ has the following effects on the firm's decision on transparency:

1. *Within region C2, the conditions on β to make the transparent algorithm preferred to the opaque algorithm is not affected.*
2. *Within regions C2, C3, and N3 the conditions on C_L to make the transparent algorithm preferred to the opaque algorithm become stricter (i.e. a larger C_L is needed).*

Proposition 13 can be proved by checking

$$\begin{aligned}\frac{\partial C_L^{C2}}{\partial \lambda} &> 0, \frac{\partial C_L^{C2}}{\partial \theta} > 0 \\ \frac{\partial C_L^{C3}}{\partial \lambda} &> 0, \frac{\partial C_L^{C3}}{\partial \theta} > 0 \\ \frac{\partial C_L^{N3}}{\partial \lambda} &> 0, \frac{\partial C_L^{N3}}{\partial \theta} > 0,\end{aligned}$$

In general, as β and λ increase, the transparent algorithm becomes more attractive compared with the opaque algorithm. This is because when β and C_L increase, the firm's equilibrium payoff in the full transparency scenario increases sharply, while the equilibrium payoff in the opaque scenario is less affected by these two parameters.

A.4 Mathematical Appendix (Proofs of All Results)

A.4.1 Proof of Lemma 1

In the proof, we will proceed as follows: First, we identify different regions in the C_H - C_L space in which a certain class of strategies can be sustained as an equilibrium; then we analyze the corresponding payoffs for the firm and agents.

In the opaque scenario, agents move first and then the firm moves after observing individuals' actions. Even though the players in this game move sequentially, we can show that the PBE of this game coincides with the Bayesian Nash Equilibrium (BNE) of the corresponding simultaneously move game. The reasons are as follows: The strategies of all the agents will influence the firm's choice of algorithms. However, the influence of a single agent's action on the firm's strategy can be neglected. Once an equilibrium is reached, an agent (first mover) assesses the profitability of a deviation by assuming that the firm's (second mover) strategy will not change accordingly. In other words, all the PBEs that can be sustained in this scenario are actually the BNEs of the corresponding simultaneous move game.

Opaque scenario 1. We first look at the case where neither H type nor L agents improve education (i.e., $q_H = q_L = 0$). In this case, $\gamma_S^e = \gamma_S^b$. The firm's beliefs are: $\mu(T = H|S) = \gamma_S^e, \forall S \in A, C, \mu(T = H|S) = 1, \forall S \in B, D$. The firm's best response is: $P_A = P_B = P_D = 1, P_C = 0$.² For this strategy combination to be a PBE, the following conditions need to be satisfied:

$$\begin{aligned}
C_H &\geq (1 - \lambda)R && (H \text{ type agents will not deviate}) \\
C_L &\geq \lambda R && (L \text{ type agents will not deviate}) \\
\gamma_A^e &\geq \gamma_{th0} && (\text{The firm will not deviate on } P_A) \\
\gamma_C^e &\leq \gamma_{th0} && (\text{The firm will not deviate on } P_C).
\end{aligned}$$

(The first two conditions ensures that the H type and L type agents are better off (in terms of utility) not improving education, and the last two conditions ensure that the firm is better off (in terms of total payoffs) not deviating from its current P_A and P_C .) The first two conditions specify the regions in C_H - C_L space that this combination of strategies can be sustained as an equilibrium (Figure A.6 shows this region). The last two conditions are the direct consequences of Assumption 2:

$$\frac{\lambda\theta}{\lambda\theta + (1 - \lambda)(1 - \theta)} \geq \frac{R}{\alpha} = \gamma_{th0} \quad \frac{(1 - \lambda)\theta}{(1 - \lambda)\theta + \lambda(1 - \theta)} \leq \theta \leq \frac{R}{\alpha} = \gamma_{th0}.$$

The total payoffs for the firm and each type of agents are given by:

$$\begin{aligned}
\Pi_{firm_{O1}} &= \lambda\theta\alpha - (\lambda\theta + (1 - \lambda)(1 - \theta))R \\
\Pi_{H_{O1}} &= \theta\lambda R \\
\Pi_{L_{O1}} &= (1 - \theta)(1 - \lambda)R.
\end{aligned}$$

where we use $\Pi_{H_{O1}}$ and $\Pi_{L_{O1}}$ to denote the total payoff of H type and L type agents,

²Any $P_B \in [0, 1]$ and $P_D \in [0, 1]$ can be sustained since they are off the equilibrium path. Here we pick

respectively. (In the remaining of the proof, we will use Π_{H0i} and Π_{L0i} to denote the total payoff of H type and L type agents under case i , respectively.)




Figure A.6: Opaque scenario 1

Opaque scenario 2. In this case, only H type agents improve education but L type agents do not (i.e., $q_H = 1, q_L = 0$), and we have: $\gamma_B^e = \gamma_D^e = 1$ and $\gamma_A^e = \gamma_C^e = 0$. The firm's beliefs are: $\mu(T = H|S) = 0, \forall S \in A, C$, $\mu(T = H|S) = 1, \forall S \in B, D$. The firm's best response is $P_A = P_C = 0$, $P_B = P_D = 1$. For this strategy combination to be a PBE, the following conditions on the parameters need to be satisfied:

$$C_H \leq R \quad (H \text{ type agents will not deviate})$$

$$C_L \geq R \quad (L \text{ type agents will not deviate})$$

$$\gamma_B^e \geq \gamma_{th1}, \gamma_D^e \geq \gamma_{th1} \quad (\text{The firm will not deviate on } P_B \text{ and } P_D)$$

$$\gamma_A^e \leq \gamma_{th0}, \gamma_C^e \leq \gamma_{th0} \quad (\text{The firm will not deviate on } P_A \text{ and } P_C).$$

The first two conditions specify the regions (see Figure A.7) where this combination of strategies can be sustained as an equilibrium. The last two constraints are trivially satisfied.

The total payoffs for the firm and each type of agents are given by:

the maximum value of P_B and P_D to make sure the equilibrium could survive the intuitive criterion.

$$\Pi_{firmO_2} = \theta(\alpha + \beta) - \theta R$$

$$\Pi_{HO_2} = \theta(R - C_H)$$

$$\Pi_{LO_2} = 0.$$




Figure A.7: Opaque scenario 2

Opaque scenario 3. In this case, both H type and L type agents improve education (i.e., $q_H = q_L = 1$). The values of γ^e 's are given below:

$$\begin{aligned} \gamma_A^e &= \gamma_C^e = 0 \\ \gamma_B^e &= \frac{\lambda\theta}{\lambda\theta + (1-\lambda)(1-\theta)} \\ \gamma_D^e &= \frac{(1-\lambda)\theta}{(1-\lambda)\theta + \lambda(1-\theta)}. \end{aligned}$$

The firm's beliefs are: $\mu(T = H|S) = 0, \forall S \in A, C$, $\mu(T = H|S) = \gamma_S^e, \forall S \in B, D$. The firm's best response is to set $P_A = P_C = P_D = 0, P_B = 1$. For this strategy combination to

be a PBE, we need the following conditions:

$$\begin{aligned}
 C_H &\leq \lambda R && (H \text{ type agents will not deviate}) \\
 C_L &\leq (1 - \lambda)R && (H \text{ type agents will not deviate}) \\
 \gamma_B^e &\geq \gamma_{th1} && (\text{The firm will not deviate on } P_B) \\
 \gamma_D^e &\leq \gamma_{th1} && (\text{The firm will not deviate on } P_D).
 \end{aligned}$$

The first two conditions specify the regions (see Figure A.8) where this combination of strategies can be sustained as an equilibrium. The last two conditions are direct consequences of Assumption 1 and 2. The total payoffs for the firm and each type of agents are given by:

$$\begin{aligned}
 \Pi_{firm_{O3}} &= \lambda\theta(\alpha + \beta) + (1 - \lambda)(1 - \theta)\beta - (\lambda\theta + (1 - \lambda)(1 - \theta))R \\
 \Pi_{H_{O3}} &= \theta\lambda R - \theta C_H \\
 \Pi_{L_{O3}} &= (1 - \theta)(1 - \lambda)R - (1 - \theta)C_L.
 \end{aligned}$$




Figure A.8: Opaque scenario 3

Opaque scenario 4. In this case, H type agents improve education with probability q_H and

L type agents do not improve education (i.e., $q_L = 0$). The values of γ^e 's are given below:

$$\begin{aligned}\gamma_A^e &= \frac{\theta(1 - q_H)\lambda}{\theta(1 - q_H)\lambda + (1 - \theta)(1 - \lambda)} \\ \gamma_C^e &= \frac{\theta(1 - q_H)(1 - \lambda)}{\theta(1 - q_H)(1 - \lambda) + (1 - \theta)\lambda} \\ \gamma_B^e &= \gamma_D^e = 1.\end{aligned}$$

The firm's beliefs are: $\mu(T = H|S) = \gamma_S^e, \forall S \in A, C$, $\mu(T = H|S) = 1, \forall S \in B, D$. The firm's best response is to use $P_A = p_4, P_C = 0, P_B = P_D = 1$. For this strategy combination to be a PBE, the following conditions should hold:

$$\begin{aligned}C_H &= (1 - \lambda p_4)R && (H \text{ type agents are indifferent}) \\ C_L &\geq (1 - (1 - \lambda)p_4)R && (L \text{ type agents will not deviate}) \\ \gamma_A^e &= \gamma_{th0} && (\text{The firm is indifferent on } P_A) \\ \gamma_C^e &\leq \gamma_{th0} && (\text{The firm will not deviate on } P_C).\end{aligned}$$

The last condition is satisfied following Assumption 2:

$$\gamma_C^e = \frac{\theta(1 - q_H)(1 - \lambda)}{\theta(1 - q_H)(1 - \lambda) + (1 - \theta)\lambda} \leq \frac{(1 - \lambda)\theta}{(1 - \lambda)\theta + \lambda(1 - \theta)} \leq \frac{R}{\alpha} = \gamma_{th0}.$$

The first condition could be used to express p as a function of the other parameters and, similarly, the third condition can be used to express q_H as a function of the other parameters. Specifically, we have:

$$p_4 = \frac{1}{\lambda} \left(1 - \frac{C_H}{R} \right) \tag{A.1}$$

$$q_H = 1 - \frac{R(1 - \theta)(1 - \lambda)}{(\alpha - R)\theta\lambda}. \tag{A.2}$$

Given the fact that p and q_H are values between 0 and 1, we can calculate the range for C_H

and C_L in which this combination of strategies can be sustained as an equilibrium (see Figure A.9):

$$(1 - \lambda)R \leq C_H \leq R$$

$$\frac{R - C_H}{R - C_L} \leq \frac{\lambda}{1 - \lambda}.$$

where the first inequality follows from the first condition and the second inequality follows from the second condition.




Figure A.9: Opaque scenario 4

The total payoffs for the firm and each type of agents are given by:

$$\Pi_{firmO4} = \theta q_H(\alpha + \beta) - \theta q_H R$$

$$\Pi_{HO4} = \theta(R - C_H)$$

$$\Pi_{LO4} = p_4(1 - \theta)(1 - \lambda)R.$$

Opaque scenario 5. In this case, H type agents improve education and L type agents

improve education with probability q_L (i.e., $q_H = 1$). We have:

$$\begin{aligned}\gamma_A^e &= \gamma_C^e = 0 \\ \gamma_B^e &= \frac{\theta\lambda}{\theta\lambda + (1-\theta)q_L(1-\lambda)} \\ \gamma_D^e &= \frac{\theta(1-\lambda)}{\theta(1-\lambda) + (1-\theta)q_L\lambda}.\end{aligned}$$

The firm's beliefs are: $\mu(T = H|S) = 0, \forall S \in A, C$, $\mu(T = H|S) = \gamma_S^e, \forall S \in B, D$. The firm's best response is to use $P_A = P_C = 0, P_B = 1, P_D = p_5$. For this strategy combination to be a PBE, the following conditions should hold:

$$\begin{aligned}C_H &\leq (\lambda + p_5(1-\lambda))R && (H \text{ type agents will not deviate}) \\ C_L &= ((1-\lambda) + p_5\lambda)R && (L \text{ type agents are indifferent}) \\ \gamma_B^e &\geq \gamma_{th1} && (\text{The firm will not deviate on } P_B) \\ \gamma_D^e &= \gamma_{th1} && (\text{The firm will not deviate on } P_D).\end{aligned}$$

The second condition can be used to express p as a function of the other parameters while the last condition can be used to express q_L as a function of the other parameters. Specifically,

$$p_5 = \frac{C_L - R}{\lambda R} + 1 \tag{A.3}$$

$$q_L = \frac{(\alpha + \beta - R)\theta(1-\lambda)}{(1-\theta)\lambda(R-\beta)}. \tag{A.4}$$

The third condition is satisfied following Equations 2 and 1.3.4:

$$\gamma_B^e = \frac{\theta\lambda}{\theta\lambda + (1-\theta)q_L(1-\lambda)} \geq \gamma_{th0}.$$

Given the fact that p and q_L are values between 0 and 1, we can calculate the range for C_H

and C_L in which this combination of strategies can be sustained as an equilibrium:

$$(1 - \lambda)R \leq C_L \leq R$$

$$\frac{R - C_H}{R - C_L} \geq \frac{1 - \lambda}{\lambda}.$$




Figure A.10: Opaque scenario 5

The total payoffs for the firm and each type of agents are given by:

$$\begin{aligned} \Pi_{firm_{O5}} &= \theta\lambda(\alpha + \beta - R) + (1 - \theta)q_L(1 - \lambda)(\beta - R) \\ &= \frac{2\lambda - 1}{\lambda}\theta(\alpha + \beta - R) \\ \Pi_{H_{O5}} &= (\theta\lambda + \theta(1 - \lambda)p_5)R - \theta C_H \\ \Pi_{L_{O5}} &= 0. \end{aligned}$$

Dealing with multiple equilibria. Per our analysis of the above five cases, there are several regions where multiple equilibria exist. According to the dynamics of the game, in the opaque scenario, the agents move first and the firm moves next. Thus, we first narrow down to the equilibria that survive the intuitive criterion (Cho and Kreps, 1987) and then

select the equilibrium outcome which gives the largest total utilities for each agent type.³ (In theory, finding such an equilibrium is not always possible; fortunately, it is possible in our case.)

- In the region where Case 4 and Case 5 overlap, Case 4 always gives higher payoffs to both H type and L type agents:

$$\begin{aligned}\Pi_{HO4} &= \theta(R - C_H) > (\theta\lambda + \theta(1 - \lambda)p_5)R - \theta C_H = \Pi_{HO5} \\ \Pi_{LO4} &= p_4(1 - \theta)(1 - \lambda)R > 0 = \Pi_{LO5}.\end{aligned}$$

The inequalities hold since $p_4 = \frac{R - C_H}{\lambda R}$ and $p_5 = \frac{C_L - R}{\lambda R} + 1$ are values between 0 and 1.

- In the region where Case 4 and Case 1 overlap, Case 1 always gives higher payoffs to both H type and L type agents:

$$\begin{aligned}\Pi_{HO1} &= \theta\lambda R \geq \theta(R - C_H) = \Pi_{HO4} \\ \Pi_{LO1} &= (1 - \theta)(1 - \lambda)R \geq p_4(1 - \theta)(1 - \lambda)R = \Pi_{LO4}.\end{aligned}$$

The inequalities hold since $\frac{C_H}{R} \geq 1 - \lambda$ in the overlapped region and $p_4 = \frac{R - C_H}{\lambda R}$ is between 0 and 1.

- In the region where Case 1 and Case 2 overlap, Case 1 always gives higher payoffs to both H type and L type agents:

$$\begin{aligned}\Pi_{HO1} &= \theta\lambda R \geq \theta(R - C_H) = \Pi_{HO2} \\ \Pi_{LO1} &= (1 - \theta)(1 - \lambda)R \geq 0 = \Pi_{LO2}.\end{aligned}$$

The first inequality holds since $\frac{C_H}{R} \geq 1 - \lambda$ in the overlapped region.

- In the region where Case 1 and Case 5 overlap, Case 1 always gives higher payoffs to

³We select the equilibrium that gives the largest total utility for each agent type, which shares the same spirit of the undefeated criterion introduced by Mailath et al. (1993). (See also (Schmidt and Buell, 2017))

both H type and L type agents:

$$\Pi_{HO1} \geq \Pi_{HO4} \geq \Pi_{HO5}$$

$$\Pi_{LO1} \geq \Pi_{LO4} \geq \Pi_{LO5}.$$

The first inequality holds since $\frac{C_H}{R} \geq 1 - \lambda$ in the overlapped region.

A.4.2 Proof of Lemma 2

Similar to the proof of the opaque scenario, we proceed by analyzing the same five cases analyzed in the opaque scenario. We show that only the equilibrium outcomes corresponding to cases 1 to 3 are sustainable. Recall that in Lemma 2 we study the case where the firm does not have commitment power, thus the firm's action in the first stage (announcing the hiring probabilities) is not binding. Similar to what we have done in Lemma 1, we can show the PBEs of this game coincide with the BNEs of the corresponding simultaneous move game.

Transparent scenario 1. We first look at the case where neither H type nor L type agents improve education (i.e., $q_H = q_L = 0$). In this case, $\gamma_E^e = \theta$ and $\gamma_F^e = 0$. The firm's beliefs are: $\mu(T = H|E) = \theta$, $\mu(T = H|F) = 1$. The firm's best response is: $P_E = 0$, $P_F = 1$. For this strategy combination to be a PBE, the following conditions should be satisfied:

$$C_H \geq R \quad (H \text{ type agents will not deviate})$$

$$C_L \geq R \quad (L \text{ type agents will not deviate})$$

$$\gamma_E^e \leq \gamma_{th0} \quad (\text{The firm will not deviate on } P_E).$$

The last condition is the direct consequence of Assumption 2:

$$\gamma_E^e = \theta < \frac{R - \beta}{\alpha} < \frac{R}{\alpha} = \gamma_{th0}.$$

The first two conditions specify the regions in C_H - C_L space that this combination of strategies can be sustained as an equilibrium (Figure A.11 shows this region). The total payoffs for the firm and each type of agents are given by:

$$\Pi_{firm_{T1}} = 0$$

$$\Pi_{H_{T1}} = 0$$

$$\Pi_{L_{T1}} = 0.$$

where we use $\Pi_{H_{T1}}$ and $\Pi_{L_{T1}}$ to denote the total payoff of H type and L type agents, respectively. (In the remaining of the proof, we will use $\Pi_{H_{Ti}}$ and $\Pi_{L_{Ti}}$ to denote the total payoff of H type and L type agents under transparent case i , respectively.)




Figure A.11: Transparent scenario 1

Transparent scenario 2. In this case, only H type agents improve education but L type agents do not (i.e., $q_H = 1, q_L = 0$), and we have: $\gamma_E^e = 0$ and $\gamma_F^e = 1$. The firm's beliefs are: $\mu(T = H|E) = 0, \mu(T = H|F) = 1$. The firm's best response is $P_E = 0, P_F = 1$. For this strategy combination to be a PBE, the following conditions on the parameters need to be

satisfied:

$$\begin{aligned}
 C_H &\leq R && (H \text{ type agents will not deviate}) \\
 C_L &\geq R && (L \text{ type agents will not deviate}) \\
 \gamma_E^e &\leq \gamma_{th0} && (\text{The firm will not deviate on } P_E) \\
 \gamma_F^e &\geq \gamma_{th1} && (\text{The firm will not deviate on } P_F).
 \end{aligned}$$

The last two conditions are trivially satisfied (by Assumption 1, we have $0 < \gamma_{th1} < \gamma_{th0} < 1$). The first two conditions specify the regions in C_H - C_L space that this combination of strategies can be sustained as an equilibrium (Figure A.12 shows this region). The total payoffs for the firm and each type of agents are given by:

$$\begin{aligned}
 \Pi_{firmT2} &= \theta(\alpha + \beta - R) \\
 \Pi_{HT2} &= \theta(R - C_H) \\
 \Pi_{LT2} &= 0.
 \end{aligned}$$




Figure A.12: Transparent scenario 2

Transparent scenario 3. In this case, both H type and L type agents improve education

(i.e., $q_H = q_L = 1$), and we have: $\gamma_E^e = 0$ and $\gamma_F^e = \theta$. The firm's beliefs are: $\mu(T = H|E) = 0$, $\mu(T = H|F) = \theta$. The firm's best response is $P_E = 0$, $P_F = 1$. For this strategy combination to be a PBE, the following conditions on the parameters need to be satisfied:

$$\begin{aligned}
 C_H &\leq R && (H \text{ type agents will not deviate}) \\
 C_L &\leq R && (L \text{ type agents will not deviate}) \\
 \gamma_F^e &\geq \gamma_{th1} && (\text{The firm will not deviate on } P_F).
 \end{aligned}$$

The third condition is a direct consequence of Assumption 1.3.4. The first two conditions specify the regions in C_H - C_L space that this combination of strategies can be sustained as an equilibrium (Figure A.13 shows this region). The total payoffs for the firm and each type of agents are given by:

$$\Pi_{firm_{T3}} = \theta(\alpha + \beta) + (1 - \theta)\beta - R$$

$$\Pi_{HT3} = \theta(R - C_H)$$

$$\Pi_{LT3} = (1 - \theta)(R - C_L).$$




Figure A.13: Transparent scenario 3

Transparent scenario 4. In this case, H type agents improve education with probability q_H and L type agents do not improve education (i.e., $q_L = 0$). The values of γ^e 's are given below:

$$\begin{aligned}\gamma_E^e &= \frac{(1 - q_H)\theta}{(1 - q_H)\theta + (1 - \theta)} \\ \gamma_F^e &= 1.\end{aligned}$$

The firm's beliefs are: $\mu(T = H|E) = \gamma_E^e$, $\mu(T = H|F) = 1$. The firm's best response is $P_E = p$, $P_F = 1$. For this strategy combination to be a PBE, the following conditions on the parameters need to be satisfied:

$$\begin{aligned}C_H &= (1 - p)R && (H \text{ type agents are indifferent}) \\ C_L &\geq (1 - p)R && (L \text{ type agents will not deviate}) \\ \gamma_E^e &= \gamma_{th0} && (\text{The firm is indifferent on } P_E).\end{aligned}$$

Since $q_H \in [0, 1]$, $0 < \gamma_E^e < \theta$. The last condition requires $0 < \frac{R}{\alpha} < \theta$, or equivalently $\alpha > \frac{R}{\theta}$. In the range of α that we are considering (i.e., $\frac{(\theta\lambda + (1-\theta)(1-\lambda))R}{\theta\lambda} < \alpha < \frac{R}{\theta}$, by Assumption 2), this combination of strategies cannot be sustained as an equilibrium.

