Data Valuation in Marketing Collaborations

Ziyao Tang Guang Zeng Paul B. Ellickson [*]

May 27, 2025

Abstract

This paper demonstrates that the value of data can depend on incentive structure. We study a co-branded credit card partnership between a retailer and a bank, focusing on approval decisions. By analyzing treatment effect heterogeneity, we find customers profitable to the bank reduce the retailer's profit, and vice versa, revealing incentive misalignment. Using counterfactual analysis, we show that retail data benefits the bank (+0.72 local dollar per applicant), but harms the retailer (-0.88) because it helps the bank identify customers that are aligned with its objectives but not retailer's. When a participation constraint is added to ensure both parties benefit, joint gains are positive but modest (+0.73). In contrast, when examining a partnership using a linear contract structure, the value of data is over 40 times greater (+32.77). These findings demonstrate that data's value is not intrinsic but shaped by how decisions are made and how gains are allocated between partners.

Keywords: Data sharing; Loyalty program; Co-Branded Credit Card

[*]Ziyao Tang (lead and corresponding author): Simon Business School, University of Rochester, ziyao.tang@simon.rochester.edu. Guang Zeng: Simon Business School, University of Rochester, guang.zeng@simon.rochester.edu. Paul B. Ellickson: Simon Business School, University of Rochester, paul.ellickson@simon.rochester.edu. We are grateful for feedback from participants at seminars at Simon Business School and Wharton AI & Analytics for Business. We also thank Wharton AI & Analytics for Business for connecting us with the data provider. All remaining errors are our own.

1 Introduction

Data is widely regarded as a critical asset that fuels firm profitability — often described as "the new oil" of the digital economy[1] — highlighting its perceived centrality to modern business strategy. This view has strong empirical support in incentive-aligned settings, such as when a single firm controls both data and decision-making. In such contexts, access to more data typically enables more precise targeting, better personalization, and improved profitability. However, as marketing strategies increasingly depend on coordination across business units — whether between firms or across divisions within the same organization — this narrative becomes more complex.

In multi-party collaborations, benefits of additional data become less clear because the decision maker and data owner often have distinct and potentially conflicting objectives. When a single firm makes decisions using its own data, more information typically leads to better outcomes, as it enhances the firm's ability to choose the correct action — whether identifying the right customers to target or determining optimal prices. However, in collaborative settings, the "correct" action for one party may be suboptimal for another. This tension means the impact of additional data becomes ambiguous. The value of data is not intrinsic but shaped by the contract that governs the partnership, specifically how decision rights are allocated and how profits are distributed between parties. When one entity controls decisions, but outcomes affect multiple stakeholders, misaligned incentives can emerge, potentially undermining the benefits of data sharing rather than enhancing them.

Studying the value of data in multi-party settings is particularly important today, as collaborative marketing efforts are now commonplace, with business units jointly leveraging data to drive targeting, product recommendations, and customer engagement. For example, co-branded credit cards — partnerships between merchants and banks — now comprise 62% of consumer credit card portfolios among 12 major issuers. [2] Meanwhile, 64% of U.S. retail executives plan to implement

[1] *The world's most valuable resource is no longer oil, but data*, The Economist. `https://www.economist.com/leaders/2017/05/06/the-worlds-most-valuable-resource-is-no-longer-oil-but-data`

[2] *Why Co-Branded Credit Cards Are Enjoying a Moment*, PaymentsJournal. `https://www.paymentsjournal.com/why-co-branded-credit-cards-are-enjoying-a-moment/`

a Retail Media Network — platforms that allow brands to target ads to shoppers using a retailer's first-party data — by the end of 2024.[3] Beyond that, 70% of global data executives report expanding use of external data sources to overcome internal blind spots.[4] These trends reflect growing enthusiasm for partner-based data strategies. But does access to more data necessarily benefit all parties?

Co-branded credit lending is an ideal setting in which to study how data valuation becomes complicated in collaborative marketing settings. While these arrangements offer mutual benefits —retailers gain enhanced customer loyalty through exclusive promotions and banks expand their cardholder base—they create inherent tensions in value creation and distribution. Banks typically fund the rewards tied to partner retailers' purchases[5] but earn interchange fees on all card transactions regardless of merchant. Retailers, conversely, benefit only from in-store spending. This creates an incentive misalignment problem: banks controlling approval decisions may prioritize customers generating high interchange fee revenue with minimal discount redemption, while retailers prefer those who concentrate spending in their stores. If banks leverage retailer transaction data for approvals, it is unclear whether this benefits the retail partner or potentially makes them worse off. This uncertainty makes co-branded cards an ideal setting to study the value of data sharing in coordinated marketing settings. Despite the importance of understanding these dynamics, empirical research on data valuation in collaborative settings remains scarce, primarily due to limited access to comprehensive data spanning the multiple entities in a partnership.

To bridge this empirical gap, we leverage a unique dataset from a major South American retail-banking conglomerate operating a co-branded credit card program. Given the conglomerate's ownership of both retail and banking divisions, we obtain detailed data on consumer transactions and financial outcomes across the enterprise. This enables us to directly observe how credit decisions affect customer