Transparent scenario 5. In this case, H type agents improve education and L type agents improve education with probability q_L (i.e., $q_H = 1$). The values of γ^e 's are given below:

$$\begin{aligned}\gamma_E^e &= 0 \\ \gamma_F^e &= \frac{\theta}{\theta + (1 - \theta)q_L}.\end{aligned}$$

The firm's beliefs are: $\mu(T = H|E) = 0$, $\mu(T = H|F) = \gamma_F^e$. The firm's best response is $P_E = 0$, $P_F = p$. For this strategy combination to be a PBE, the following conditions on the

parameters need to be satisfied:

$$\begin{aligned}
C_H &\leq pR && (H \text{ type agents will not deviate}) \\
C_L &= pR && (L \text{ type agents are indifferent}) \\
\gamma_F^e &= \gamma_{th1} && (\text{The firm is indifferent on } P_F).
\end{aligned}$$

Note that, the second condition implies $p = \frac{C_L}{R}$. Given that $p \in [0, 1]$, any value of $C_L \in [0, R]$ is valid. As for the last condition, $q_L \in [0, 1]$ implies $\theta < \gamma_F^e < 1$. However, by Assumption 1.3.4, $\beta > R - \theta\alpha$, which implies $\gamma_{th1} = \frac{R-\beta}{\alpha} < \theta$. Thus, the last condition cannot be satisfied and, therefore, this combination of strategies cannot be sustained as an equilibrium.

A.4.3 Proof of Lemma 3

The proof of Lemma 3 is straightforward: $\Pi_{agents_{O_i}} = \Pi_{H_{O_i}} + \Pi_{L_{O_i}} \forall i \in \{1, 2, 3, 4, 5\}$, $\Pi_{agents_{T_i}} = \Pi_{H_{T_i}} + \Pi_{L_{T_i}} \forall i \in \{1, 2, 3\}$, where $\Pi_{H_{O_i}}$, $\Pi_{L_{O_i}}$, $\Pi_{H_{T_i}}$, $\Pi_{L_{T_i}}$ are defined in Appendix A.4.1 and A.4.2.

A.4.4 Proof of Lemma 4 and Lemma 5

We start from the equilibrium strategy in the partial transparency case and see whether there exists an ex-post sub-optimal strategy that the firm can commit to that can further increase the firm's payoff. In region 1 of Figure 1.5, any commitment on hiring strategy will not change agents' behavior on the causal feature, because in this region, neither type of agents will improve education due to the large values of C_H and C_L . Therefore, the firm will not commit to an ex-post sub-optimal strategy. The equilibrium payoff in this region is equal to that in the partial transparency case $\Pi_{FT1} = 0$. In region 2, the firm has no incentive to commit to an ex-post sub-optimal strategy because the H type and L type agents are already perfectly separated on the causal feature, and the firm has achieved the maximum payoff. Thus the equilibrium payoff in this region is $\Pi_{FT2} = \theta(\alpha + \beta - R)$. However, in region 3, the

firm may have an incentive to commit to a strategy that is ex-post sub-optimal. Given a (C_H, C_L) combination in this region, the firm could announce hiring probabilities $P_E = 0$ and $P_F = \frac{C_L}{R}$.⁴ Observing this hiring strategy, H type agents improve education but L type agents choose not to. The firm then follows the pre-announced hiring strategy and hires $\frac{C_L}{R}$ portion of H type agents in State F , and this will result in a payoff of $\Pi_{Stackelberg} = \frac{C_L \theta (\alpha + \beta - R)}{R}$. This payoff is greater than the equilibrium payoff when the firm does not have the commitment power (i.e., $\Pi_{firm_{T3}} = \theta(\alpha + \beta) + (1 - \theta)\beta - R$ in Lemma 2) if and only if $\frac{C_L}{R} \geq p_0$ where $p_0 = \frac{\theta\alpha + \beta - R}{\theta(\alpha + \beta - R)}$. The region in which the condition $\frac{C_L}{R} \geq p_0$ holds is denoted as region S in Figure 1.5.

As for the agents welfare, it is easy to check that in region S , all H type agents improve education but only $\frac{C_L}{R}$ fraction of them will be hired, thus $\Pi_{agents_{FTS}} = \theta(C_L - C_H)$. In other regions, agents welfare stays the same as in the partial transparency case because neither agents' nor the firm's strategy changes.

A.4.5 Proof of Lemma 11 and Lemma 12

We start with the 'no-advantage case'. We first consider the case where $c \geq R$. In this case, no agents will improve the correlational feature. If $C_H > R$ and $C_L > R$, the firm sets $P_A = 1$, $P_B = P_C = P_D = 0$, no agents will move, and the equilibrium payoff to the firm is given by $\Pi_{1a}^{ex1} = \lambda\theta(\alpha - R) - (1 - \lambda)(1 - \theta)R$. If $C_H \leq R$ and $C_L > R$, the firm sets $P_A = P_C = 0$, $P_B = P_D = 1$, all H type agents improve the causal feature, and the equilibrium payoff to the firm is $\Pi_{2a}^{ex1} = \theta(\alpha + \beta - R)$. If $C_H \leq R$ and $C_L \leq R$, the firm sets $P_B = 1$, $P_A = P_C = P_D = 0$, all agents in state A move to state B, and the equilibrium payoff to the firm is $\Pi_{3a}^{ex1} = \Pi_{3b}^{ex1} = \Pi_{3c}^{ex1} = \lambda\theta(\alpha + \beta - R) + (1 - \lambda)(1 - \theta)(\beta - R)$.

We next consider the 'no-advantage case' with $c < R$. We identify an equilibrium for each of the 5 regions denoted in Figure A.1. In region 1a, the firm sets $P_A = \frac{c}{R}$, $P_C = 0$, $P_B \in [0, 1]$, $P_D \in [0, 1]$, a fraction of agents in State C move to State A to make the firm

⁴Here we only consider the case where $\theta < \frac{R}{\alpha + \beta}$, the results in for the case where $\theta \geq \frac{R}{\alpha + \beta}$ could be derived

indifferent between hiring and not hiring in State A. The equilibrium payoff to the firm is given by $\Pi_{1a}^{ex2} = 0$. In region 2a, the firm sets $P_A = P_C = 0$, $P_B = P_D = 1$, all H type agents improve the causal feature, and the equilibrium payoff to the firm is $\Pi_{2a}^{ex2} = \theta(\alpha + \beta - R)$. In region 3a, the firm sets $P_A = P_C = 0$, $P_B = P_D = 1$, all agents improve the causal feature, and the equilibrium payoff to the firm is $\Pi_{3a}^{ex2} = \theta(\alpha + \beta - R) + (1 - \theta)(\beta - R)$. In region 3b, the firm sets $P_A = P_C = P_D = 0$, $P_B = 1$, H type agents in State A and C move to State B, and the equilibrium payoff to the firm is $\Pi_{3b}^{ex2} = \theta(\alpha + \beta - R) + (1 - \lambda)(1 - \theta)(\beta - R)$. In region 3c, the firm sets $P_A = P_C = P_D = 0$, $P_B = 1$, H type agents in State A move to State B, and the equilibrium payoff to the firm is $\Pi_{3c}^{ex2} = \lambda\theta(\alpha + \beta - R) + (1 - \lambda)(1 - \theta)(\beta - R)$.

By comparing the equilibrium payoffs in the ‘costly’ and the ‘costless’ setting, we have $\Pi_{1a}^{ex1} > \Pi_{firmT1}$, $\Pi_{2a}^{ex1} = \Pi_{firmT2}$, $\Pi_{3a}^{ex1} = \Pi_{3b}^{ex1} = \Pi_{3c}^{ex1} > \Pi_{firmT3}$. Moreover, $\Pi_{1a}^{ex2} = \Pi_{firmT1}$, $\Pi_{2a}^{ex2} = \Pi_{firmT2}$, $\Pi_{3a}^{ex2} = \Pi_{firmT3}$, $\Pi_{3b}^{ex2} > \Pi_{firmT3}$, $\Pi_{3c}^{ex2} > \Pi_{firmT3}$. Thus, the firm will receive a weakly better payoff than in the ‘costless’ setting for any combination of C_H , C_L , and c .

A.4.6 Proof of Proposition 3

We first identify the conditions under which the firm will prefer the fully transparent algorithm over the opaque algorithm; this can be done by comparing the firm’s equilibrium payoff in each scenario (shown in Lemma 1 and Lemma 4). The conditions are summarized in Theorem 4.

Theorem 4 *The firm will prefer the fully transparent algorithm over the opaque algorithm if and only if any of the following conditions is satisfied:*

- Value pair $(\frac{C_L}{R}, \frac{C_H}{R})$ falls in region C1.
- Value pair $(\frac{C_L}{R}, \frac{C_H}{R})$ falls in region C4.
- Value pair $(\frac{C_L}{R}, \frac{C_H}{R})$ falls in region C2 and either $\beta > \beta_2$ or $C_L > C_L^{C2}$.
- Value pair $(\frac{C_L}{R}, \frac{C_H}{R})$ falls in region C3 and $C_L > C_L^{C3}$.
- Value pair $(\frac{C_L}{R}, \frac{C_H}{R})$ falls in region N3 and $C_L > C_L^{N3}$.

similarly.

where $C_L^{C2} = R - \frac{R^2(1-\theta)(1-\lambda)}{(\alpha-R)\theta\lambda}$, $C_L^{C3} = \frac{(2\lambda-1)R\theta}{\lambda}$, $C_L^{N3} = R\lambda\theta + \frac{R(1-\lambda)(1-\theta)(\beta-R)}{\theta(\alpha+\beta-R)}$.

Now we prove Theorem 4. Recall that, in Appendix A.4.4, we have shown that the firm's equilibrium payoffs are the same in the partial transparency scenario and the full transparency scenario if the value pair $(\frac{C_L}{R}, \frac{C_H}{R})$ falls in region $N1$, $N2$, and $C4$. Thus, in these regions, the firm's preference of transparency will not change when we switch from partial transparency to full transparency. If the value pair instead falls in regions $C1$, $C2$, $C3$, and $N3$, the firm may get a strictly higher payoff when we consider full transparency instead of partial transparency since these regions might overlap with region S . Consequently, in these regions, the firm's decision on algorithmic transparency might be altered if we study full transparency instead of partial transparency. First, notice that the left boundary of region S is $\frac{C_L}{R} = p_0 = \frac{\theta\alpha+\beta-R}{\theta(\alpha+\beta-R)}$. It is not difficult to check that $p_0 \in [0, 2 - \frac{1}{\lambda}]$ and region $C1$ will be included in region S since the left boundary of region $C1$ is $\frac{C_L}{R} = \lambda$ and $\lambda \geq 2 - \frac{1}{\lambda}$. Within region $C1$, the firm's payoff is constant in the opaque scenario ($\Pi_O^{C1} = \lambda\theta\alpha - (\lambda\theta + (1-\lambda)(1-\theta))R$) but increases in C_L in the full transparency scenario. The minimum payoff for the firm in the full transparency scenario in region $C1$ is reached at the left boundary and its value is $\min \Pi_{FT}^{C1} = \lambda\theta(\alpha + \beta - R)$. Since $\min \Pi_{FT}^{C1} - \Pi_O^{C1} = \lambda\theta\beta + (1-\lambda)(1-\theta)R > 0$, the firm will always prefer full transparency over opacity in region $C1$. For region $C2$, $\beta > \beta_2$ is a sufficient condition for the firm to prefer partial transparency over opacity and, thus, is also a sufficient condition for the firm to prefer full transparency over opacity since the firm's payoff is weakly higher in the full transparency scenario than in the partial transparency scenario. Another sufficient condition on $C_L > C_L^{C2}$ could be found by letting $\Pi_{FT}^{C2} = \Pi_O^{C2}$. Adding together, $(\beta > \beta_2 \text{ or } C_L > C_L^{C2})$ constitutes a necessary and sufficient condition for the firm to prefer transparency over opacity in region $C2$. The necessary and sufficient conditions for the firm to prefer transparency in region $C3$ and $N3$ can be found similarly and the details are omitted here.

A.4.7 Proof of Proposition 11

Our discussion on the ‘advantage’ case where $c_l > c_h$ is built on the previous case where $c_h = c_l = c$ shown in Lemma 11 and Lemma 12. We show that an increase in c_l can weakly increase the firm’s payoff in the transparent scenario. We will show this case by case in Figure A.1. In region 1a, increasing c_l will not impact agents’ behavior on the causal feature. It can only affect L type agents’ incentive to game the correlational feature, and this is always beneficial to the firm since the two types of agents will become more separated on the correlational feature compared with the ‘no advantage’ case. In regions 2a, since L type agents will not improve on the correlational feature even in the ‘no-advantage’ case, increasing c_l will not affect the equilibrium outcome and the payoff for the firm. In region 3a, before c_l reaches R , it would have no impact on the equilibrium outcome. If c_l increases to a value greater than R , the firm will be better off compared with the ‘no advantage’ case since L type agents in State C will not be willing to move to State B and thus H type and L type agents can be better separated apart. In regions 3b and 3c, increasing c_l will have no impact on the equilibrium outcome since L type agents in State C will not improve the causal feature even in the ‘no advantage’ case. To summarize, the firm gets weakly better payoffs in the ‘advantage’ setting than in the ‘no advantage’ setting.

A.4.8 Proof of Proposition 12

To make our discussion concrete, we assume the firm starts from a fixed wage R_0 and can strategically lower the wage to R_O and R_T in the opaque and transparent scenario, respectively. On the one hand, the firm has an incentive to decrease R to reduce the cost of hiring; on the other hand, the degree of separation among the agents on the causal feature may become worse if R is too small. The firm sets R by considering the trade-off between these two effects. We assume that R_O and R_T satisfy the conditions in Assumptions 1-1.3.4.




Figure A.14: Illustration on how the firm chooses optimal wage

Specifically, $R_O > R$ and $R_T > R$, where $R = \max(\theta\alpha, \beta + \frac{\theta(1-\lambda)\alpha}{\theta(1-\lambda)+(1-\theta)\lambda})^5$. This assumption helps us avoid several uninteresting scenarios such as the case where the firm sets R_O or R_T at a very low level such that it does not care about the true type of an applicant. Below, we analyze the firm's optimal choice of R_O and R_T and compare the equilibrium payoff in the opaque and transparent scenarios.

We start with the opaque scenario. Recall that, in Section 1.4.1, we divided the C_H - C_L space into different regions according to agents' strategies on the causal feature. Now, since R_O can be decreased by the firm, the (x, y) coordinates of any point $(\frac{C_L}{R_O}, \frac{C_H}{R_O})$ in Figure A.14a can be increased proportionally. For example, the yellow point in Figure A.14a will be moving along the yellow line, and the two ends of the line are defined as $(\frac{C_L}{R_0}, \frac{C_H}{R_0})$ and $(\frac{C_L}{R}, \frac{C_H}{R})$. As discussed before, we restrict our attention to the case where the firm prefers the transparent scenario under the fixed wage setting (R_0); in other words, we only consider the case where the left end of the line $(\frac{C_L}{R_0}, \frac{C_H}{R_0})$ lands in region $C1$, $C2$, and $C4$. We first study the case where the left end of the line lies in region $C1$ or $C4$. In this case, reducing R_O will not change agents' behavior on the causal feature (i.e., the whole yellow line lies in region 'case 1' defined in Section 1.4.1). Consequently, the firm will set $R_{O^*} = R$, and using the payoff function we have derived in Section 1.4.1, we have that the firm's equilibrium payoff

⁵This assumption can be relaxed and it will not affect the intuition provided by the analysis

is $\Pi_{C1}^O = \Pi_{C4}^O = \lambda\theta\alpha - (\lambda\theta + (1 - \lambda)(1 - \theta))R$. We next move to the case where $(\frac{C_L}{R_0}, \frac{C_H}{R_0})$ lands in region $C2$. In this case, if R_O is set lower than $\frac{C_L}{\lambda}$, the yellow point will cross the boundary of case 1 and case 4 so agents' strategies on the causal feature will be altered. The optimal point is either at the boundary of case 1 and case 4 or at the right end of the yellow line, depending on which of them gives the firm a higher payoff. That is,

$$R_{O^*} = \begin{cases} \frac{C_L}{R_0} & \text{if } R < \frac{C_L}{R_0} \text{ and } \Pi_{firm_{O4}}(\frac{C_L}{R_0}) \geq \Pi_{firm_{O1}}(R) \\ R & \text{if } R \leq \frac{C_L}{R_0} \text{ or } \Pi_{firm_{O4}}(\frac{C_L}{R_0}) < \Pi_{firm_{O1}}(R) \end{cases}$$

where $\Pi_{firm_{O1}}(R)$ and $\Pi_{firm_{O4}}(R)$ are the firm's payoff functions as defined in Section 1.4.1:

$$\Pi_{firm_{O1}}(R) = \lambda\theta\alpha - (\lambda\theta + (1 - \lambda)(1 - \theta))R, \Pi_{firm_{O4}}(R) = \theta(\alpha + \beta - R) \left(1 - \frac{R(1-\theta)(1-\lambda)}{(\alpha-R)\theta\lambda}\right).$$

The firm's optimal payoff in region $C2$ can then be shown as

$$\Pi_{C2}^O = \begin{cases} \lambda\theta\alpha - (\lambda\theta + (1 - \lambda)(1 - \theta))R_{O^*} & \text{if } R_{O^*} \leq \frac{C_L}{R_0} \\ \theta(\alpha + \beta - R_{O^*}) \left(1 - \frac{R_{O^*}(1-\theta)(1-\lambda)}{(\alpha-R_{O^*})\theta\lambda}\right) & \text{otherwise} \end{cases}$$

We next move to the transparent scenario. Figure A.14b illustrates how $R_T \in [R, R_0]$ affects the agents' behavior on the causal feature in the transparent scenario. Similar to our discussions in the opaque scenario, we only consider the case where $(\frac{C_L}{R_0}, \frac{C_H}{R_0})$ falls in $C1$, $C2$, and $C4$. We first consider the case where $R \geq C_H$. It is straightforward to see that, in this case, $R_T^* = R$ regardless of which region $(\frac{C_L}{R_0}, \frac{C_H}{R_0})$ is in. Next, we consider the case where $R < C_H$: it can be shown that in this case $R_T^* = C_H$ regardless of which region $(\frac{C_L}{R_0}, \frac{C_H}{R_0})$ is in. The intuition is as follows: If the yellow line crosses the border between case 1 and case 2, the optimal point will be exactly on the border because a further decrease in R_T will change the agents' strategies on the causal feature from perfect separation to polling, and the latter

will result in a zero payoff. The firm's optimal payoff can then be shown as

$$\Pi_{C1}^T = \Pi_{C2}^T = \Pi_{C4}^T = \begin{cases} \theta(\alpha + \beta - R_{T*}) & \text{if } R_{T*} \leq C_L \\ \theta(\alpha + \beta) + (1 - \theta)\beta - R_{T*} & \text{otherwise} \end{cases}$$

We next compare the optimal wage and the maximum profit in the opaque and the transparent scenarios. It is straightforward to see $R_{T*} \leq R_{O*}$ if $R \geq C_H$. Comparing the optimal payoff functions, we have $\Pi_{C_i}^T \geq \Pi_{C_i}^O, \forall i \in 1, 2, 4$. There is an intuitive explanation behind this result: If we force the firm to choose R_{O*} (a sub-optimal choice) in the transparent scenario, the firm can still get a better payoff than choosing R_{O*} (the optimal choice) in the opaque scenario, which means that the firm will get an even better payoff when choosing R_{T*} in the transparent scenario. We next consider the case $R < C_H$. In this case, when $R_{T*} = C_H$ and $R_{O*} = R$, $R_{T*} \leq R_{O*}$ does not hold, and therefore, the above logic does not work. We need to directly compare the firm's payoff in the transparent and the opaque scenarios. Note that in this case $(\frac{C_L}{R_{T*}}, \frac{C_H}{R_{T*}})$ falls in region $C4$, which belongs to transparent scenario 2; $(\frac{C_L}{R_{O*}}, \frac{C_H}{R_{O*}})$ falls in region $N1$, which belongs to opaque scenario 1. Using the firm's payoff function shown in Lemma 1 and Lemma 2, we get $\Pi_{C_i}^T = \theta(\alpha + \beta - R_{T*})$ and $\Pi_{C_i}^O = \lambda\theta\alpha - (\lambda\theta + (1 - \lambda)(1 - \theta))R$. Since $R_{T*} < \beta + \theta\alpha$ (according to Assumption 1.3.4) and $R_{O*} = R \geq \theta\alpha$ (according to Assumption 2), the following inequalities must hold: $\Pi_{C_i}^T = \theta(\alpha + \beta - R_{T*}) > \theta(\alpha + \beta - (\beta + \theta\alpha)) = \alpha\theta(1 - \theta)$, and $\Pi_{C_i}^O = \lambda\theta\alpha - (\lambda\theta + (1 - \lambda)(1 - \theta))R < \lambda\theta\alpha - (\lambda\theta + (1 - \lambda)(1 - \theta))\theta\alpha = \alpha\theta(1 - \theta)(2\lambda - 1)$. Considering the fact that $2\lambda - 1 \leq 1$, we have $\Pi_{C_i}^T > \Pi_{C_i}^O, \forall i \in 1, 2, 4$.

Appendix B

Appendix for Chapter 2

B.1 Notation Summary

B.2 Model Extension

B.2.1 Multiple Applications

In the main model, we restrict each borrower to apply to at most one lender for simplification. In this appendix, we consider the possibility that a lender with a unit demand may apply to multiple lenders sequentially. This scenario arises when, for example, a borrower believes she will be approved by both lenders and applies to lender 1, but is rejected by lender 1 due to incorrect beliefs. She then proceeds to apply to lender 2. Considering the multi-stage application process is more complicated, but it is worth examining to assess the robustness of our result. Specifically, in this extension, borrowers have the opportunity to re-apply to another lender if they are rejected by one lender.

Notice that this adjustment does not change the analysis for the R-R scenario, since the approval decisions that borrowers receive are always consistent with their beliefs. In the N-N scenario, this adjustment will change borrowers' behavior and the lenders' equilibrium pricing strategy accordingly. In particular, given our assumption of binary beliefs, borrowers who

Notation	Meaning
P_b	Accuracy of lenders' algorithms
P_c	Accuracy of borrowers' beliefs in absence of pre-approval tools
θ	Percentage of H type borrowers in the population
m	Borrowers' cost of applying
M_h	Utility of being approved for a H type borrower
M_l	Utility of being approved for a L type borrower
b	Interest rate set by the lenders
\bar{b}	Maximum interest rate can be set by the lenders
$Y_{k,b}^i$	Lender i 's prediction of borrower k 's type
$Y_{k,c}^i$	borrower k 's belief of lender i 's approval decision
$N_{hh}^{rr}(N_{hh}^{nn})$	The number of borrowers in a lender's common segment in the R-R (N-N) case
$N_{hl}^{rr}(N_{hl}^{nn})$	The number of borrowers in a lender's captive segment in the R-R (N-N) case
$V_{hh}^{rr}(V_{hh}^{nn})$	The percentage of H type borrowers in a lender's common segment in the R-R (N-N) case
$V_{hl}^{rr}(V_{hl}^{nn})$	The percentage of H type borrowers in a lender's captive segment in the R-R (N-N) case
$N_{hh}^1(N_{hh}^2)$	The number of borrowers in Lender 1(2)'s common segment in the R-N case
$N_{hl}^1(N_{hl}^2)$	The number of borrowers in Lender 1(2)'s captive segment in the R-N case
$V_{hh}^1(V_{hh}^1)$	The percentage of H type borrowers in Lender 1(2)'s common segment in the R-N case
$V_{hl}^1(V_{hl}^2)$	The percentage of H type borrowers in Lender 1(2)'s captive segment in the R-N case

Table B.1: Notation summary

believe they will not be approved by either lender will not apply in the first place. Similarly, borrowers who believe they will be approved by only one of the two lenders but end up being rejected will not apply to the other lender either. The possibility of re-application only arises for borrowers who believe they will be approved by both lenders but end up being rejected by one of them. From lender 1's perspective, this refers to the *HHLH* (Row 3, Column 1) and *LHLH* (Row 3, Column 3) segments depicted in Table 2.2. Moreover, the *LHLH* segment does not affect lender 1's pricing strategy since they will not be approved by lender 1 eventually. The only segment that affects lenders' equilibrium strategies is the *HHLH* segment: now, even when a lender's price is higher than the competitor's, it will still attract the *HHLH* segment eventually. Consequently, competition is less intense compared to the main model. The lenders' equilibrium profits are shown in Lemma 13.

Lemma 13 *With borrowers' multi-stage applications, the lenders' equilibrium payoff in the N-N scenario is:*

$$\begin{aligned}\Pi_{nn-MA} &= N_{hl}^{nn} (\bar{b}V_{hl}^{nn} + V_{hl}^{nn} - 1) + N_x^{nn} (\bar{b}V_x^{nn} + V_x^{nn} - 1) \\ &= \Pi_{nn} + N_x^{nn} (\bar{b}V_x^{nn} + V_x^{nn} - 1)\end{aligned}$$

Where N_{hl}^{nn} and V_{hl}^{nn} are defined in Section 2.4.4, Π_{nn} is the equilibrium payoff in the main case, and $N_x^{nn} = P_b(1 - P_b)P_c(1 - P_c)$ and $V_x^{nn} = \theta$ are the number of borrowers and the fraction of *H* type borrowers in the *HHLH* segment, respectively.

Comparing Lemma 13 with Lemma 7, we observe that both lenders' equilibrium profits are higher than those in the main model: the existence of this *HHLH* segment softens the competition.

In the R-N scenario, following similar logic, the *HLH* segment in Table 2.3 will be affected: they will apply to lender 1 after the initial rejection from lender 2 (not the other way around since if they first apply to lender 1, they will get approved). The lenders' equilibrium profits in this case are summarized in Lemma 14.

Lemma 14 *With borrowers' multi-stage applications, the lenders' equilibrium payoff in the R-N scenario is:*

$$\begin{aligned}\Pi_{rn-MA}^1 &= \Pi_{rn}^1 + N_x^{rn} (\bar{b}V_x^{rn} + V_x^{rn} - 1) \\ \Pi_{rn-MA}^2 &= \Pi_{rn}^2 + N_x^{rn} (\bar{b}V_x^{rn} + V_x^{rn} - 1)(N_{hh}^2 V_{hh}^2 + N_{lh}^2 V_{lh}^2)/(N_{hh}^1 V_{hh}^1 + N_{hl}^1 V_{hl}^1)\end{aligned}$$


Where N_{hh}^1 , V_{hh}^1 , N_{hl}^1 , V_{hl}^1 , N_{lh}^2 , and V_{lh}^2 are defined in Section 2.4.4, Π_{rn}^1 and Π_{rn}^2 are the equilibrium payoff for the R-N scenario in the main model. And $N_x^{rn} = P_b(1 - P_b)(1 - P_c)$ and $V_x^{rn} = \theta$ are the number of borrowers and the fraction of H type borrowers in the HLH segment, respectively.

It's straightforward to see that the equilibrium payoffs for the two lenders here are strictly higher than those in the main model (i.e., $\Pi_{rn-MA}^1 > \Pi_{rn}^1$ and $\Pi_{rn-MA}^2 > \Pi_{rn}^2$), showing that the inclusion of segment *HLH* softens the competition.


By comparing the equilibrium outcomes in each case, the lenders' decisions on pre-approval revelation are summarized in Proposition 14.

Proposition 14 *Anticipating borrowers' multi-stage applications, the condition on P_b for the asymmetric pre-approval revealing equilibrium is less strict compared with the main case in Proposition 4.*

The proof is straightforward. Notice that N-N will again never be sustained in equilibrium, given 1) $\Pi_{rn}^1 > \Pi_{nn}$ (from Proposition 4) and 2) $\Pi_{rn-MA}^1 - \Pi_{rn}^1 > \Pi_{nn-MA} - \Pi_{nn}$ (which can be derived by comparing Lemma 13 and 14). Now, whether the equilibrium is R-R or R-N depends on the non-revealing lender's equilibrium payoff in each case. The equilibrium payoff in the R-R scenario is the same as in the main case, as discussed before. However, the equilibrium payoff in the R-N scenario is higher than that in the main case: $\Pi_{rn-MA}^2 > \Pi_{rn}^2$ (the extra term in the expression of Π_{rn-MA}^2 is strictly positive). Thus, for any parameter configurations that will sustain R-N as an equilibrium in the main case, they will also sustain R-N as an equilibrium in this extended model. The main results of our paper are strengthened



(a) SPNE in the main model



(b) SPNE in the extended model

Figure B.1: Pre-approval revealing equilibrium (depending on P_b and P_c) comparison between the main model and the extended model. Other parameters are set at: $\theta = 0.8$, $\bar{b} = 0.5$. In region R-R, both lenders reveal. In region R-N, only one lender reveals

when considering multi-stage applications.

A visual comparison of the SPNE in the main case and in this extension is shown in Figure B.1.

B.2.2 Vertically Differentiated Products

In this appendix, we relax the assumption of symmetric lenders and consider the case where lender 1's quality is higher than lender 2's. This can be thought of as lender 1's financial product having better non-price features or the overall service being of higher quality. Specifically, we use M_h^1 (M_h^2) to denote the overall benefit that an H type borrower can get from lender 1(2)'s product, and similarly use M_l^1 (M_l^2) to denote the overall benefit that an L type borrower can get from lender 1(2)'s product. For ease of discussion, we assume $M_h^1 - M_h^2 = M_l^1 - M_l^2 = \delta b$, where $\delta b > 0$.

With asymmetric lenders, solving for the equilibrium pricing strategy and equilibrium price is more complicated, as we cannot simply focus on the symmetric mixed strategy equilibrium as in Appendix B.3.1 and B.3.2. The detailed derivations are omitted; instead, we directly report the equilibrium outcomes. Since the two lenders are not symmetric, we have to study 4 subgames: the R-R case where both lenders reveal the pre-approval outcomes,

the N-N case where neither lender reveals, the R-N case where only lender 1 reveals, and the N-R case where only lender 2 reveals.

The equilibrium profits for each of the 4 subgames are shown in Lemma 15, 16, 17, 18, respectively.

Lemma 15 *With vertically differentiated products, if both lenders reveal the pre-approval outcomes (R-R case), lender 1 and 2's equilibrium payoffs are:*

$$\Pi_{rr-VD}^1 = \max(1[F_{rr}^{1A}(\bar{b}) \leq 1] \cdot \Pi_{rr}^{1A}, 1[F_{rr}^{2B}(\bar{b}) < 1] \cdot \Pi_{rr}^{1B})) \quad (\text{B.1})$$

$$\Pi_{rr-VD}^2 = \max(1[F_{rr}^{1A}(\bar{b}) \leq 1] \cdot \Pi_{rr}^{2A}, 1[F_{rr}^{2B}(\bar{b}) < 1] \cdot \Pi_{rr}^{2B})) \quad (\text{B.2})$$

Where

$$\Pi_{rr}^{1A} = P_b (-\bar{b}P_b\theta + \bar{b}\theta - P_b\theta + P_b + \theta - 1 + \delta b\theta(1 - P_b))$$

$$\Pi_{rr}^{2A} = P_b (-\bar{b}P_b\theta + \bar{b}\theta - P_b\theta + P_b + \theta - 1 - \delta bP_b\theta)$$

$$\Pi_{rr}^{1B} = P_b (-\bar{b}P_b\theta + \bar{b}\theta - P_b\theta + P_b + \theta - 1 + \delta bP_b\theta)$$

$$\Pi_{rr}^{2B} = P_b (-\bar{b}P_b\theta + \bar{b}\theta - P_b\theta + P_b + \theta - 1 + \delta b\theta(P_b - 1))$$

$$F_{rr}^{1A}(b) = \frac{P_b\theta(b - \delta b) - P_b\theta - P_b(-\bar{b}P_b\theta + \bar{b}\theta - \delta bP_b\theta - P_b\theta + P_b + \theta - 1) + P_b + \theta - 1}{P_b^2\theta(b - \delta b) + P_b^2\theta - P_b^2 - 2P_b\theta + 2P_b + \theta - 1}$$

$$F_{rr}^{2B}(b) = \frac{P_b\theta(b + \delta b) - P_b\theta - P_b(-\bar{b}P_b\theta + \bar{b}\theta + \delta bP_b\theta - P_b\theta + P_b + \theta - 1) + P_b + \theta - 1}{P_b^2\theta(b + \delta b) + P_b^2\theta - P_b^2 - 2P_b\theta + 2P_b + \theta - 1}$$

Lemma 16 *With vertically differentiated products, if neither lenders reveal the pre-approval outcomes (N-N case), lender 1 and 2's equilibrium payoffs are:*

$$\Pi_{nn-VD}^1 = \max(1[F_{nn}^{1A}(\bar{b}) \leq 1] \cdot \Pi_{nn}^{1A}, 1[F_{nn}^{2B}(\bar{b}) < 1] \cdot \Pi_{nn}^{1B})) \quad (\text{B.3})$$

$$\Pi_{nn-VD}^2 = \max(1[F_{nn}^{1A}(\bar{b}) \leq 1] \cdot \Pi_{nn}^{2A}, 1[F_{nn}^{2B}(\bar{b}) < 1] \cdot \Pi_{nn}^{2B})) \quad (\text{B.4})$$

Where

$$\begin{aligned}
\Pi_{nn}^{1A} &= N_{hl} (\bar{b}V_{hl} + \delta bV_{hl} + V_{hl} - 1) \\
\Pi_{nn}^{2A} &= \bar{b}N_{hl}V_{hl} - \delta bN_{hh}V_{hh} + N_{hl}V_{hl} - N_{hl} \\
\Pi_{nn}^{1B} &= \bar{b}N_{hl}V_{hl} + \delta bN_{hh}V_{hh} + N_{hl}V_{hl} - N_{hl} \\
\Pi_{nn}^{2B} &= N_{hl} (\bar{b}V_{hl} - \delta bV_{hl} + V_{hl} - 1) \\
F_{nn}^{1A}(b) &= \frac{-\bar{b}N_{hl}V_{hl} + \delta bN_{hh}V_{hh} + N_{hh}V_{hh} (b - \delta b) + N_{hh}V_{hh} - N_{hh} + N_{hl}V_{hl} (b - \delta b)}{N_{hh} (V_{hh} (b - \delta b) + V_{hh} - 1)} \\
F_{nn}^{2B}(b) &= \frac{-\bar{b}N_{hl}V_{hl} - \delta bN_{hh}V_{hh} + N_{hh}V_{hh} (b + \delta b) + N_{hh}V_{hh} - N_{hh} + N_{hl}V_{hl} (b + \delta b)}{N_{hh} (V_{hh} (b + \delta b) + V_{hh} - 1)}
\end{aligned}$$

Lemma 17 *With vertically differentiated products, if only lender 1 reveals the pre-approval outcomes (R-N case), lender 1 and 2's equilibrium payoffs are:*

$$\Pi_{rn-VD}^1 = \max(1[F_{rn}^{1A}(\bar{b}) \leq 1] \cdot \Pi_{rn}^{1A}, 1[F_{rn}^{2B}(\bar{b}) < 1] \cdot \Pi_{rn}^{1B})) \quad (\text{B.5})$$

$$\Pi_{rn-VD}^2 = \max(1[F_{rn}^{1A}(\bar{b}) \leq 1] \cdot \Pi_{rn}^{2A}, 1[F_{rn}^{2B}(\bar{b}) < 1] \cdot \Pi_{rn}^{2B})) \quad (\text{B.6})$$

Where

$$\Pi_{rn}^{1A} = N_{hl}^1 (\bar{b}V_{hl}^1 + \delta bV_{hl}^1 + V_{hl}^1 - 1)$$

$$\begin{aligned} \Pi_{rn}^{2A} = & (\bar{b}N_{hh}^2 N_{hl}^1 V_{hh}^2 V_{hl}^1 + \bar{b}N_{hl}^1 N_{lh}^2 V_{hl}^1 V_{lh}^2 - \delta bN_{hh}^1 N_{hh}^2 V_{hh}^1 V_{hh}^2 - \delta bN_{hh}^1 N_{lh}^2 V_{hh}^1 V_{lh}^2 \\ & - N_{hh}^1 N_{hh}^2 V_{hh}^1 + N_{hh}^1 N_{hh}^2 V_{hh}^2 - N_{hh}^1 N_{lh}^2 V_{hh}^1 + N_{hh}^1 N_{lh}^2 V_{lh}^2 + N_{hh}^2 N_{hl}^1 V_{hh}^2 V_{hl}^1 \\ & - N_{hh}^2 N_{hl}^1 V_{hl}^1 + N_{hl}^1 N_{lh}^2 V_{hl}^1 V_{lh}^2 - N_{hl}^1 N_{lh}^2 V_{hl}^1) / (N_{hh}^1 V_{hh}^1 + N_{hl}^1 V_{hl}^1) \end{aligned}$$

$$\begin{aligned} \Pi_{rn}^{1B} = & (\bar{b}N_{hh}^1 N_{lh}^2 V_{hh}^1 V_{lh}^2 + \bar{b}N_{hl}^1 N_{lh}^2 V_{hl}^1 V_{lh}^2 + \delta bN_{hh}^1 N_{hh}^2 V_{hh}^1 V_{hh}^2 + \delta bN_{hh}^2 N_{hl}^1 V_{hh}^2 V_{hl}^1 \\ & + N_{hh}^1 N_{hh}^2 V_{hh}^1 - N_{hh}^1 N_{hh}^2 V_{hh}^2 + N_{hh}^1 N_{lh}^2 V_{hh}^1 V_{lh}^2 - N_{hh}^1 N_{lh}^2 V_{lh}^2 - N_{hh}^2 N_{hl}^1 V_{hh}^2 \\ & + N_{hh}^2 N_{hl}^1 V_{hl}^1 + N_{hl}^1 N_{lh}^2 V_{hl}^1 V_{lh}^2 - N_{hl}^1 N_{lh}^2 V_{lh}^2) / (N_{hh}^2 V_{hh}^2 + N_{lh}^2 V_{lh}^2) \end{aligned}$$

$$\Pi_{rn}^{2B} = N_{lh}^2 (\bar{b}V_{lh}^2 - \delta bV_{lh}^2 + V_{lh}^2 - 1)$$

$$F_{rn}^{1A}(b) = ((N_{hh}^2 V_{hh}^2 + N_{lh}^2 V_{lh}^2) (b - \delta b + 1) - N_{hh}^2 - N_{lh}^2 - \Pi_{rn}^{2A}) / (N_{hh}^2 (V_{hh}^2 (b - \delta b + 1) - 1))$$

$$F_{rn}^{2B}(b) = ((N_{hh}^1 V_{hh}^1 + N_{hl}^1 V_{hl}^1) (b + \delta b + 1) - N_{hh}^1 - N_{hl}^1 - \Pi_{rn}^{1B}) / (N_{hh}^1 (V_{hh}^1 (b + \delta b + 1) - 1))$$

Lemma 18 *With vertically differentiated products, if only lender 2 reveals the pre-approval outcomes (N-R case), lender 1 and 2's equilibrium payoffs are:*

$$\Pi_{nr-VD}^1 = \max(1[F_{nr}^{1A}(\bar{b}) \leq 1] \cdot \Pi_{nr}^{1A}, 1[F_{nr}^{2B}(\bar{b}) < 1] \cdot \Pi_{nr}^{1B})) \quad (\text{B.7})$$

$$\Pi_{nr-VD}^2 = \max(1[F_{nr}^{1A}(\bar{b}) \leq 1] \cdot \Pi_{nr}^{2A}, 1[F_{nr}^{2B}(\bar{b}) < 1] \cdot \Pi_{nr}^{2B})) \quad (\text{B.8})$$

Where

$$\begin{aligned}
\Pi_{nr}^{1A} &= (\bar{b}N_{hh}^2N_{hl}^1V_{hh}^2V_{hl}^1 + \bar{b}N_{hl}^1N_{lh}^2V_{hl}^1V_{lh}^2 + \delta bN_{hh}^1N_{hh}^2V_{hh}^1V_{hh}^2 + \delta bN_{hh}^1N_{lh}^2V_{hh}^1V_{lh}^2 \\
&\quad - N_{hh}^1N_{hh}^2V_{hh}^1 + N_{hh}^1N_{hh}^2V_{hh}^2 - N_{hh}^1N_{lh}^2V_{hh}^1 + N_{hh}^1N_{lh}^2V_{lh}^2 + N_{hh}^2N_{hl}^1V_{hh}^2V_{hl}^1 \\
&\quad - N_{hh}^2N_{hl}^1V_{hl}^1 + N_{hl}^1N_{lh}^2V_{hl}^1V_{lh}^2 - N_{hl}^1N_{lh}^2V_{hl}^1)/(N_{hh}^1V_{hh}^1 + N_{hl}^1V_{hl}^1) \\
\Pi_{nr}^{2A} &= N_{hl}^1(\bar{b}V_{hl}^1 - \delta bV_{hl}^1 + V_{hl}^1 - 1) \\
\Pi_{nr}^{1B} &= N_{lh}^2(\bar{b}V_{lh}^2 + \delta bV_{lh}^2 + V_{lh}^2 - 1) \\
\Pi_{nr}^{2B} &= (\bar{b}N_{hh}^1N_{lh}^2V_{hh}^1V_{lh}^2 + \bar{b}N_{hl}^1N_{lh}^2V_{hl}^1V_{lh}^2 - \delta bN_{hh}^1N_{hh}^2V_{hh}^1V_{hh}^2 - \delta bN_{hh}^2N_{hl}^1V_{hh}^2V_{hl}^1 \\
&\quad + N_{hh}^1N_{hh}^2V_{hh}^1 - N_{hh}^1N_{hh}^2V_{hh}^2 + N_{hh}^1N_{lh}^2V_{hh}^1V_{lh}^2 - N_{hh}^1N_{lh}^2V_{lh}^2 - N_{hh}^2N_{hl}^1V_{hh}^2 \\
&\quad + N_{hh}^2N_{hl}^1V_{hl}^1 + N_{hl}^1N_{lh}^2V_{hl}^1V_{lh}^2 - N_{hl}^1N_{lh}^2V_{lh}^2)/(N_{hh}^2V_{hh}^2 + N_{lh}^2V_{lh}^2) \\
F_{nn}^{1A} &= ((N_{hh}^2V_{hh}^2 + N_{lh}^2V_{lh}^2)(b + \delta b + 1) - N_{hh}^2 - N_{lh}^2 - \Pi_{nr}^{1A})/(N_{hh}^2(V_{hh}^2(b + \delta b + 1) - 1)) \\
F_{nr}^{2B} &= ((N_{hh}^1V_{hh}^1 + N_{hl}^1V_{hl}^1)(b - \delta b + 1) - N_{hh}^1 - N_{hl}^1 - \Pi_{nr}^{2B})/(N_{hh}^1(V_{hh}^1(b - \delta b + 1) - 1))
\end{aligned}$$

The equilibrium outcomes in all the subgames constitute the payoff metric that two lenders face in the pre-approval revealing stage:

		Lender 2	
		Not Reveal	Reveal
Lender 1	Not Reveal	$(\Pi_{nn-V_D}^1, \Pi_{nn-V_D}^2)$	$(\Pi_{nr-V_D}^1, \Pi_{nr-V_D}^2)$
	Reveal	$(\Pi_{rn-V_D}^1, \Pi_{rn-V_D}^2)$	$(\Pi_{rr-V_D}^1, \Pi_{rr-V_D}^2)$

The comparison of each term in the payoff matrix is not straightforward, here we numerically show the equilibrium outcome under each parameter pair. Consistent with the parameters used in Figure 2.4, in Figure B.2, we use $\theta = 0.8$ and $\bar{b} = 0.5$, and vary P_b , P_c , and δb .

Compare Figure B.2 with Figure 2.4, the result on asymmetric equilibrium stays unchanged qualitatively: at least one of the two lenders will reveal the pre-approval outcome, and both will reveal only if P_b is relatively small.




Figure B.2: Pre-approval revealing equilibrium depending on P_b , P_c , and δb when $\theta = 0.8$, $\bar{b} = 0.5$. In region R-R, both lenders reveal. In region R-N, lender 1 reveals. In region N-R, lender 2 reveals. In region R-N/N-R, there exists two asymmetric equilibrium

B.2.3 Asymmetric Lender Accuracy

In this section, we relax the symmetric algorithm accuracy assumption and allow the two lenders' algorithms to have different predictive accuracy. Denote lender 1's algorithm's accuracy as P_{b1} , lender 2's as P_{b2} . We assume $P_{b1} \geq P_{b2}$ without loss of generality. Specifically, we use a parameter η to capture the difference in the two lenders' algorithm accuracy: $P_{b2} = 0.5 + \eta(P_{b1} - 0.5)$ where $\eta \in (0, 1]$. As mentioned in Appendix ??, with asymmetric lenders, solving for the equilibrium pricing strategy is more complicated, since we cannot simply focus on the symmetric mixed strategies. We study 4 subgames in this extension: R-R case where both lenders reveal the pre-approval outcomes. N-N case where neither lender reveals. R-N case where only lender 1 reveals. N-R case where only lender 2 reveals.

We omit the detailed derivations of equilibrium strategies in each subgame and directly show the results visually. Consistent with the parameters used in Figure 2.4, in Figure B.3, we use $\theta = 0.8$ and $\bar{b} = 0.5$, and vary P_{b1} , P_c , and η .

Compare Figure B.3 with Figure 2.4, the result on asymmetric equilibrium stays unchanged qualitatively: at least one of the two lenders will reveal the pre-approval outcome, and both will reveal only if P_b is relatively small. However, as the difference between the two lenders'




Figure B.3: Pre-approval revealing equilibrium depending on P_{b1} , P_c , and η when $\theta = 0.8$, $\bar{b} = 0.5$. In region R-R, both lenders reveal. In region R-N, lender 1 reveals. In region N-R, lender 2 reveals. In region R-N/N-R, there exists two asymmetric equilibrium

algorithm accuracy increases (i.e., as η decreases), the region where only the high-accuracy lender (i.e., lender 1) reveals emerges and enlarges. To explain this, let's recap the logic for the existence of asymmetric equilibrium: the non-revealing lender does not want to reveal because it will lead to head-head competition with the other lender. However, if the non-revealing lenders' accuracy is high, its concerns about the intensified competition in the R-R case will be reduced thanks to its advantage in accuracy. The high accuracy non-revealing lender will have the incentive to reveal, eliminating the N-R case being an equilibrium. Moreover, this happens only when P_b is not very large, otherwise, the competition intensity in the R-R case will be so high that even the high-accuracy lender wants to avoid it. In such a case, we will get both N-R and R-N as an equilibrium (the dark grey regions in each graph).

B.2.4 Correlated Predictions

In this subsection, we relax the independence assumption and allow the two lenders' predictions on borrowers' types to be positively correlated. Specifically, we use a parameter ρ to capture the amount of correlation between the two lenders: the probability that both lenders predict an H-type borrower as H-type is $P_b^2 + \rho(P_b - P_b^2)$, the probability that both lenders predict an H-type borrower as L-type is $(1 - P_b)^2 + \rho(1 - P_b - (1 - P_b)^2)$, the probability that lender 1(2)




Figure B.4: Pre-approval revealing equilibrium depending on P_b , P_c , and ρ when $\theta = 0.8$, $\bar{b} = 0.5$. In region R-R, both lenders reveal. In region R-N/N-R, there exists two asymmetric equilibrium

predicts an H-type borrower as H-type but lender 2(1) predicts as L-type is $P_b(1 - P_b)(1 - \rho)$. When $\rho = 0$, it degenerates to the independence case we studied in the main model.

Again, we omit the detailed derivations of equilibrium strategies in each subgame and directly show the results visually. Consistent with the parameters used in Figure 2.4, in Figure B.4, we use $\theta = 0.8$ and $\bar{b} = 0.5$, and vary P_{b1} , P_c , and ρ .

Relaxing the independence assumption strengthens our main results. As the correlation between the two lenders' predictions increases, asymmetric revealing equilibria appear in larger regions. The logic is clear: as the correlation of predictions increases, the two lenders tend to favor the same group of borrowers, intensifying the competition. If these borrowers know they will be approved by both lenders upfront, it will make things even worse for the two lenders. That's why with correlated predictions, the two lenders have a higher incentive to avoid the R-R case, making the asymmetric case more favorable.

B.3 Mathematical Appendix

B.3.1 Proof of Lemma 6

Let $F^{rr}(b)$ denote the CDF of both lender's equilibrium mixed strategy for the pricing decision. Facing Lender 2's mixed strategy characterized by $F^{rr}(b)$, Lender 1's expected profit is:

$$\begin{aligned} E[\Pi_1^{rr}](b) = & (1 - F^{rr}(b))(N_{hh}^{rr}(V_{hh}^{rr}b - 1 + V_{hh}^{rr}) + N_{hl}^{rr}(V_{hl}^{rr}b - 1 + V_{hl}^{rr})) \\ & + F^{rr}(N_{hl}^{rr}(V_{hl}^{rr}b - 1 + V_{hl}^{rr})) \end{aligned} \quad (\text{B.9})$$

The first line in Equation (B.9) corresponds to the case where Lender 1 sets a lower interest rate b than Lender 2 and thus Lender 1 gets segments A_{hh}^{rr} and A_{hl}^{rr} . The second line corresponds to the case where Lender 1 sets a higher b than Lender 2 and thus Lender 1 gets only segment A_{hl}^{rr} .

To make sure that Lender 1 is using the same mixed strategy characterized by $F^{rr}(b)$, Lender 1 has to be indifferent in setting any b on the support of $F^{rr}(b)$, mathematically,

$$E[\Pi_1^{rr}](b) = k \quad (\text{B.10})$$

where k is a constant. The maximum b that can be set by the lenders is \bar{b} , i.e.,

$$F^{rr}(\bar{b}) = 1 \quad (\text{B.11})$$

Using Equations (B.9), (B.10), and (B.11), we can solve for $F^{rr}(b)$. The mixed strategy equilibrium in this sub-game is as shown in Lemma 6.

B.3.2 Proof of Lemma 7

From Table 2.2, we can calculate:

$$\begin{aligned} N_{hh}^{nn} &= P_c^2 P_b^2 \theta + P_c(1 - P_c)P_b(1 - P_b) + P_c^2(1 - P_b)^2(1 - \theta) \\ V_{hh}^{nn} &= \frac{P_c^2 P_b^2 \theta + P_c(1 - P_c)P_b(1 - P_b)\theta}{N_{hh}^{nn}} > \theta \\ N_{hl}^{nn} &= P_c(1 - P_c)P_b^2 \theta + P_c(1 - P_c)(1 - P_b)^2(1 - \theta) + P_c^2 P_b(1 - P_b) \\ V_{hl}^{nn} &= \frac{P_b^2 P_c(1 - P_c)\theta + P_c^2 P_b(1 - P_b)\theta}{N_{hl}^{nn}} > \theta \end{aligned} \quad (\text{B.12})$$

Facing Lender 2's pricing strategy characterized by the CDF $F^{nn}(b)$, Lender 1's expected profit is:

$$\begin{aligned} E[\Pi_1](b) = & (1 - F^{nn}(b))(N_{hh}^{nn}(V_{hh}^{nn}b - (1 - V_{hh}^{nn}))b + N_{hl}^{nn}(V_{hl}^{nn}b - (1 - V_{hl}^{nn}))) \\ & + F^{nn}(b)N_{hl}^{nn}(V_{hl}^{nn}b - (1 - V_{hl}^{nn})) \end{aligned} \quad (\text{B.13})$$

The first line corresponds to the case where Lender 1 sets a lower b than Lender 2, and

thus Lender 1 gets segments A_{hh}^{nn} and A_{hl}^{nn} . The second line corresponds to the case where Lender 1 sets a higher b than Lender 2, and thus Lender 1 only gets the A_{hl}^{nn} segment.

In a symmetric equilibrium, Lender 1 is using the same mixed strategy characterized by $F^{nn}(b)$, therefore, Lender 1 has to be indifferent in setting any b on the support of $F^{nn}(b)$. Mathematically,

$$E[\Pi_1](b) = k \quad (\text{B.14})$$

where k is a constant. The maximum b that can be set by the lenders is \bar{b} , i.e.,

$$F^{nn}(\bar{b}) = 1 \quad (\text{B.15})$$

Combining Equations (B.13), (B.14), and (B.15), we can solve for $F^{nn}(b)$. The mixed strategy equilibrium in this sub-game is summarized in Lemma 7.

B.3.3 Proof of Lemma 8

From Table 2.3, we can calculate:

$$\begin{aligned} N_{hh}^1 &= P_b^2 P_c \theta + P_b \theta (-P_b + 1) (-P_c + 1) + P_b (-P_b + 1) (-P_c + 1) (-\theta + 1) \\ &\quad + P_c (-P_b + 1)^2 (-\theta + 1) \\ N_{hh}^2 &= P_b^2 P_c \theta + P_c (-P_b + 1)^2 (-\theta + 1) \\ N_{hl}^1 &= P_b^2 \theta (-P_c + 1) + P_b P_c \theta (-P_b + 1) + P_b P_c (-P_b + 1) (-\theta + 1) \\ &\quad + (-P_b + 1)^2 (-P_c + 1) (-\theta + 1) \\ N_{lh}^2 &= P_b P_c \theta (-P_b + 1) + P_b P_c (-P_b + 1) (-\theta + 1) \\ V_{hh}^1 &= \frac{P_b^2 P_c \theta + P_b \theta (-P_b + 1) (-P_c + 1)}{P_b^2 P_c \theta + P_b \theta (-P_b + 1) (-P_c + 1) + P_b (-P_b + 1) (-P_c + 1) (-\theta + 1) + P_c (-P_b + 1)^2 (-\theta + 1)} \\ V_{hh}^2 &= \frac{P_b^2 P_c \theta}{P_b^2 P_c \theta + P_c (-P_b + 1)^2 (-\theta + 1)} \\ V_{hl}^1 &= \frac{P_b^2 \theta (-P_c + 1) + P_b P_c \theta (-P_b + 1)}{P_b^2 \theta (-P_c + 1) + P_b P_c \theta (-P_b + 1) + P_b P_c (-P_b + 1) (-\theta + 1) + (-P_b + 1)^2 (-P_c + 1) (-\theta + 1)} \\ V_{lh}^2 &= \frac{P_b P_c \theta (-P_b + 1)}{P_b P_c \theta (-P_b + 1) + P_b P_c (-P_b + 1) (-\theta + 1)} \end{aligned} \quad (\text{B.16})$$

If $b_1 > b_2$, Lender 1 will only get the A_{hl}^1 segment, and Lender 2 will get the A_{hh}^2 and A_{lh}^2 segments. If $b_1 < b_2$, Lender 1 will get the A_{hh}^1 and A_{hl}^1 segments, and Lender 2 will get only the A_{lh}^2 segment of borrowers. We do not need to consider the case where $b_1 = b_2$ since this

will happen with probability 0 under the mixed strategy setting.¹

Lender 1's expected payoff can be written as:

$$\begin{aligned} E[\Pi_1^{rn}] &= (1 - F_2(b)) (N_{hh}^1 (V_{hh}^1 b - 1 + V_{hh}^1) + N_{hl}^1 (V_{hl}^1 b - 1 + V_{hl}^1)) \\ &\quad + F_2(b) N_{hl}^1 (V_{hl}^1 b - 1 + V_{hl}^1) \end{aligned} \quad (\text{B.17})$$

Lender 2' expected payoff can be written as:

$$\begin{aligned} E[\Pi_2^{rn}] &= (1 - F_1(b)) (N_{hh}^2 (V_{hh}^2 b - 1 + V_{hh}^2) + N_{lh}^2 (V_{lh}^2 b - 1 + V_{lh}^2)) \\ &\quad + F_1(b) N_{lh}^2 (V_{lh}^2 b - 1 + V_{lh}^2) \end{aligned} \quad (\text{B.18})$$

Facing the competitor's strategy, Lender 1 (Lender 2) should be indifferent in any b in the support of $F_1(b)$ and $F_2(b)$. Mathematically:

$$E[\Pi_1^{rn}](b) = k_1 \quad (\text{B.19})$$

$$E[\Pi_2^{rn}](b) = k_2 \quad (\text{B.20})$$

Additionally, consider the maximum b that can be set by both firm is \bar{b} , we have

$$F_1(\bar{b}) = 1 \quad (\text{B.21})$$

$$F_2(\bar{b}) = 1 \quad (\text{B.22})$$

Another condition that need to be satisfied is that the CDF $F_1(b)$ should be greater than or equal to 0 in any region where a positive probability is assigned to $F_2(b)$. Similarly, the CDF $F_2(b)$ should be greater than or equal to 0 in any region where a positive probability is assigned to $F_1(b)$. The logic is that setting $b_1 = \min(b_2)$ ensures that Lender 1 gets the common segment with probability 1, so Lender 1 has no incentive to further reduce b_1 . Similarly, setting $b_2 = \min(b_1)$ ensures that Lender 2 gets the common segment with probability 1 so Lender 2 has no incentive to further reduce b_2 . This requires:

$$F_1(\underline{b}) = 0 \quad (\text{B.23})$$

$$F_2(\underline{b}) = 0 \quad (\text{B.24})$$

Solving Equations (B.17) - (B.24), we get each lender's equilibrium strategy in the R-N sub-game. The mixed strategy equilibrium in this sub-game is summarized in Lemma 8.

¹It can be shown that there is no mass point on b_2 's distribution.

B.3.4 Proof of Proposition 4

We first prove that $\Pi_{rn}^1 > \Pi_{nn}$. Expand and rewrite Π_{nn} and Π_{rn}^1 as

$$\Pi_{nn} = -P_c (\bar{b}P_b\theta (P_b (P_c - 1) + P_c (P_b - 1)) + P_b P_c (P_b - 1) (\theta - 1) + (P_b - 1)^2 (P_c - 1) (\theta - 1))$$

$$\Pi_{rn}^1 = -\bar{b}P_b\theta (P_b (P_c - 1) + P_c (P_b - 1)) - P_b P_c (P_b - 1) (\theta - 1) - (P_b - 1)^2 (P_c - 1) (\theta - 1)$$

It is straightforward to see that $\Pi_{rn}^1 > P_c \Pi_{rn}^1 = \Pi_{nn}$.

We next prove that $\Pi_{rn}^2 > \Pi_{rr}$ if and only if $P_b > P_b^0$: Expand and rewrite the two terms as:

$$\Pi_{rn}^2 = -P_c (\bar{b}P_b\theta (P_b (P_c - 1) + P_c (P_b - 1)) + P_b P_c (P_b - 1) (\theta - 1) + (P_b - 1)^2 (P_c - 1) (\theta - 1))$$

$$\Pi_{rr} = P_b (-\bar{b}P_b\theta + \bar{b}\theta - P_b\theta + P_b + \theta - 1)$$

We first prove that the two functions, Π_{rn}^2 and Π_{rr} cross at most once. Take the first derivative w.r.t. P_b . The goal is to show $\frac{\partial \Pi_{rn}^2}{\partial P_b} \geq \frac{\partial \Pi_{rr}}{\partial P_b}$ for any value of P_b .

$$\frac{\partial \Pi_{rn}^2}{\partial P_b} = P_c (-4\bar{b}P_b P_c \theta + 2\bar{b}P_b \theta + \bar{b}P_c \theta - 4P_b P_c \theta + 4P_b P_c + 2P_b \theta - 2P_b + 3P_c \theta - 3P_c - 2\theta + 2)$$

$$\frac{\partial \Pi_{rr}}{\partial P_b} = -2\bar{b}P_b \theta + \bar{b}\theta - 2P_b \theta + 2P_b + \theta - 1$$

Since it is difficult to compare these two values directly, we go one step further and take the second derivative.

$$\frac{\partial^2 \Pi_{rn}^2}{\partial P_b^2} = 2P_c (-2\bar{b}P_c \theta + \bar{b}\theta - 2P_c \theta + 2P_c + \theta - 1)$$

$$\frac{\partial^2 \Pi_{rr}}{\partial P_b^2} = -2\bar{b}\theta - 2\theta + 2$$

We next prove that $\frac{\partial^2 \Pi_{rn}^2}{\partial P_b^2} \leq \frac{\partial^2 \Pi_{rr}}{\partial P_b^2}$: rewrite $\frac{\partial^2 \Pi_{rn}^2}{\partial P_b^2} = 2P_c(2P_c - 1)(1 - \theta - \bar{b})$, so we have $\frac{\partial^2 \Pi_{rn}^2}{\partial P_b^2} / \frac{\partial^2 \Pi_{rr}}{\partial P_b^2} = P_c(2P_c - 1) \leq 1$. Now since $\frac{\partial^2 \Pi_{rn}^2}{\partial P_b^2} \leq \frac{\partial^2 \Pi_{rr}}{\partial P_b^2}$, we only need to show $\frac{\partial \Pi_{rn}^2}{\partial P_b} \geq \frac{\partial \Pi_{rr}}{\partial P_b}$ holds at the right most point, that is, $\frac{\partial \Pi_{rn}^2}{\partial P_b} |_{P_b=1} \leq \frac{\partial \Pi_{rr}}{\partial P_b} |_{P_b=1}$. Expand and rewrite these two terms:

$$\frac{\partial \Pi_{rn}^2}{\partial P_b} |_{P_b=1} = P_c (-3\bar{b}P_c \theta + 2\bar{b}\theta - P_c \theta + P_c)$$

$$\frac{\partial \Pi_{rr}}{\partial P_b} |_{P_b=1} = -\bar{b}\theta - \theta + 1$$

The latter is invariant in P_c while the former is a quadratic function of P_c . Specifically, the former first increases and then decreases in P_c , and therefore, the minimum value is reached at either $P_c = 0.5$ or at $P_c = 1$. Since $\frac{\partial \Pi_{rn}^2}{\partial P_b} |_{P_b=1, P_c=0.5} = (\bar{b}\theta - \theta + 1)/4 > 0 > \frac{\partial \Pi_{rn}^2}{\partial P_b} |_{P_b=1}$ and $\frac{\partial \Pi_{rn}^2}{\partial P_b} |_{P_b=1, P_c=1} = -2\bar{b}\theta - 2\theta + 2 = \frac{\partial \Pi_{rn}^2}{\partial P_b} |_{P_b=1}$, $\frac{\partial \Pi_{rn}^2}{\partial P_b} |_{P_b=1} \geq \frac{\partial \Pi_{rn}^2}{\partial P_b} |_{P_b=1}$ for any value of P_c .

Combine the facts that $\frac{\partial^2 \Pi_{rn}^2}{\partial P_b^2} \leq \frac{\partial^2 \Pi_{rr}}{\partial P_b^2}$ and $\frac{\partial \Pi_{rn}^2}{\partial P_b} |_{P_b=1} \geq \frac{\partial \Pi_{rr}}{\partial P_b} |_{P_b=1}$, we get $\frac{\partial \Pi_{rn}^2}{\partial P_b} \geq \frac{\partial \Pi_{rr}}{\partial P_b}, \forall P_b$. This is sufficient to ensure that the two curves intersect at most once. We then plug the expression of P_b^0 into Π_{rn}^2 and Π_{rr} and get

$$\Pi_{rn}^2 |_{P_b=P_b^0} = \Pi_{rr} |_{P_b=P_b^0}$$

Thus we have $\Pi_{rn}^2 > \Pi_{rr}$ if and only if $P_b > P_b^0$

B.3.5 Proof of Proposition 5

We first prove that $\frac{\partial P_b^0}{\partial P_c} < 0$. Rewrite the expression of P_b^0 by collecting the P_c terms as:

$$P_b^0 = \frac{P_c (\bar{b}\theta + 3\theta - 3) + \bar{b}\theta + \theta - 1 + \sqrt{P_c^2 (\bar{b}\theta - \theta + 1)^2 + 2P_c (\bar{b}\theta + \theta - 1)^2 + (\bar{b}\theta + \theta - 1)^2}}{4P_c (\bar{b}\theta + \theta - 1) + 2(\bar{b}\theta + \theta - 1)} \quad (\text{B.25})$$

Let $x = \bar{b}\theta - \theta + 1$ and $y = \bar{b}\theta + \theta - 1$. Rewrite Equation B.25 as:

$$P_b^0 = \frac{\frac{P_c(2y-x)}{y} + 1 + \sqrt{\frac{P_c^2 x^2}{y^2} + 2P_c + 1}}{2(2P_c + 1)} \quad (\text{B.26})$$

Now take the derivative w.r.t. P_c . Note that P_c is not contained in either x or y .

$$\frac{\partial P_b^0}{\partial P_c} = \frac{x^2 P_c - x \sqrt{x^2 P_c^2 + 2y^2 P_c + y^2} - 2y^2 P_c - y^2}{2y \sqrt{x^2 P_c^2 + 2y^2 P_c + y^2} (2P_c + 1)^2} \quad (\text{B.27})$$

Also notice that according to Assumption 1, $y > 0$. Furthermore, $x = y + 2(1 - \theta) > y > 0$. It is easy to see that the denominator of the right-hand side expression in Equation B.27 is greater than 0. Rewrite the numerator as $x(\sqrt{x^2 P_c^2} - \sqrt{x^2 P_c^2 + 2y^2 P_c + y^2}) + y^2(-2P_c - 1)$, and we can see both terms in the numerator are negative. Thus, we have $\frac{\partial P_b^0}{\partial P_c} < 0$.

We next prove that $\frac{\partial P_b^0}{\partial \theta} > 0$. Rewrite the expression of P_b^0 while collecting the terms involving θ as

$$P_b^0 = \frac{\theta(\bar{b}P_c + \bar{b} + 3P_c + 1) - 3P_c - 1 + \sqrt{\Delta'}}{\theta(2\bar{b}P_c + \bar{b} + 2P_c + 1) - 2P_c - 1} \quad (\text{B.28})$$

where $\Delta' = \theta^2(\bar{b}^2(P_c^2 + 2P_c + 1) + 2\bar{b}(-P_c^2 + 2P_c + 1) + (P_c^2 + 2P_c + 1)) + \theta(-2\bar{b}(-P_c^2 + 2P_c + 1) - 2(P_c^2 + 2P_c + 1)) + P_c^2 + 2P_c + 1$. Rewrite Δ' as

$$\Delta' = (\theta\bar{b}P_c - \theta P_c + P_c)^2 + (2P_c + 1)(\bar{b}\theta + \theta - 1)^2 \quad (\text{B.29})$$

Note that for the second term $\frac{\partial((2P_c+1)(\bar{b}\theta+\theta-1)^2)}{\partial\theta} = 2(2P_c + 1)(\bar{b}\theta + \theta - 1)(\bar{b} + 1) > 0$ according to Assumption 1. Combine the fact that the partial derivative of the denominator in Equation B.28 w.r.t. θ equals $2\bar{b}P_c + \bar{b} + 2P_c + 1$, which is greater than 0, we have

$$\begin{aligned} \frac{\partial P_b^0}{\partial\theta} &= \frac{\partial \frac{\theta(\bar{b}P_c + \bar{b} + 3P_c + 1) - 3P_c - 1 + \sqrt{\Delta'}}{\theta(2\bar{b}P_c + \bar{b} + 2P_c + 1) - 2P_c - 1}}{\partial\theta} \\ &> \frac{\partial \frac{\theta(\bar{b}P_c + \bar{b} + 3P_c + 1) - 3P_c - 1 + \sqrt{(\theta\bar{b}P_c - \theta P_c + P_c)^2}}{\theta(2\bar{b}P_c + \bar{b} + 2P_c + 1) - 2P_c - 1}}{\partial\theta} \\ &= \frac{\partial \frac{\theta(\bar{b}P_c + \bar{b} + 3P_c + 1) - 3P_c - 1 + (\theta\bar{b}P_c - \theta P_c + P_c)}{\theta(2\bar{b}P_c + \bar{b} + 2P_c + 1) - 2P_c - 1}}{\partial\theta} \\ &= \frac{\partial \frac{\theta(2\bar{b}P_c + \bar{b} + 2P_c + 1) - 2P_c - 1}{\theta(2\bar{b}P_c + \bar{b} + 2P_c + 1) - 2P_c - 1}}{\partial\theta} \\ &= 0 \end{aligned}$$

Thus $\frac{\partial P_b^0}{\partial\theta} > 0$.

B.3.6 Proof of Proposition 6

In this proof, we first compare total surplus in each of the sub-games resulting from the two lenders' algorithm revealing decisions, and then subtract the lenders' equilibrium payoff from the total surplus and compare the borrower (consumer) surplus. We calculate the total surplus by adding together the social surplus generated by borrowers from different segments. When a H type borrower gets the loan, a ($TS_h = M_h - m$) amount of social welfare is generated. When a L type borrower gets the loan, a ($TS_l = M_l - m - 1$) amount of social welfare is generated (i.e., the lender loses 1, the loan amount, while the borrower loses m , the negative impact on the credit score). When a borrower is rejected, a $-m$ amount of social surplus is generated. To calculate the social surplus generated in each sub-game, we need to find out the number of H type borrowers who are approved ($N_{H,a}$), the number of L type borrowers who are approved ($N_{L,a}$), and the number of borrowers who are rejected (N_r). We do not care about the borrowers who do not apply since they will generate 0 surplus to the social welfare. For the R-R case, there are 4 segments of borrowers: HH HL LH and LL .

Borrowers in the HH HL and LH segments will apply and will be approved. Thus we have $N_{H,a}^{rr} = N_{hh}^{rr}V_{hh}^{rr} + 2N_{hl}^{rr}V_{hl}^{rr}$, $N_{L,a}^{rr} = N_{hh}^{rr}(1 - V_{hh}^{rr}) + 2N_{hl}^{rr}(1 - V_{hl}^{rr})$, and $N_r^{rr} = 0$. (We use the superscript rr , nn , and rn to denote the R-R, N-N, and R-N case respectively.) The total surplus generated in the R-R case is

$$TS^{rr} = N_{H,a}^{rr}TS_h + N_{L,a}^{rr}TS_l \quad (\text{B.30})$$

Similarly, in the N-N case, we have

$$TS^{nn} = N_{H,a}^{nn}TS_h + N_{L,a}^{nn}TS_l + N_r^{nn}(-m) \quad (\text{B.31})$$

where

$$\begin{aligned} N_{H,a}^{nn} &= P_b P_c \theta (-2P_b P_c + P_b + P_c + 1) \\ N_{L,a}^{nn} &= P_c (P_b - 1) (\theta - 1) (2P_b P_c - P_b (P_c - 1) - P_c (P_b - 1) + 2(P_b - 1)(P_c - 1)) \\ N_r^{rr} &= P_b^2 (-3P_c^2 + 5P_c - 2) + 4P_b P_c \theta (P_c - 1) + P_b P_c (P_c - 3) + 2P_b - 2P_c \theta (P_c - 1) \end{aligned}$$

In the R-N case, we have

$$TS^{rn} = N_{H,a}^{rn}TS_h + N_{L,a}^{rn}TS_l + N_r^{rn}(-m) \quad (\text{B.32})$$

where

$$\begin{aligned} N_{H,a}^{rn} &= P_b \theta (-P_b P_c + P_c + 1) \\ N_{L,a}^{rn} &= (P_b - 1) (\theta - 1) (2P_b P_c - P_b (P_c - 1) - P_c (P_b - 1) + (P_b - 1)(P_c - 1)) \\ N_r^{rn} &= (P_c - 1) (P_b^2 (\theta - 1) - \theta (P_b - 1)^2) \end{aligned}$$

Next, we show that $TS^{rr} > TS^{rn} > TS^{nn}$. The comparison is straightforward but cumbersome and thus the details are omitted here. The intuition is as follows. Note that $N_r^{nn} > N_r^{rn} > N_r^{rr}$ and $N_{H,a}^{nn} + N_{L,a}^{nn} < N_{H,a}^{rn} + N_{L,a}^{rn} < N_{H,a}^{rr} + N_{L,a}^{rr}$. In other words, in the N-N case, the number of borrowers who are approved is the lowest and the number of borrowers who are rejected is the highest. In the R-R case, the number of borrowers who are approved is the highest and the number of borrowers who are rejected is the lowest. The R-N case is between the two extremes. Since rejected borrowers will generate a negative social surplus, and accepted borrowers on average will generate a positive social surplus, we have $TS^{rr} > TS^{rn} > TS^{nn}$.

We next subtract the lenders' equilibrium profits from the social surplus to compute

borrower surplus:

$$CS^{rr} = TS^{rr} - 2\Pi_{rr} \quad (\text{B.33})$$

$$CS^{nn} = TS^{nn} - 2\Pi_{nn} \quad (\text{B.34})$$

$$CS^{rn} = TS^{rn} - \Pi_{rn}^1 - \Pi_{rn}^1 \quad (\text{B.35})$$

where TS^{rr} , TS^{nn} , and TS^{rn} are defined in Equations B.30, B.31, and ?? respectively. Π_{rr} , Π_{nn} , and Π_{rn} are defined in Lemma 6, 7, and 8 respectively. We can plug the expression of total surplus and lenders' profit into these expressions and compare them, and we will get $CS^{rr} > CS^{rn} > CS^{nn}$. Again, the comparison is straightforward but cumbersome, so we omitted the details here. The intuition of $CS^{rr} > CS^{rn}$ is simple: (1) The number of borrowers who are approved in the R-R case is larger than in the R-N case. (2) $F^{rr}(b) > F_1(b)$ and $F^{rr}(b) > F_2(b)$, which means on average borrowers face lower interest rates in the R-R case than in the R-N case. Points (1) and (2) above ensure that borrowers' welfare in the R-R case is higher than in the R-N case. We next discuss why $CS^{rn} > CS^{nn}$. Comparing the difference in total surplus and the difference in lenders' profit in the N-N and R-N cases, we have $TS^{rn} - TS^{nn} > \Pi_{rn}^1 + \Pi_{rn}^2 - 2\Pi_{nn} = \Pi_{rn}^1 - \Pi_{nn}$ (Note that $\Pi_{rn}^2 = \Pi_{nn}$), which means the lenders only take a portion of the increased surplus, and the remaining is left to the borrowers.

B.3.7 Proof of Proposition 7

We consider sub-games resulting from the lenders' decisions on algorithm upgrade (NU - $NU, NU-U, U-NU, U-U$), where NU stands for "not upgrade" and the algorithm's accuracy stays at P_b , and U stands for "upgrade" and the algorithm's accuracy rises to P_b^* . The equilibrium payoff in the two symmetric cases, $NU-NU$ and $U-U$, have already been calculated and shown in Lemma 6. Since the two lenders' decisions on revealing is symmetric under mandatory revealing, the lenders' equilibrium payoff in $NU-U$ case mirrors the payoff in the $U-NU$ case. Compared with the $NU-NU$ case, in the $NU-U$ case, the only things that will change are the number of borrowers in each segment. Similar to the analysis in Appendix B.3.1, we can find the lenders' equilibrium payoff when their algorithms' accuracy are P_b and P_b^* respectively. The payoff matrix of the "algorithm upgrading" game is shown in the following table:

		Lender 2	
		Not Upgrade	Upgrade
Lender 1	Not Upgrade	$(\Pi_{NU-NU}^{rr}, \Pi_{NU-NU}^{rr})$	$(\Pi_{NU-U}^{rr}, \Pi_{U-NU}^{rr})$
	Upgrade	$(\Pi_{U-NU}^{rr}, \Pi_{NU-U}^{rr})$	$(\Pi_{U-U}^{rr}, \Pi_{U-U}^{rr})$

Where

$$\Pi_{NU-NU}^{rr} = P_b (-\bar{b}P_b\theta + \bar{b}\theta - P_b\theta + P_b + \theta - 1)$$

$$\Pi_{U-U}^{rr} = P_b^* (-\bar{b}P_b^*\theta + \bar{b}\theta - P_b^*\theta + P_b^* + \theta - 1)$$

$$\Pi_{NU-U}^{rr} = P_b (-\bar{b}P_b\theta + \bar{b}\theta - P_b\theta + P_b + \theta - 1)$$

$$\Pi_{U-NU}^{rr} = P_b^* (-\bar{b}P_b^*\theta + \bar{b}\theta - P_b^*\theta + P_b^* + \theta - 1)$$

We focus on the symmetric mixed strategy equilibrium and assume both lenders choose to upgrade with probability P^M . The following equation must hold to ensure both lenders are indifferent between upgrading or not:

$$(1 - P^M)\Pi_{NU-NU} + P^M\Pi_{NU-U}^1 = (1 - P^M)\Pi_{NU-U}^2 + P^M\Pi_{U-U} \quad (\text{B.36})$$

Solve Equation B.36 and we get $P^M = \frac{1-P_b}{P_b^*}$.

B.3.8 Proof of Proposition 8

Assume in this case Lender 1 reveals its algorithm but Lender 2 does not. We solve for lenders equilibrium payoff in sub-games followed by every combination of the two lenders' upgrading decisions, the payoff matrix of the game is shown in the following table:

		Lender 2	
		Not Upgrade	Upgrade
Lender 1	Not Upgrade	$(\Pi_{NU-NU}^{rn1}, \Pi_{NU-NU}^{rn2})$	$(\Pi_{NU-U}^{rn1}, \Pi_{U-NU}^{rn2})$
	Upgrade	$(\Pi_{U-NU}^{rn1}, \Pi_{NU-U}^{rn2})$	$(\Pi_{U-U}^{rn1}, \Pi_{U-U}^{rn2})$

where

$$\left\{ \begin{array}{l}
\Pi_{NU-NU}^{rn1} = -\bar{b}P_b\theta (P_b(P_c - 1) + P_c(P_b - 1)) - (P_b + P_c - 1)(\theta - 1)(P_b - 1) \\
\Pi_{NU-NU}^{rn2} = P_c(P_b^2(-2\bar{b}P_c\theta + \bar{b}\theta - 2P_c\theta + 2P_c + \theta - 1) + P_b(\bar{b}P_c\theta + (\theta - 1)(3P_c - 2)) \\
\quad - P_c\theta + P_c + \theta - 1) \\
\Pi_{U-NU}^{rn1} = -\bar{b}P_b^*\theta (P_b(P_c - 1) + P_c(P_b - 1)) - P_bP_c(P_b^* - 1)(\theta - 1) \\
\quad - (P_b - 1)(P_b^* - 1)(P_c - 1)(\theta - 1) \\
\Pi_{NU-U}^{rn2} = P_c(P_b^2(P_b^*(-2\bar{b}P_c\theta + \bar{b}\theta - 2P_c\theta + 2P_c + \theta - 1) + 2P_c\theta - 2P_c - \theta + 1) \\
\quad + P_b(P_b^*(\bar{b}P_c\theta + P_c\theta - P_c - \theta + 1) - P_c\theta + P_c) + P_b^*(\theta - 1))/P_b^* \\
\Pi_{NU-U}^{rn1} = -\bar{b}P_b\theta (P_b^*(P_c - 1) + P_c(P_b^* - 1)) - P_b^*P_c(P_b - 1)(\theta - 1) \\
\quad - (P_b - 1)(P_b^* - 1)(P_c - 1)(\theta - 1) \\
\Pi_{U-NU}^{rn2} = P_c(P_b(P_b^{*2}(-2\bar{b}P_c\theta + \bar{b}\theta + (\theta - 1)(1 - 2P_c)) + P_b^*(\bar{b}P_c\theta + (\theta - 1)(P_c - 1)) \\
\quad + \theta - 1) + P_b^{*2}(2P_c\theta - 2P_c - \theta + 1) + P_b^*(-P_c\theta + P_c))/P_b \\
\Pi_{U-U}^{rn1} = -\bar{b}P_b^*\theta (P_b^*(P_c - 1) + P_c(P_b^* - 1)) - P_b^*P_c(P_b^* - 1)(\theta - 1) \\
\quad - (P_b^* - 1)^2(P_c - 1)(\theta - 1) \\
\Pi_{U-U}^{rn2} = P_c(-2\bar{b}P_b^{*2}P_c\theta + \bar{b}P_b^{*2}\theta + \bar{b}P_b^*P_c\theta - 2P_b^{*2}P_c\theta + 2P_b^{*2}P_c + P_b^{*2}\theta - P_b^{*2} \\
\quad + 3P_b^*P_c\theta - 3P_b^*P_c - 2P_b^*\theta + 2P_b^* - P_c\theta + P_c + \theta - 1)
\end{array} \right.$$

We first prove that Lender 1 has a dominate strategy “U” by showing $\Pi_{U-NU}^{rn1} > \Pi_{NU-NU}^{rn1}$ and $\Pi_{U-U}^{rn1} - \Pi_{NU-U}^{rn1}$.

$$\Pi_{U-NU}^{rn1} - \Pi_{NU-NU}^{rn1} = (P_b - P_b^*) (\bar{b}\theta (P_b(P_c - 1) + P_c(P_b - 1)) + (\theta - 1)(P_bP_c + (P_b - 1)(P_c - 1)))$$

Since $P_b < P_b^*$, $P_b < 1$, $P_c < 1$ and $\theta < 1$, $\Pi_{U-NU}^{rn1} - \Pi_{NU-NU}^{rn1} > 0$. Further,

$$\Pi_{U-U}^{rn1} - \Pi_{NU-U}^{rn1} = (P_b - P_b^*) (\bar{b}\theta (P_b^*(P_c - 1) + P_c(P_b^* - 1)) + (\theta - 1)(P_b^*P_c + (P_b^* - 1)(P_c - 1)))$$

Again, since $P_b < P_b^*$, $P_b^* < 1$, $P_c < 1$ and $\theta < 1$, $\Pi_{U-U}^{rn1} - \Pi_{NU-U}^{rn1} > 0$.

As Lender 1 always chooses to upgrade the algorithm, Lender 2 makes the upgrading decision by comparing Π_{NU-U}^{rn2} with Π_{U-U}^{rn2} . Let $\delta = \Pi_{U-U}^{rn2} - \Pi_{NU-U}^{rn2}$. It is easy to check that δ is linearly decreasing in θ . Since $\delta|_{\theta=\theta^0} = 0$, $\delta > 0$ when $\theta < \theta^0$. Notice that θ^0 is decreasing in P_c , and $\theta^0|_{P_c=P_c^0} = 1$, which means $\theta^0 > 1$ when $P_c < P_c^0$. Thus $\theta < \theta^0$ will always hold.

B.3.9 Proof of Proposition 9

In the mandatory revealing scenario, both lenders choose to upgrade the algorithm with probability P_M , that is, with probability P_M^2 both lenders' algorithms' accuracy will be P_b^* , with probability $(1 - P_M)^2$ both lenders' algorithms' accuracy will be P_b , and with probability $2P_M(1 - P_M)$ one lender's accuracy will be P_b and the other's will be P_b^* . Borrower surplus is thus the expected total surplus minus the lenders' expected profit, as shown in Equation B.37.

$$CS^{rr*} = (1 - P_M)^2 (TS_{P_b, P_b}^{rr} - 2\Pi_{P_b, P_b}^{rr}) + P_M^2 (TS_{P_b^*, P_b^*}^{rr} - 2\Pi_{P_b^*, P_b^*}^{rr}) + 2P_M(1 - P_M) (TS_{P_b, P_b^*}^{rr} - \Pi_{P_b, P_b}^{rr} - \Pi_{P_b, P_b^*}^{rr}) \quad (\text{B.37})$$

where

$$\left\{ \begin{array}{l} TS_{P_b, P_b}^{rr} = P_b(2 - P_b)(-\theta(m - M_h) + (\theta - 1)(m - M_l + 1)) \\ TS_{P_b, P_b^*}^{rr} = (\theta(m - M_h) - (\theta - 1)(m - M_l + 1))(-P_bP_b^* + P_b(P_b^* - 1) + P_b^*(P_b - 1)) \\ TS_{P_b^*, P_b^*}^{rr} = P_b^*(2 - P_b^*)(-\theta(m - M_h) + (\theta - 1)(m - M_l + 1)) \\ \Pi_{P_b, P_b}^{rr} = P_b(-\bar{b}P_b\theta + \bar{b}\theta - P_b\theta + P_b + \theta - 1) \\ \Pi_{P_b, P_b^*}^{rr} = P_b^*(-\bar{b}P_b\theta + \bar{b}\theta - P_b\theta + P_b + \theta - 1) \\ \Pi_{P_b^*, P_b^*}^{rr} = P_b^*(-\bar{b}P_b^*\theta + \bar{b}\theta - P_b^*\theta + P_b^* + \theta - 1) \end{array} \right.$$

In the voluntary scenario, when $\theta < \theta^0$, both lenders upgrade their algorithms to an accuracy of P_b^* . Thus borrower surplus can be calculated as:

$$CS^{rn*} = TS_{P_b^*, P_b^*}^{rn} - \Pi_{P_b^*, P_b^*}^{rn1} - \Pi_{P_b^*, P_b^*}^{rn2} \quad (\text{B.38})$$

where

$$\left\{ \begin{array}{l} TS_{P_b^*, P_b^*}^{rn} = -m(P_c - 1)(P_b^{*2}(\theta - 1) - \theta(P_b^* - 1)^2) \\ \quad + P_b^*\theta(m - M_h)(-P_b^*P_c + P_b^*(P_c - 1) + 2P_c(P_b^* - 1)) \\ \quad - (P_b^* - 1)(\theta - 1)(m - M_l + 1)(2P_b^*P_c - P_c(P_b^* - 1) + (P_b^* - 1)(P_c - 1)) \\ \Pi_{P_b^*, P_b^*}^{rn1} = -\bar{b}P_b^*\theta(P_b^*(P_c - 1) + P_c(P_b^* - 1)) - P_b^*P_c(P_b^* - 1)(\theta - 1) \\ \quad - (P_b^* - 1)^2(P_c - 1)(\theta - 1) \\ \Pi_{P_b^*, P_b^*}^{rn2} = P_c(-2\bar{b}P_b^{*2}P_c\theta + \bar{b}P_b^{*2}\theta + \bar{b}P_b^*P_c\theta - 2P_b^{*2}P_c\theta + 2P_b^{*2}P_c + P_b^{*2}\theta - P_b^{*2} \\ \quad + 3P_b^*P_c\theta - 3P_b^*P_c - 2P_b^*\theta + 2P_b^* - P_c\theta + P_c + \theta - 1) \end{array} \right.$$

Define $\delta = CS^{rr*} - CS^{rn*}$. Combining and rearranging terms, we can see that δ is a quadratic function in P_c : $\delta = AP_c^2 + BP_c + C$, where A , B , and C are defined in Proposition

9. We can then check that $\delta|_{P_c=0} = C < 0$, and $\delta|_{P_c=1} = A + B + C > 0$ in the parameter ranges defined in Assumptions 1 and 2. Thus one of the roots of the quadratic function, $P_c^0 = \frac{-B + \sqrt{B^2 - 4AC}}{2A}$, must be in range $(0, 1)$. Thus $\delta > 0$ if $P_c > P_c^0$.

Appendix C

Appendix for Chapter 3

C.1 Experiment Outcomes: Alternative Settings

C.1.1 More Advanced RL Algorithms

We examine whether our results are robust when we replace the Q-learning algorithm with two types of more advanced RL algorithms: *Deep Q Network* (DQN) and *Advanced Actor-Critic Method* (A2C). Specifically, we let seller 1 use an advanced RL algorithm, and compare seller 2's equilibrium price and profit when using the same advanced RL algorithm vs using a rule-based algorithm to compete. For each advanced algorithm (DQN or A2C), we run the experiment 100 times and report the mean and standard deviation of equilibrium price and profit in Table C.1. The result shows that our main result holds regardless of which specific type of reinforcement learning algorithm is used.

C.1.2 Multiple Players in the Market

In this section, we extend our baseline model to an oligopoly setting ($N > 2$). We study the case where there are three sellers in the market and examine when seller 1 and seller 2 use Q-learning pricing algorithms, how would seller 3's equilibrium price and equilibrium profit

Panel A: Deep Q Network (DQN)		
	DQN	Rule-based Algorithm
Price	1.680(0.024)	1.894(0.011)
Profit 2	0.291(0.015)	0.335(0.002)

Panel B: Advanced Actor-Critic Method (A2C)		
	A2C	Rule-based Algorithm
Price	1.586(0.043)	1.844(0.050)
Profit	0.265(0.019)	0.356(0.013)

Table C.1: Simple Rule-based Algorithm Competes with More Advanced Reinforcement Learning Algorithms (DQN and A2C)

differ when it decides between a Q-learning algorithm and a simple rule-based algorithm.

In this extension, we compare seller 3’s equilibrium price and equilibrium profit across 4 scenarios: when it uses 1) the same Q-learning algorithm as sellers 1 and 2, 2) a simple rule-based algorithm that undercuts seller 1’s price by one price grid, 3) a simple rule-based algorithm that undercuts seller 2’s price by one price grid, and 4), a simple rule-based algorithm that undercuts the lowest price (between sellers 1 and 2) by one price grid.

We run each experiment 100 times and the outcomes are reported in Table C.2.

	Q-learning	Undercut seller 1	Undercut seller 2	Undercut the lowest
Price	1.464(0.030)	1.663(0.059)	1.667(0.069)	1.654(0.055)
Profit	0.148(0.009)	0.175(0.015)	0.172(0.017)	0.202(0.014)

Table C.2: Equilibrium Outcomes When Seller 3 Uses Different Algorithm

C.2 Structural Demand Model Details

C.2.1 Data

The dataset employed in our study is transaction-level data from a prominent Chinese e-commerce platform, JD.com. The platform’s operational mechanism involves presenting consumers with products from multiple sellers in each product category, accompanied by pertinent details such as prices, ratings, and key product features. Upon selecting a product

listing by clicking on it, consumers can obtain more comprehensive information. Specifically, our dataset captures the click and purchase behavior of over 2.5 million consumers within a specific product category throughout March 2018.

The data contains 4 tables: *SKU Table* describes the characteristics of all 31,868 SKUs that belong to a single product category receiving at least one click during March 2018. For each SKU, it records the unique identifier SKU ID, whether this SKU is from a first-party seller or a third-party seller. *Consumer Table* describes the characteristics of each consumer, including the unique identifier consumer ID and a bunch of consumer demographics like age, gender, marital status etc. *Click Table* records the information of all the click events, including the consumer ID that initiates the click, the SKU ID that she clicks on, and time of the click. *Purchase Table* records the information of all the purchase events, including consumer ID, SKU ID, time of the purchase, purchase quantity, listed prices, and effective prices.

Instead of studying the competition among all SKUs, we focus on a small subset of them. The motivation of this choice is two-fold: 1), the sales of these products follow a ‘long-tail distribution’, where over 96% of SKUs have single-digit daily sales. We aim to concentrate on the subset of products that have a sizeable volume of sales. 2), we focus on a group of products that are close substitutes, rather than analyzing the competition between dissimilar products. For example, we would like to study the competition between electric shavers from two brands, instead of the competition between an electric shaver and a water flosser. We thus pre-process the data and identify a small subset in which all SKUs have decent sales volume and are close substitutes. In Appendix C.3, we provide details on how we determine if two products are close substitutes and construct the subset.

Our study concentrates on a specific subset of products that exhibit substantial sales volume and are identified as close substitutes, to derive an accurate demand model. This category comprises 5 distinct products, and these products have been clicked by more than 25,000 consumers. The basic information of these 5 products and the involved consumers are

shown in Table C.3 and C.4, respectively. For each purchase, we observe the purchase time and the original listed price, and the price discount that a consumer gets. Interestingly, we find the effective price (listed price minus discounts)¹ is almost identical for all consumers at any given time, indicating very little first- or third-degree price discrimination. The most common age segment is 26-35 years old and the most common education level is a Bachelor’s degree. There are more females in the sample than males. The sample is approximately equally split between single and married customers.

	P1	P2	P3	P4	P5
Average price (¥)	68.84	52.27	57.94	71.54	69.05
1st/3rd party seller	1st	1st	3rd	3rd	1st
Unique clicks	14,171	15,478	5,351	4,512	4,683
Unique orders	3,942	3,490	1,247	1,179	726

Table C.3: Product Information

Table C.5 provides a summary of the click and purchase behavior of consumers. Several key observations can be made from this table. First, 20.1% of consumers click on 3 or more products and 25.5% of consumers click on 2 products before making a purchase decision, suggesting that a substantial search cost is involved in the search-then-purchase process. Second, less than 1% of consumers purchase more than one product, indicating that it is highly unlikely that the product subset we have selected includes complementary items.

C.2.2 Testing Search Models

In this study, we examine two major models that describe consumers’ search-then-purchase behavior: the sequential search model (McCall, 1970, Weitzman, 1979) and the non-sequential search model (Stigler, 1961, Mehta et al., 2003, De Los Santos et al., 2012). While consumers’ actual search behavior may fall somewhere between these two models, it’s important to determine which model is more appropriate based on the context (De Los Santos et al., 2012).

¹Here we only consider direct discount but not volume discount since very few consumers buy more than one product.

Key Characteristics	Count(Percentage)
Number of consumers	25,339
Age	
≤ 25 years old	6,335 (25.0%)
26 - 35 years old	11,567 (45.6%)
36 - 45 years old	5,604 (22.1%)
46 - 55 years old	1,107 (4.4%)
≥ 56 years old	726 (2.9%)
Gender	
Male	5,657 (22.3%)
Female	19,682 (77.7%)
Marital Status	
Single	13,046 (51.5%)
Married	12,293 (48.5%)
Education Level	
Less Than High School	912 (3.6%)
High School or Equivalent	8,705 (34.4%)
Bachelor's Degree	13,620 (53.8%)
Post-graduate Degree	2,102 (8.3%)

Table C.4: Descriptive Statistics of Consumer Characteristics

Key Characteristics	Count(Percentage)
Number of consumers	25,339
Choice Set Size	
1	13,775 (54.4 %)
2	6,459 (25.5 %)
≥ 3	5,105 (20.1 %)
Number of products purchased	
0	14,953 (59.0%)
1	10,192 (40.2%)
≥ 2	194 (0.8%)

Table C.5: Descriptive Statistics of Consumer Clicks and Orders

In our scenario, where consumers search for product fit rather than prices, we employ two empirical metrics to evaluate the search model. Our findings suggest that a non-sequential search model provides a better fit for our online search and purchase context. We use metrics such as recall and the conditioned probability of purchase to support our conclusion. We find that 58% of people show recall behavior and that the order in which products are searched has no significant impact on their purchase probability. Based on these results, we construct a structural model under the non-sequential search assumption to describe the market environment.

C.2.3 A Non-sequential Search Structural Model

In this section, we present a non-sequential search structural model that captures consumer search and purchase behavior within our empirical context. There are I consumers and J sellers in the market, each seller (seller j) supplies one product (product j) with marginal cost mc_j . Each consumer demands at most one product. We specify consumer i 's utility of purchasing product j as:

$$u_{ij} = \eta_j + \beta_i p_j + \varepsilon_{ij} \quad (\text{C.1})$$

where η_j is the product (or seller) fixed effect, p_j is the price of product j . β_i is individual specific price sensitivity: $\beta_i = \Gamma z_i + \varepsilon_\beta$ where z_i is a vector of individual characteristics for consumer i , Γ is a coefficient vector where the k th entry represents the part-worth of the k th characteristics on β_i , and ε_β is an unobserved heterogeneity component which is assumed to follow a normal distribution with mean 0 and variance Σ_β . ε_{ij} is the fit value between consumer i and product j . We assume ε_{ij} is *i.i.d* distributed and follows a Type I extreme value distribution with a location parameter of zero and a scale parameter σ_ε ($\varepsilon_{ij} \sim EVT1(0, \sigma_\varepsilon)$).

Online marketplaces such as JD.com typically display important product information such as prices, ratings, number of reviews, and key features through thumbnail images or

product listings. This allows users to quickly browse and access this information without the need to click on individual products. It is assumed that consumers have access to prices (p_j) and product values (i.e., product fixed effects, η_j) for all products upon entering the market. Further, consumers know the distribution of the fit value (ε) but cannot observe its realization at this point. To obtain the exact fit value, consumers must conduct a search by clicking on a specific product and examining the detailed product description, images, videos, and consumer reviews.

We model a non-sequential search process: in the first stage, consumers form a consideration set which is a subset of all the J products, based on the information available to them when they enter the market. In the second stage, consumers click on each of the products in the consideration set, learn the fit values, and decide which product to purchase or not to purchase at all (i.e., the outside option). Next, we explain these two stages in detail.

First Stage: Form Consideration Sets

Consumers determine their consideration set by balancing the expected benefit of inspecting all products within it against the search costs involved. On the one hand, having more products in the consideration set will increase the chances of finding great fit later on. On the other hand, consumers have to incur larger search costs for larger consideration sets (e.g., time and effort spent on visiting the product information pages). We use c to denote the cost of inspecting one product, and for simplicity, we assume the search cost is homogeneous across consumers and products. We use \mathcal{S} to denote the universe of products, and consumers' consideration set is denoted as $S \subseteq \mathcal{S}$. Consumer i 's expected utility of forming a consideration set S is

$$m_{iS} = \mathbb{E}[\max_{j \in S} \{u_{ij}\}] - c|S| \quad (\text{C.2})$$

We can rewrite u_{ij} as $u_{ij} = \delta_{ij} + \varepsilon_{ij}$, where $\delta_{ij} = \eta_j + \beta_i p_j$. If we denote the CDF of the distribution of ε_{ij} as F , then the CDF of u_{ij} is $F(u - \delta_{ij})$, and the CDF of $\max_{j \in S} \{u_{ij}\}$ is

$\prod_{j \in S} F(u - \delta_{ij})$. Then,

$$\mathbb{E}[\max_{j \in S}\{u_{ij}\}] = \int_{-\infty}^{\infty} u \rho(u) du \quad (\text{C.3})$$

where $\rho(u)$ is the pdf of $\max_{j \in S}\{u_{ij}\}$ and can be expressed as

$$\rho(u) = \frac{d}{du} \left[\prod_{j \in S} F(u - \delta_{ij}) \right]$$

Plug $F(u) = \exp[-\exp(-u)]$ into Equation C.3, and we get²

$$\mathbb{E}[\max_{j \in S}\{u_{ij}\}] = \gamma + \log\left[\sum_{j \in S} \exp(\delta_{ij})\right] \quad (\text{C.4})$$

where γ is the Euler-Mascheroni constant.³

To smooth the choice set formation probability and allow consumers with similar levels of price sensitivity facing the similar product price to form different consideration sets, we follow De Los Santos et al. (2012) and assume that consumers are subject to an assessment error ζ when assessing the expected net benefit of forming consideration set S , where ζ_{iS} follows another Type I extreme value distribution with a location parameter of zero and a scale parameter of σ_{ζ} . Consumer i determine her consideration set by maximizing the expected

²Specifically, let $v = \sum_{j \in S} \exp(-u + \delta_{ij})$, then we have $dv = -v du$ and $u = \log(\sum_{j \in S} \exp(\delta_{ij})) - \log(v)$

$$\begin{aligned} \mathbb{E}[\max_{j \in S}\{u_{ij}\}] &= \int_{-\infty}^{\infty} u \frac{d}{du} [\exp[-\sum_{j \in S} \exp(-u + \delta_{ij})]] du \\ &= \int_{-\infty}^{\infty} \frac{u \sum_{j \in S} \exp(-u + \delta_{ij})}{\exp[-\sum_{j \in S} \exp(-u + \delta_{ij})]} du \\ &= \int_{-\infty}^{\infty} \frac{uv}{\exp(v)} du \\ &= - \int_0^{\infty} \frac{\log(\sum_{j \in S} \exp(\delta_{ij}))}{\exp(v)} dv + \int_0^{\infty} \frac{\log(v)}{\exp(v)} dv \\ &= \gamma + \log\left[\sum_{j \in S} \exp(\delta_{ij})\right] \end{aligned}$$

³We ignore this constant in the rest of the analysis since the constant term does not affect choice probability.

utility associated with the consideration set plus the assessment error.

$$S = \arg \max_{S \in \mathcal{S}} [m_{iS} + \zeta_{iS}] \quad (\text{C.5})$$

Consumer i 's probability of forming a consideration set S follows a logit form:

$$P_{iS} = \frac{\exp[m_{iS}/\sigma_\zeta]}{\sum_{S' \in \mathcal{S}} \exp[m_{iS'}/\sigma_\zeta]} \quad (\text{C.6})$$

Second Stage: Purchase

Upon identifying the consideration set, a consumer proceeds to examine each product within the set and evaluates their respective fit values before deciding on which product to purchase, including the outside option. The likelihood of consumer i purchasing product j , conditioned on the formation of consideration set S , can be expressed as follows:

$$P_{ij|S} = \Pr(u_{ij} > u_{ik}, \forall k \neq j \in S) \quad (\text{C.7})$$

Note that the fit value ε_{ij} is revealed to consumer i *after* she forms a consideration set and inspects all products within it. From the researchers' perspective, conditioned on the consideration set consumer i forms, the probability of the consumer purchasing product j , denoted as $P_{ij|S}$, follows a logit form:

$$P_{ij|S} = \frac{\exp[\delta_{ij}]}{\sum_{k \in S} \exp[\delta_{ik}] + \exp[\delta_0]} \quad (\text{C.8})$$

The joint probability of consumer i forming consideration set S and purchase product j is

$$P_{ijS} = P_{ij|S} P_{iS} \quad (\text{C.9})$$

The log-likelihood function of observing all consumers' consideration sets and their purchase decisions can be specified as

$$LL = \sum_i \log \hat{P}_{ijS} = \sum_i \log \hat{P}_{ij|S} \hat{P}_{iS} \quad (\text{C.10})$$

Where \hat{P}_{ijS} is the likelihood of consumer i selecting the observed consideration set and making the corresponding purchase decision. Consistent with the conventions of discrete choice demand models, the coefficients for the observable variables are identified in reference to the variability of unobserved factors. In order to achieve this, we first standardize the variance of the random utility component ε_{ij} to 1, followed by the estimation of the coefficients for observed variables and the variance of the assessment error term σ_{ζ} .

Demand Estimation

We estimate the model using the Hierarchical Bayes approach (Rossi and Allenby, 2003) to account for observed and unobserved consumer heterogeneity. The parameters to estimate are: product (seller) fixed effects η_j , search cost c , the variance of the assessment error term σ_{ζ}^2 , individual specific price sensitivity β_i , a coefficient vector that connects price sensitivity with individual characteristics Γ , the variance of the unobserved heterogeneity component Σ_{β} . The observables are: the price of the product p_j , consumer characteristics Z , consumers' choice sets Y_S , and consumers' purchase decisions Y_P . A hierarchy graph of all parameters and observables is shown in Figure C.1.

We use Gibbs sampling to estimate the Hierarchical Bayes model. The idea of this procedure is to recursively draw each variable from the conditional distributions given other variables⁴. In each iteration, variables are drawn from the corresponding conditional distributions in the following order and are updated to form conditional distributions for

⁴More precisely, given all other variables in its Markov blanket




Figure C.1: A Hierarchical Graph of the Non-sequential Search Model

In the graph, latent variables (e.g., parameters need to be estimated) are represented by white nodes, and observed outcomes are represented by grey nodes. The directed edges represent the causal relationship between nodes. Subscripts are omitted.

other variables:

$$\beta_i \mid \mathbf{p}_i, Y_{S_i}, Y_{P_i}, z_i, \boldsymbol{\eta}, \sigma_\zeta, c, \Gamma, \Sigma_\beta, \forall i \quad (\text{C.11})$$

$$\Gamma \mid Z, \boldsymbol{\beta}, \Sigma_\beta \quad (\text{C.12})$$

$$\Sigma_\beta \mid Z, \boldsymbol{\beta}, \Gamma \quad (\text{C.13})$$

$$\boldsymbol{\eta} \mid \mathbf{p}, Y_S, Y_P, \boldsymbol{\beta}, \sigma_\zeta, c \quad (\text{C.14})$$

$$\sigma_\zeta \mid \mathbf{p}, Y_S, Y_P, \boldsymbol{\eta}, \boldsymbol{\beta}, c \quad (\text{C.15})$$

$$c \mid \mathbf{p}, Y_S, Y_P, \boldsymbol{\eta}, \boldsymbol{\beta}, \sigma_\zeta \quad (\text{C.16})$$

The posterior (conditional) distribution of Γ is a multivariate normal distribution, and the posterior distribution distribution of Σ_β is an Inverse-Wishart distribution. These two parameters can be drawn directly from the corresponding posterior distributions. However, the posterior distribution of β_i , $\boldsymbol{\eta}$, σ_ζ and c does not have closed forms. Thus a Markov Chain Monte Carlo (MCMC) approach is used to obtain the draws. The exact forms of these distributions and the details of generating the draws in each iteration are specified in Appendix C.5. We recursively generate the draws by looping through Equation C.11 to C.16,

until the distributions are stable. After convergence, we keep running the program for 1000 iterations to obtain statistics of interest such as mean and variance of the variables.

Identification and Results

Before presenting the estimation results, we explain how the parameters in our structural model can be identified.

As mentioned previously, we normalize the variance of the random utility component ε_{ij} in the utility function to 1. The remaining parameters in the consumer utility function, namely the product (seller) fixed effects (η_j) and consumer price sensitivity (β_i), can be identified from the consumers' purchase decisions conditioned on their consideration sets, which are observable in the available data. The systematic variation in purchase probabilities, conditioned on being considered, across different products help identify η_j . In our dataset, sellers on average change the price of their product 4 times during the one-month period, and the average price change is between 10% to 20%. With those price changes, β_i can be best identified if we can observe how each consumer's purchase behavior changes at different price points, as the product values remain constant over time. Although repeated purchases from the same consumers are infrequent in our data, we can leverage multiple consumers from the same demographic group and with the same characteristics.⁵ The average impact of price changes on the probability of each product being purchased conditioned on being searched enable us to identify the mean price sensitivity of that group (Γz_i). The extent of variation in this effect across consumers within a demographic group helps identify the variance of the part of price sensitivity that cannot be explained by demographic characteristics (ϵ_β). Additionally, how Γz_i vary by z_i provides identification of Γ .

Once η_j and the distribution of β_i are identified, the search cost per product (c) is identified from the average size of the consideration set: a smaller average size indicates a larger value for c . Furthermore, by examining the extent of the variation in the composition of

⁵(Note that in our data, demographic characteristics happen to be all categorical).

the consideration set among consumers within the same demographic group, we can identify the variance of the first-stage assessment error (σ_ζ).

Variable	Notation	Coeff.	SE
Price Coefficient	Γ		
Intercept		-0.074	0.001
Age			
≤ 25 years old		-0.008	0.001
36 - 45 years old		0.005	0.001
46 - 55 years old		0.007	0.001
≥ 56 years old		0.008	0.002
Is Male		0.003	0.001
Is Married		0.000	0.001
Education Level			
Less Than High School		0.008	0.002
High School or Equivalent		0.010	0.001
Post-graduate Degree		-0.001	0.001
Seller Fixed Effect	η		
Seller 1		4.004	0.088
Seller 2		2.970	0.070
Seller 3		2.180	0.079
Seller 4		2.910	0.084
Seller 5		2.326	0.085
Search Cost	c	1.208	0.011
Unobserved Heterogeneity	Σ_β	0.001	0.000
Stochastic term of choice set	σ_ζ	0.751	0.009

Table C.6: Estimation Results

The estimation results are shown in Table C.6. On average, 1st party sellers have higher product (seller) fixed effects than 3rd party sellers. Average search cost is estimated to be 18.5 CNY (2.7 USD), about 25% of product prices. As mentioned earlier, the scale of ε_{ij} in the utility function is normalized to 1, and the scale parameter of the choice set-specific stochastic term (σ_ζ) is estimated to be 0.724, which suggests the impact of optimization error on the expected net benefit of the choice sets is relatively small. Consumer price sensitivity tends to decrease with age⁶, while showing an increasing trend based on consumer education

⁶Note that in the utility function, we did not include a negative sign before the price term. As a result, the greater the negativity of β_i , the higher the price sensitivity of consumer i becomes.

level (with Bachelor's Degree as the baseline group). In general, male consumers exhibit lower levels of price sensitivity.

Supply Side Estimation

In order to run counterfactual simulations, we need supply-side parameters, such as the marginal cost of the products. Following the common practice of the supply-side estimation, we assume sellers maximize their profit by taking into account the demand-side information and supply-side cost information. Denote seller j 's marginal cost as mc_j , seller j ' pricing strategy as σ_j (possibly mixed), and seller j ' strategy space as $\Delta(\Phi_j)$, then a strategy profile $(\sigma_1, \dots, \sigma_J)$ constitutes a Nash equilibrium if it satisfies the following conditions:

$$\Pi_j(\sigma_j, \sigma_{-j}) \geq \Pi_j(\sigma'_j, \sigma_{-j}), \quad \forall j, \quad \forall \sigma'_j \in \Delta(\Phi_j) \quad (\text{C.17})$$

where $\Pi(\cdot)$ is the profit function given all sellers' strategies. Let $\Phi_j^+ \subset \Phi_j$ denote the set of prices that seller j has set at least once (pure strategies played with positive probabilities), then the profit function can be expressed as:

$$\Pi_j(\sigma_j, \sigma_{-j}) = (p_j - mc_j)D(p_j, \sigma_{-j}), \quad \forall p_j \in \Phi_j^+ \quad (\text{C.18})$$

The term $D(p_j, \sigma_{-j})$ above represents the demand for product j when the focal seller sets its price at p_j and the opponents set their prices according to strategy σ_{-j} , which can be derived from the previously estimated demand model. The conditions in Equation (C.17) require that $\forall j$:

$$(p_j - mc_j)D(p_j, \sigma_{-j}) = (p'_j - mc_j)D(p'_j, \sigma_{-j}), \quad \forall p_j, p'_j \in \Phi_j^+ \quad (\text{C.19})$$

$$(p_j - mc_j)D(p_j, \sigma_{-j}) \geq (p'_j - mc_j)D(p'_j, \sigma_{-j}), \quad \forall p_j \in \Phi_j^+, \quad \forall p'_j \notin \Phi_j^+ \quad (\text{C.20})$$

Since the exact formula of the demand function $D(\cdot)$ is complicated, we solve for the

equilibrium prices that satisfy the above conditions numerically when estimating the marginal costs of the products.

C.3 SKU Subset Selection Details

Unfortunately, we cannot see the detailed product information from the data set so we cannot manually select a subset of close substitutes. Instead, we use consumers' clicking data to infer the subsets which contain close substitutes: if two products are close substitutes, then consumers are very likely to click on both of them to compare. We say two products are 'co-clicked' by a consumer if this consumer clicks on both products within a relatively short period of time (for example, 5 hours). If two products are 'co-clicked' by many consumers, then these two products are very likely to be close substitutes. With this idea, we first create an un-directed weighted graph using the 2.5 million click events in the *Click Table* to describe the 'co-click' relationship among all SKUs. Specifically, each node in the graph represents an SKU, the weight of each edge in the graph represents the number of consumers who have co-clicked on these two SKUs. The adjacency matrix of this graph turns out to be very sparse which means that most SKU pairs are never 'co-clicked' by any consumer. Next, we would like to find a subset of SKUs within which all SKUs are close substitutes (i.e., the weight of all edges connecting these nodes are high). In graph theory, such activity is called 'community detection', which means detecting local communities in a graph that each member in it has a close relationship with others (The SKUs within a community are more densely connected than the SKUs across communities). Considering the size of our graph, we adopt the *Louvain method for community detection* (Blondel et al., 2008) to detect local communities which have algorithm complexity of $O(n \log n)$. One potential issue of this method is that if two products are complementary instead of substitutes, they will also be included in the same cluster. Later in the paper, we rule out this possibility by showing that very few consumers end up purchasing more than one product.

C.4 Theoretical Framework Details

We study the competition of algorithms in oligopoly pricing, which is set within a repeated game framework. The market consists of n differentiated products, each sold by one of n sellers. Without any capacity constraints, each seller offers their own product in every period. Sellers concurrently determine the prices of their respective products in each period, and previous price history is common knowledge. We assume sellers use a size m discrete price grid and have one-period memory, and we restrict our attention to the case $n = 2$. Specifically, the action space for the Q-learning algorithm is $A = \{p_i\}, i \in \{0, \dots, m - 1\}$. The state contains the last period prices of all sellers $S = \{s_{ij} = (p_i, p_j)\}, i \in \{0, \dots, m - 1\}, j \in \{0, \dots, m - 1\}$.

In essence, the procedure of learning of the algorithm is the procedure of updating its Q-table. The size of the Q-table is $\|A\| \times \|S\| = m \times m^2$, that is, for each state, there are m Q-values representing the value of taking each of the m actions. We start with the Q-Q scenario. Denote p_k^t as the price charged by seller k at time t , respectively. Denote s^t as the state in time t , where $s^t = (p_1^{t-1}, p_2^{t-1})$. Seller k receives a reward $r_k(p_1^t, p_2^t)$, and the state transits to $s^{t+1} = (p_1^t, p_2^t)$ in the next time period. We focus on the case where the two sellers are symmetric and we derive everything from seller 1's perspective. We denote the seller 1's Q-value at time t that corresponds to taking action $p_{i'}$ in state (p_i, p_j) as $Q_{(p_i, p_j) - p_{i'}}^1(t)$. According to the Q-learning updating rule, at time t , one of the Q-values in the Q-table, $Q_{(p_1^{t-1}, p_2^{t-1}) - p_1^t}^1(t)$ will be updated as follows:

$$Q_{(p_1^{t-1}, p_2^{t-1}) - p_1^t}^1(t+1) = (1 - \alpha)Q_{(p_1^{t-1}, p_2^{t-1}) - p_1^t}^1(t) + \alpha(r_1(p_1^t, p_2^t) + \delta \max_{p_{i'}}(Q_{(p_1^t, p_2^t) - p_{i'}}^1)) \quad (\text{C.21})$$

Where δ is the discount factor and α is the learning rate. From the above, we can see that the reward and state transition depend on both sellers' actions, which are not always deterministic (e.g., random exploration is needed in the early stage of learning). We are interested in how the Q-value evolves *in expectation*. We now re-formalize Equation C.21 using the two sellers' *probabilities of taking each action* instead of the actions that they *actually*

take. Assume that given state $s = (p_i, p_j)$, seller k takes action $a_k = p_{i'}$ with probability $x_{(p_i, p_j) - p_{i'}}^k, \forall p_{i'} \in A$. Denote the probability that the state (p_i, p_j) appears as $P_{(p_i, p_j)}$, then $Q_{(p_i, p_j) - p_{i'}}^k$ will be updated *in expectation* in the following way:

$$Q_{(p_i, p_j) - p_{i'}}^1(t+1) = (1 - \alpha)Q_{(p_i, p_j) - p_{i'}}^1(t) + \alpha P_{(p_i, p_j)} x_{(p_i, p_j) - p_{i'}}^1 \sum_{j'} (x_{(p_i, p_j) - p_{j'}}^2 r(p_{i'}, p_{j'})) + \delta \sum_{j'} (x_{(p_i, p_j) - p_{j'}}^2 \max_{p_{i''}} (Q_{(p_{i'}, p_{j'}) - p_{i''}}^1)) \quad (\text{C.22})$$

The basic logic for Equation C.22 is: Q-values are state and action specific, and we want to weight the learning rate of a Q-value that corresponds to a certain state and a certain action with the probability that this specific state occurs and the specific action is taken. In essence, for the state that is frequently visited and the action that is frequently taken, the corresponding Q-value will be updated faster than those less frequently visited states and actions.

Now we transform the discrete-time learning problem into a continuous time version by compressing the time interval between two consecutive time periods from 1 to Δt . As Δt approaches 0, the Q-learning dynamics can be captured by the following system of differential equations:

$$\left\{ \begin{array}{l} \frac{dQ_{(p_i, p_j) - p_{i'}}^1(t)}{dt} = \alpha P_{(p_i, p_j)} x_{(p_i, p_j) - p_{i'}}^1 (\sum_{j'} (x_{(p_i, p_j) - p_{j'}}^2 r_1(p_{i'}, p_{j'})) + \delta \sum_{j'} (x_{(p_i, p_j) - p_{j'}}^2 \max_{i''} (Q_{(p_{i'}, p_{j'}) - p_{i''}}^1)) - Q_{(p_i, p_j) - p_{i'}}^1(t)) \quad \forall i, j, i' \in \{1, \dots, m\} \end{array} \right. \quad (\text{C.23})$$

where $x_{(p_i, p_j) - p_{i'}}^k, \forall k \in \{1, 2\}$ are determined by the current Q-values and the exploration strategy. In particular, we use the ϵ greedy exploration strategy here, where the probability of taking a greedy action (the action with the largest Q-value) given a state is $1 - \epsilon$ and the probability for taking any other actions is ϵ . Then we have:

$$x_{(p_i, p_j) - p_{i'}}^k = \begin{cases} \epsilon & \text{if } p_{i'} = \arg \max_p Q_{(p_i, p_j) - p}^k \quad \forall i, j, i' \in \{1, \dots, m\}, \forall k \in \{1, 2\} \\ \frac{\epsilon}{(m-1)} & \text{otherwise} \end{cases} \quad (\text{C.24})$$

$P_{(p_i, p_j)}$ is the frequency that state (p_i, p_j) is visited under sellers' current strategies, and can be determined by the following system of linear equations

$$\begin{cases} P_{(p_i, p_j)} = \sum_{i', j'} P_{(p_{i'}, p_{j'})} x_{(p_{i'}, p_{j'}) - p_i}^1 x_{(p_{i'}, p_{j'}) - p_j}^2 & \forall i, j \in \{1, \dots, m\} \\ \sum_{i', j'} P_{(p_{i'}, p_{j'})} = 1 \end{cases} \quad (\text{C.25})$$

The complete Q-dynamics can be found by solving Equations C.23 to C.25.

C.4.1 Learning Dynamics of The Q-Q scenario (Proof of Lemma 9)

Learning Stage 0: Initialization and Random Exploration

we initialize the Q-values using the discounted payoff that would arise for the focal seller if the opponent seller chooses the price randomly, which is also the Q-values that the Q-learning algorithm will learn when both sellers fully explore. For notation simplicity, we use $Q_{ij-k}(t)$ to denote seller 1's Q-value of taking action p_k in state (p_i, p_j) , $\forall i, j, k \in \{l, h\}$.

$$\begin{aligned} Q_{ul-l}(0) = Q_{hh-l}(0) = Q_{lh-l}(0) = Q_{hl-l}(0) &= \frac{r_u + r_{lh}}{2(1 - \delta)} \\ Q_{ul-h}(0) = Q_{hh-h}(0) = Q_{lh-h}(0) = Q_{hl-h}(0) &= \frac{\delta(r_u + r_{lh})}{2(1 - \delta)} + \frac{r_{hl} + r_{hh}}{2} \end{aligned} \quad (\text{C.26})$$

Since our initialization is consistent with the Q-values that the Q-learning algorithm will actually learn, the Q-values don't change in this stage. At the end of this stage, both sellers' greedy strategies are to play p_l at all of the states.

Learning Stage 1: The Drop of Q_{u-l}

After the exploration is shut down at $t = t_0$, both sellers start to take their greedy actions: play p_l repeatedly. In this stage, only the stage (p_l, p_l) is visited so only Q_{u-l} is updated. According to Equation C.21, the Q-dynamic for Q_{u-l} in this stage is:

$$\frac{dQ_{u-l}(t)}{dt} = \alpha(r_u + \delta Q_{u-l}(t) - Q_{u-l}(t)) \quad (\text{C.27})$$

With initial condition:

$$Q_{u-l}(t_0) = \frac{r_u + r_{lh}}{2(1 - \delta)} \quad (\text{C.28})$$

Therefore,

$$Q_{u-l}(t) = \frac{r_{lh} - r_u}{2(1 - \delta)} \exp(-\alpha(t - t_0)(1 - \delta)) + \frac{r_u}{1 - \delta} \quad (\text{C.29})$$

It's straightforward to check that $Q_{u-l}(t)$ decreases in t , the intuition is that: the Q-values learned in stage 0 is for the case when the opponent seller randomly pick between p_l and p_h , however, in stage 1, the opponent seller becomes more aggressive in setting the price (it charges p_l repeatedly). Thus, the Q values learned in stage 0 overestimate the true Q-values that appear in stage 1. As new experiences are gathered, the Q-values are updated downwards.

A sufficient condition that the two sellers are trapped in the (p_l, p_l) equilibrium is:

$$\lim_{t \rightarrow \infty} Q_{u-l}(t) > Q_{u-h}(t_0) \quad (\text{C.30})$$

The intuition is that if the Q-value for taking p_l in state (p_l, p_l) never drops below the Q-value for taking p_h in state (p_l, p_l) , the two sellers will never reach to other states in the first place, thus they will end up in state (p_l, p_l) forever. Take Equations C.39 and C.26 in Equation C.30, we get the sufficient condition for equilibrium outcome (p_l, p_l) :

$$\delta < \frac{b}{2-b} \quad (\text{C.31})$$

Learning Stage 2: The Drop of Q_{hh-l}

If Condition C.31 is not met, Q_{ll-l} will drop below Q_{ll-h} at t_1 , the two sellers begins to charge p_h in state (p_l, p_l) but charge p_l in state (p_h, p_h) . Thus the state will jump between (p_l, p_l) and (p_h, p_h) , the two Q-values, Q_{ll-h} and Q_{hh-l} will be updated at the same time. The Q-dynamics for these two Q-values are:

$$\begin{aligned} \frac{dQ_{ll-h}(t)}{dt} &= \frac{\alpha}{2}(r_{hh} + \delta Q_{hh-l}(t) - Q_{ll-h}(t)) \\ \frac{dQ_{hh-l}(t)}{dt} &= \frac{\alpha}{2}(r_{ll} + \delta Q_{ll-h}(t) - Q_{hh-l}(t)) \end{aligned} \quad (\text{C.32})$$

With initial conditions:

$$\begin{aligned} Q_{ll-h}(t_1) &= \frac{\delta(r_{ll} + r_{lh})}{2(1-\delta)} + \frac{r_{hl} + r_{hh}}{2} \\ Q_{hh-l}(t_1) &= \frac{r_{ll} + r_{lh}}{2(1-\delta)} \end{aligned} \quad (\text{C.33})$$

The solutions turn out to be:

$$\begin{aligned}
Q_{u-h}(t) &= (((1 + \delta)(1 - \delta)(r_u + r_{lh} - r_{hh} - r_{hl}) + 2(1 - \delta)(r_{hh} - r_u)) \exp(\frac{-(t - t_1)(1 + \delta)}{2}) \\
&\quad + (1 + \delta)((r_u + r_{lh} - r_{hh} - r_{hl})(1 - \delta) + 2(r_{hh} - r_{lh})) \exp(\frac{-(t - t_1)(1 - \delta)}{2}) \\
&\quad - 4(r_{hh} + r_u \delta)) / (4(\delta - 1)(\delta + 1)) \\
Q_{hh-l}(t) &= - ((1 + \delta)(1 - \delta)(r_u + r_{lh} - r_{hh} - r_{hl}) + 2(1 - \delta)(r_{hh} - r_u)) \exp(\frac{-(t - t_1)(1 + \delta)}{2}) \\
&\quad + (1 + \delta)((r_u + r_{lh} - r_{hh} - r_{hl})(1 - \delta) + 2(r_{hh} - r_{lh})) \exp(\frac{-(t - t_1)(1 - \delta)}{2}) \\
&\quad - 4(r_u + r_{hh} \delta) / (4(\delta - 1)(\delta + 1))
\end{aligned} \tag{C.34}$$

A sufficient condition that the two sellers end up with alternating between (p_l, p_l) and (p_h, p_h) in equilibrium is:

$$\begin{aligned}
\lim_{t \rightarrow \infty} Q_{u-h}(t) &> Q_{u-l}(t_1) \\
\lim_{t \rightarrow \infty} Q_{hh-l}(t) &> Q_{hh-h}(t_1)
\end{aligned} \tag{C.35}$$

Take Equations C.34 and C.26 in Equation C.35, we get the sufficient condition for equilibrium outcome $(p_l, p_l) - (p_h, p_h)$:

$$\frac{b}{2 - b} \leq \delta < \sqrt{\frac{b}{2 - b}} \tag{C.36}$$

Learning Stage 3: Rise of Q_{hh-h}

If neither Condition C.31 nor C.36 is met, Q_{hh-l} will drop below Q_{hh-h} at t_2 , the two sellers switch to charge p_h in state (p_h, p_h) . Then only Q_{hh-h} will be updated afterward. The Q-dynamics can be shown:

$$\frac{dQ_{hh-h}(t)}{dt} = \alpha(r_{hh} + \delta Q_{hh-h}(t) - Q_{hh-h}(t)) \quad (\text{C.37})$$

With initial condition:

$$Q_{hh-h}(t_2) = \frac{\delta(r_{ul} + r_{lh})}{2(1-\delta)} + \frac{r_{hl} + r_{hh}}{2} \quad (\text{C.38})$$

Therefore the solution:

$$Q_{hh-h}(t) = \frac{r_{hh}}{1-\delta} - \frac{r_{hh} - r_{hl} + \delta(r_{hh} - r_{ul} + r_{hl} - r_{lh})}{2(1-\delta)} \exp(-\alpha(t-t_2)(1-\delta)) \quad (\text{C.39})$$

It's straightforward to check that $\lim_{t \rightarrow \infty} Q_{hh-h}(t) = \frac{r_{hh}}{1-\delta} > Q_{hh-h}(t_2)$, Therefore, a sufficient condition for the (p_h, p_h) equilibrium is:

$$\delta \geq \sqrt{\frac{b}{2-b}} \quad (\text{C.40})$$

C.4.2 Learning Dynamics of The Q-R scenario (Proof of Lemma 10)

The Q-learning dynamic can be found by solving Equations C.41 to C.43.

$$\left\{ \begin{array}{l} \frac{dQ_{(p_i, p_j) - p_{i'}}^R(t)}{dt} = \alpha P_{(p_i, p_j)} x_{(p_i, p_j) - p_{i'}}^R (r(p_{i'}, p_{\max(i-1, 1)}) \\ + \delta \max_{j''} (Q_{(p_{i'}, p_{\max(i-1, 1)}) - p_{j''}}^R)) \quad \forall i, j, i' \in \{1, \dots, m\} \\ - Q_{(p_i, p_j) - p_{i'}}^R(t) \end{array} \right. \quad (\text{C.41})$$

$$x_{(p_i, p_j) - p_{i'}}^R = \begin{cases} \epsilon & \text{if } p_{i'} = \arg \max_p Q_{(p_i, p_j) - p}^R \\ \frac{\epsilon}{(m-1)} & \text{otherwise} \end{cases} \quad \forall i, j, i' \in \{1, \dots, m\} \quad (\text{C.42})$$

$$\left\{ \begin{array}{l} P_{(p_i, p_j)} = \sum_{i'j'} P_{(p_{i'}, p_{j'})} x_{(p_{i'}, p_{j'}) - p_i}^R \mathbb{1}\{p_j = p_{\max(i'-1, 1)}\} \quad \forall i, j \in \{1, \dots, m\} \\ \sum_{i'j'} P_{(p_{i'}, p_{j'})} = 1 \end{array} \right. \quad (\text{C.43})$$

However, finding the learning outcome in the Q-R scenario is much simpler. Since seller 2's pricing strategy does not alter, the environment is stationary for seller 1, and the Q-learning algorithm is guaranteed to find the optimal pricing policy. The eventual Q-values can be found by solving the following system of equations:

$$\begin{aligned} Q_{u-l} &= r_u + \delta \max(Q_{u-l} + Q_{u-h}) \\ Q_{u-h} &= r_{hl} + \delta \max(Q_{hh-l} + Q_{hh-h}) \\ Q_{hh-l} &= r_{lh} + \delta \max(Q_{u-l} + Q_{u-h}) \\ Q_{hh-h} &= r_{hh} + \delta \max(Q_{hh-l} + Q_{hh-h}) \end{aligned} \quad (\text{C.44})$$

Depending on the sign of $Q_{u-l} - Q_{u-h}$ and $Q_{hh-l} + Q_{hh-h}$, we get 3 types of outcomes: When $\delta < 0.5$, $Q_{u-l} > Q_{u-h}, Q_{hh-l} > Q_{hh-h}$ and (p_l, p_l) will be the equilibrium outcome. When $0.5 \leq \delta < \frac{1}{2(1-b)}$, $Q_{u-l} > Q_{u-h}, Q_{hh-l} \leq Q_{hh-h}$ and both (p_h, p_h) and (p_l, p_l) will be the equilibrium outcome, according to our equilibrium selection criteria, (p_h, p_h) is selected since it Pareto dominates (p_l, p_l) . When $\delta \geq \frac{1}{2(1-b)}$, $Q_{u-l} \leq Q_{u-h}, Q_{hh-l} < Q_{hh-h}$ and (p_h, p_h) will be the equilibrium outcome.

C.4.3 Proof of Theorem 3

The proof is straightforward by comparing Lemma 9 and Lemma 10 thus is omitted here.

C.5 A Hierarchical Bayes Estimation Algorithm

Following (Netzer et al., 2008), we specify our Hierarchical Bayes estimation algorithm details in this section. All vectors and matrices are bold. Subscript i denotes consumer i , and N is the total number of consumers.

In each iteration, variables are drawn from corresponding conditional distributions in the following order and are updated to form subsequent conditional distributions for other variables:

$$\beta_i \mid p, Y_{S_i}, Y_{P_i}, z_i, \eta, \sigma_\zeta, c, \Gamma, \Sigma_\beta$$

$$\Gamma \mid Z, \beta, \Sigma_\beta$$

$$\Sigma_\beta \mid Z, \beta, \Gamma$$

$$\eta \mid p, Y_S, Y_P, \beta, \sigma_\zeta, c$$

$$\sigma_\zeta \mid p, Y_S, Y_P, \eta, \beta, c$$

$$c \mid p, Y_S, Y_P, \eta, \beta, \sigma_\zeta$$

(1) Generate $\beta_i, \forall i$

$$\begin{aligned} & f(\beta_i \mid p_i, Y_{S_i}, Y_{P_i}, z_i, \eta, \sigma_c, c, \Gamma, \Sigma_\beta) \\ & \propto \mathcal{N}(\beta_i \mid \eta, \sigma_c, c, \Gamma, \Sigma_\beta) L(Y_{S_i}, Y_{P_i}) \\ & \propto |\Sigma_\beta|^{-1/2} \exp(-1/2(\beta_i - \Gamma' z_i)' \Sigma_\beta^{-1} (\beta_i - \Gamma' z_i)) L(Y_{S_i}, Y_{P_i}) \end{aligned} \tag{C.45}$$

Where $L(Y_{S_i}, Y_{P_i})$ is the likelihood function from Equation C.10. We use the Metropolis-Hasting algorithm to draw new β_i from its conditional distribution. In each iteration, a new vector value β_i^{new} is drawn from a multivariate normal distribution $\mathcal{N}(\beta_i, V_D^2)$ where V_D^2 is chosen adaptively, and we accept the generated new draw with probability

$$Pr = \min\left\{\frac{\exp[-1/2(\beta_i^{new} - \Gamma' z_i)' \Sigma_\beta^{-1}(\beta_i^{new} - \Gamma' z_i)]L(Y_{S_i}, Y_{P_i}|\beta_i^{new})}{\exp[-1/2(\beta_i - \Gamma' z_i)' \Sigma_\beta^{-1}(\beta_i - \Gamma' z_i)]L(Y_{S_i}, Y_{P_i}|\beta_i)}, 1\right\} \quad (\text{C.46})$$

(2) Generate Γ

The vectorized Γ (denoted as $v\Gamma$) follows a multivariate normal distribution

$$v\Gamma \mid Z, \beta, \Sigma_\beta = \mathcal{N}(u_n, V_n) \quad (\text{C.47})$$

Where

$$V_n = [(Z'Z \otimes \Sigma_\beta^{-1}) + V_0^{-1}]^{-1},$$

$$u_n = V_n[Z' \otimes \Sigma_\beta^{-1}\beta^* + V_0^{-1}u_0],$$

$$Z = (z'_1, z'_2, \dots, z'_N) \text{ is an } N \times nz \text{ matrix,}$$

$$\beta^o = (\beta'_1, \beta'_2, \dots, \beta'_N) \text{ is an } N \times n\beta \text{ matrix,}$$

$$\beta^* = \text{vec}(\beta^o),$$

$$nz = \dim(z_i),$$

$$n\beta = \dim(\beta_i),$$

We set prior hyperparameters $u_0 = \{0\}_{n\beta \cdot nz}$, and $V_0 = 100I_{n\beta \cdot nz}$.

(3) Generate Σ_β

$$\Sigma_\beta \mid Z, \beta, \Gamma \sim \mathcal{W}_{n\beta}^{-1}(f_0 + N, G_0^{-1} + \Sigma_{i=1}^N (\beta_i - \Gamma' Z_i)' (\beta_i - \Gamma' Z_i)) \quad (\text{C.48})$$

Where \mathcal{W}^{-1} is the inverse Wishart distribution. $f_0 = n\beta + 5$ and $G_0 = I_{n\beta}$ are prior hyperparameters.

(4) Generate η

The procedure of generating η , σ_ζ , and c is similar to the procedure of generating β_i , where the Metropolis-Hasting algorithm is used to generate draws from the specified probability distribution. We use diffused priors instead of having prior distribution specified by other

higher-level variables.

$$\begin{aligned}
& f(\eta \mid p, Y_S, Y_P, \beta, \sigma_\zeta, c) \\
& \propto \mathcal{N}(\eta_0, V_{\eta_0})L(Y_S, Y_P) \\
& \propto |V_{\eta_0}|^{-1/2} \exp(-1/2(\eta - \eta_0)' V_{\eta_0}^{-1}(\eta - \eta_0))L(Y_S, Y_P)
\end{aligned} \tag{C.49}$$

In each iteration, a new draw η^{new} is generated from the multivariate normal distribution $\mathcal{N}(\eta, V_\eta)$ and will be accepted with the following acceptance probability

$$Pr = \min\left\{\frac{\exp(-1/2(\eta^{new} - \eta_0)' V_{\eta_0}^{-1}(\eta^{new} - \eta_0))L(Y_S, Y_P|\eta^{new})}{\exp(-1/2(\eta - \eta_0)' V_{\eta_0}^{-1}(\eta - \eta_0))L(Y_S, Y_P|\eta)}, 1\right\} \tag{C.50}$$

The diffusion prior hyperparameters are set as follows: $\eta_0 = \{0\}_{ns}$, $V_{\eta_0} = 100I_{ns}$.

(5) Generate σ_ζ

$$\begin{aligned}
& f(\sigma_\zeta \mid p, Y_S, Y_P, \eta, \beta, c) \\
& \propto \mathcal{N}(\sigma_{\zeta_0}, V_{\sigma_{\zeta_0}})L(Y_S, Y_P) \\
& \propto |V_{\sigma_{\zeta_0}}|^{-1/2} \exp(-1/2(\sigma_\zeta - \sigma_{\zeta_0})' V_{\sigma_{\zeta_0}}^{-1}(\sigma_\zeta - \sigma_{\zeta_0}))L(Y_S, Y_P)
\end{aligned} \tag{C.51}$$

In each iteration, a new draw σ_ζ^{new} is generated from a normal distribution $\mathcal{N}(\sigma_\zeta, V_{\sigma_\zeta})$ and will be accepted with the following acceptance probability

$$Pr = \min\left\{\frac{\exp(-1/2(\sigma_\zeta^{new} - \sigma_{\zeta_0})' V_{\sigma_{\zeta_0}}^{-1}(\sigma_\zeta^{new} - \sigma_{\zeta_0}))L(Y_S, Y_P|\sigma_\zeta^{new})}{\exp(-1/2(\sigma_\zeta - \sigma_{\zeta_0})' V_{\sigma_{\zeta_0}}^{-1}(\sigma_\zeta - \sigma_{\zeta_0}))L(Y_S, Y_P|\sigma_\zeta)}, 1\right\} \tag{C.52}$$

The diffusion prior hyperparameters are set as follows: $\sigma_{\zeta_0} = 0$, $V_{\sigma_{\zeta_0}} = 30$.

(5) Generate c

$$\begin{aligned}
& f(c \mid p, Y_S, Y_P, \eta, \beta, \sigma_\zeta) \\
& \propto \mathcal{N}(c_0, V_{c_0})L(Y_S, Y_P) \\
& \propto |V_{c_0}|^{-1/2} \exp(-1/2(c - c_0)' V_{c_0}^{-1} (c - c_0)) L(Y_S, Y_P)
\end{aligned} \tag{C.53}$$

In each iteration, a new draw c^{new} is generated from a normal distribution $\mathcal{N}(c, V_c)$ and will be accepted with the following acceptance probability

$$Pr = \min\left\{ \frac{\exp(-1/2(c^{new} - c_0)' V_{c_0}^{-1} (c^{new} - c_0)) L(Y_S, Y_P | c^{new})}{\exp(-1/2(c - c_0)' V_{c_0}^{-1} (c - c_0)) L(Y_S, Y_P | c)}, 1 \right\} \tag{C.54}$$

The diffusion prior hyperparameters are set as follows: $c_0 = 0$, $V_{c_0} = 30$.

Bibliography

- Hilke Schellmann and Jason Bellini. Artificial intelligence: The robots are now hiring. *The Wall Street Journal*, 20, 2018. Cited on page 1.
- Runshan Fu, Yan Huang, and Param Vir Singh. Crowds, lending, machine, and bias. *Information Systems Research*, 32(1):72–92, 2021a. Cited on page 1.
- Charles AE Goodhart. Problems of monetary management: the uk experience. In *Monetary Theory and Practice*, pages 91–121. Springer, 1984. Cited on page 2.
- Robert E Lucas et al. Econometric policy evaluation: A critique. In *Carnegie-Rochester conference series on public policy*, volume 1, pages 19–46, 1976. Cited on page 2.
- Danielle Keats Citron and Frank A. Pasquale. The scored society: Due process for automated predictions. *Washington Law Review*, 89, 2014a. Cited on pages 2 and 49.
- David Segal. The dirty little secrets of search. *The New York Times*, 12(02), 2011. Cited on page 2.
- Jon Kleinberg and Manish Raghavan. How do classifiers induce agents to invest effort strategically? In *Proceedings of the 2019 ACM Conference on Economics and Computation*, pages 825–844, 2019. Cited on pages 3, 11, 13, and 14.
- Tal Alon, Magdalen R. C. Dobson, Ariel D. Procaccia, Inbal Talgam-Cohen, and Jamie Tucker-Foltz. Multiagent evaluation mechanisms. In *AAAI*, 2020. Cited on pages 3, 13, and 14.
- Nika Haghtalab, Nicole Immorlica, Brendan Lucier, and Jack Wang. Maximizing welfare

- with incentive-aware evaluation mechanisms. In *29th International Joint Conference on Artificial Intelligence*, 2020. Cited on pages 3 and 13.
- Michael Spence. Job market signaling. *Quarterly Journal of Economics*, 87:355–374, 1973. Cited on pages 4, 5, 10, and 11.
- Carlos Alós-Ferrer and Julien Prat. Job market signaling and employer learning. *Journal of Economic Theory*, 147(5):1787 – 1817, 2012. Cited on pages 5 and 12.
- Brendan Daley and Brett Green. Market signaling with grades. *Journal of Economic Theory*, 151:114 – 145, 2014. Cited on pages 5, 10, and 12.
- Maxim Engers. Signalling with many signals. *Econometrica*, 55(3):663–674, 1987. Cited on pages 10 and 12.
- Andrew Weiss. A sorting-cum-learning model of education. *Journal of Political Economy*, 91(3):420–442, 1983. Cited on pages 10, 12, and 18.
- Alex Frankel and Navin Kartik. Improving information from manipulable data. *arXiv preprint arXiv:1908.10330*, 2019. Cited on pages 11, 12, 14, and 45.
- Alessandro Bonatti and Gonzalo Cisternas. Consumer Scores and Price Discrimination. *The Review of Economic Studies*, 87(2):750–791, 2019. Cited on pages 11 and 14.
- Vincent P Crawford and Joel Sobel. Strategic information transmission. *Econometrica*, pages 1431–1451, 1982. Cited on page 12.
- Moritz Hardt, Nimrod Megiddo, Christos Papadimitriou, and Mary Wootters. Strategic classification. In *Proceedings of the 2016 ACM conference on innovations in theoretical computer science*, pages 111–122, 2016. Cited on page 13.
- Reshef Meir, Ariel Procaccia, and Jeffrey Rosenschein. Algorithms for strategyproof classification. *Journal of Artificial Intelligence*, 186:123–156, 07 2012. Cited on page 13.
- Rachel Cummings, Stratis Ioannidis, and Katrina Ligett. Truthful linear regression. In *Conference on Learning Theory*, pages 448–483, 2015. Cited on page 13.

- Smitha Milli, John Miller, Anca D Dragan, and Moritz Hardt. The social cost of strategic classification. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*, pages 230–239, 2019. Cited on page 45.
- Bryce Goodman and Seth Flaxman. European union regulations on algorithmic decision-making and a “right to explanation”. *AI magazine*, 38(3):50–57, 2017. Cited on page 49.
- Dimitris Bertsimas, Arthur Delarue, Patrick Jaillet, and Sebastien Martin. The price of interpretability. *arXiv preprint arXiv:1907.03419*, 2019. Cited on page 50.
- Danielle Keats Citron and Frank Pasquale. The scored society: Due process for automated predictions. *Wash. L. Rev.*, 89:1, 2014b. Cited on pages 51 and 87.
- Sherrill Shaffer. The winner’s curse in banking. *Journal of Financial Intermediation*, 7(4): 359–392, 1998. ISSN 1042-9573. Cited on page 51.
- Don Carmichael. The winner’s curse in an online lending market. *SSRN 2931022*, 2017. Cited on page 51.
- Experian. Survey findings: How do consumers feel about credit cards? *Experian*, 2020. Cited on page 51.
- Sanford J. Grossman. The informational role of warranties and private disclosure about product quality. *The Journal of Law and Economics*, 24:461–483, 1981. Cited on page 55.
- Paul R. Milgrom. Good news and bad news: Representation theorems and applications. *The Bell Journal of Economics*, 12(2):380–391, 1981. Cited on page 55.
- Dmitri Kuksov and Yuanfang Lin. Information provision in a vertically differentiated competitive marketplace. *Marketing Science*, 29(1):122–138, 2010. Cited on page 55.
- Oliver Board. Competition and disclosure. *The Journal of Industrial Economics*, 57(1): 197–213, 2009. Cited on page 55.
- Dan Levin, James Peck, and Lixin Ye. Quality disclosure and competition. *Journal of*

- Industrial Economics*, 57(1):167–196, 2009. Cited on page 55.
- Zheyin (Jane) Gu and Ying Xie. Facilitating fit revelation in the competitive market. *Management Science*, 59(5):1196–1212, 2013. Cited on pages 55 and 56.
- V. Joseph Hotz and Mo Xiao. Strategic information disclosure: The case of multi-attribute products with heterogeneous consumers. *Economic Inquiry*, 51(1):865–881, January 2013. Cited on page 55.
- Liang Guo and Ying Zhao. Voluntary quality disclosure and market interaction. *Marketing Science*, 28(3):488–501, 2009. Cited on pages 55 and 56.
- Bikram Ghosh and Michael R. Galbreth. The impact of consumer attentiveness and search costs on firm quality disclosure: A competitive analysis. *Management Science*, 59(11):2604–2621, 2013. doi: 10.1287/mnsc.2013.1724. Cited on page 56.
- Jianqiang Zhang and Krista J. Li. Quality disclosure under consumer loss aversion. *Management Science*, 67(8):5052–5069, 2021. doi: 10.1287/mnsc.2020.3745. Cited on page 56.
- Robert Hauswald and Robert Marquez. Information technology and financial services competition. *Review of Financial Studies*, 16:921–948, 02 2003. Cited on pages 56 and 63.
- Robert Hauswald and Robert Marquez. Competition and Strategic Information Acquisition in Credit Markets. *Review of Financial Studies*, 19(3):967–1000, 2006. Cited on pages 56 and 63.
- Jan Bouckaert and Hans Degryse. Softening competition by inducing switching in credit markets. *The Journal of Industrial Economics*, 52(1):27–52, 2004. Cited on page 56.
- J. Miguel Villas-Boas and Udo Schmidt-Mohr. Oligopoly with asymmetric information: Differentiation in credit markets. *The RAND Journal of Economics*, 30(3):375–396, 1999. Cited on page 56.
- Jean Tirole and Tirole Jean. *The Theory of Industrial Organization*. MIT press, 1988. Cited

on page 57.

John R. Hauser. Competitive price and positioning strategies. *Marketing Science*, 7(1):76–91, 1988. Cited on page 57.

K. Sridhar Moorthy. Product and price competition in a duopoly. *Marketing Science*, 7(2):141–168, 1988. Cited on page 57.

Hal R Varian. A Model of Sales. *American Economic Review*, 70(4):651–659, 1980. Cited on page 68.

Runshan Fu, Yan Huang, and Param Vir Singh. Crowd, lending, machine, and bias. *Information Systems Research*, 31(1):72–92, 2021b. Cited on page 87.

Yanhao Wei, Pinar Yildirim, Christophe Van den Bulte, and Chrysanthos Dellarocas. Credit scoring with social network data. *Marketing Science*, 35(2):234–258, 2016. Cited on page 87.

Tat Chan, Naser Hamdi, Xiang Hui, and Zhenling Jiang. The value of verified employment data for consumer lending: Evidence from equifax. *Marketing Science*, 41(4):795–814, 2022. Cited on page 87.

Oded Netzer, Alain Lemaire, and Michal Herzenstein. When words sweat: Identifying signals for loan default in the text of loan applications. *Journal of Marketing Research*, 56(6):960–980, 2019. Cited on page 87.

Runshan Fu, Yan Huang, and Param Vir Singh. Artificial intelligence and algorithmic bias: Source, detection, mitigation, and implications. In *Tutorials in OR*, pages 39–63. INFORMS, 2020. Cited on page 87.

Srinivas Tunuguntla. Display ad measurement using observational data: A reinforcement learning approach. 2021. Cited on page 92.

Xiao Liu. Dynamic coupon targeting using batch deep reinforcement learning: An application to livestream shopping. *Marketing Science*, 2022. Cited on page 92.

- Wen Wang, Beibei Li, and Xueming Luo. Deep reinforcement learning for sequential targeting. *Available at SSRN 3487145*, 2021. Cited on page 92.
- Omid Rafieian. Optimizing user engagement through adaptive ad sequencing. *Marketing Science*, 2022. Cited on page 92.
- Nicolás Aramayo, Mario Schiappacasse, and Marcel Goic. A multiarmed bandit approach for house ads recommendations. *Marketing Science*, 2022. Cited on page 92.
- Jiaxi Liu, Yidong Zhang, Xiaoqing Wang, Yuming Deng, and Xingyu Wu. Dynamic pricing on e-commerce platform with deep reinforcement learning: A field experiment, 2019. Cited on page 92.
- Emilio Calvano, Giacomo Calzolari, Vincenzo Denicolo, and Sergio Pastorello. Artificial intelligence, algorithmic pricing, and collusion. *American Economic Review*, 110(10): 3267–97, 2020. Cited on pages 93, 94, 96, 97, 100, 103, and 105.
- Karsten T. Hansen, Kanishka Misra, and Mallesh M. Pai. Frontiers: Algorithmic collusion: Supra-competitive prices via independent algorithms. *Marketing Science*, 40(1):1–12, 2021. doi: 10.1287/mksc.2020.1276. Cited on pages 93, 96, 97, and 98.
- Timo Klein. Autonomous algorithmic collusion: Q-learning under sequential pricing. *The RAND Journal of Economics*, 52(3):538–558, 2021. Cited on pages 93 and 94.
- Justin Johnson, Andrew Rhodes, and Matthijs R Wildenbeest. Platform design when sellers use pricing algorithms. *Available at SSRN 3753903*, 2020. Cited on pages 93, 97, and 100.
- John Asker, Chaim Fershtman, Ariel Pakes, et al. Artificial intelligence, algorithm design and pricing. In *AEA Papers and Proceedings*, volume 112, pages 452–56. American Economic Association, 2022. Cited on pages 93, 94, 97, and 100.
- Le Chen, Alan Mislove, and Christo Wilson. An empirical analysis of algorithmic pricing on amazon marketplace. In *Proceedings of the 25th international conference on World Wide Web*, pages 1339–1349, 2016. Cited on pages 93 and 97.

- Jeanine Miklós-Thal and Catherine Tucker. Collusion by algorithm: Does better demand prediction facilitate coordination between sellers? *Management Science*, 65(4):1552–1561, 2019. Cited on page 96.
- Jason O’Connor and Nathan E. Wilson. Reduced demand uncertainty and the sustainability of collusion: How ai could affect competition. *Information Economics and Policy*, 54: 100882, 2021. Cited on page 96.
- Stephanie Assad, Robert Clark, Daniel Ershov, and Lei Xu. Algorithmic pricing and competition: Empirical evidence from the german retail gasoline market. 2020. Cited on page 97.
- Zach Y Brown and Alexander MacKay. Competition in pricing algorithms. Technical report, National Bureau of Economic Research, 2021. Cited on page 97.
- Martino Banchio and Giacomo Mantegazza. Artificial intelligence and spontaneous collusion. *Available at SSRN*, 2023. Cited on page 98.
- Ludo Waltman and Uzay Kaymak. Q-learning agents in a cournot oligopoly model. *Journal of Economic Dynamics and Control*, 32(10):3275–3293, 2008. Cited on page 100.
- In-Koo Cho and David M Kreps. Signaling games and stable equilibria. *The Quarterly Journal of Economics*, 102(2):179–221, 1987. Cited on page 150.
- George Mailath, Masahiro Okuno-Fujiwara, and Andrew Postlewaite. Belief-based refinements in signalling games. *Journal of Economic Theory*, 60(2):241–276, 1993. Cited on page 151.
- William Schmidt and Ryan W. Buell. Experimental evidence of pooling outcomes under information asymmetry. *Management Science*, 63(5):1586–1605, 2017. doi: 10.1287/mnsc.2015.2407. Cited on page 151.
- J. J. McCall. Economics of information and job search. *The Quarterly Journal of Economics*, 84(1):113–126, 1970. ISSN 00335533, 15314650. Cited on page 192.
- Martin L Weitzman. Optimal search for the best alternative. *Econometrica: Journal of the*

- Econometric Society*, pages 641–654, 1979. Cited on page 192.
- George J. Stigler. The economics of information. *Journal of Political Economy*, 69(3):213–225, 1961. ISSN 00223808, 1537534X. Cited on page 192.
- Nitin Mehta, Surendra Rajiv, and Kannan Srinivasan. Price uncertainty and consumer search: A structural model of consideration set formation. *Marketing Science*, 22(1):58–84, 2003. doi: 10.1287/mksc.22.1.58.12849. Cited on page 192.
- Babur De Los Santos, Ali Hortaçsu, and Matthijs R. Wildenbeest. Testing models of consumer search using data on web browsing and purchasing behavior. *American Economic Review*, 102(6):2955–80, May 2012. doi: 10.1257/aer.102.6.2955. Cited on pages 192 and 196.
- Peter E Rossi and Greg M Allenby. Bayesian statistics and marketing. *Marketing Science*, 22(3):304–328, 2003. Cited on page 198.
- Vincent D Blondel, Jean-Loup Guillaume, Renaud Lambiotte, and Etienne Lefebvre. Fast unfolding of communities in large networks. *Journal of statistical mechanics: theory and experiment*, 2008(10):P10008, 2008. Cited on page 203.
- Oded Netzer, James M. Lattin, and V. Srinivasan. A hidden markov model of customer relationship dynamics. *Marketing Science*, 27(2):185–204, 2008. doi: 10.1287/mksc.1070.0294. URL <https://doi.org/10.1287/mksc.1070.0294>. Cited on page 212.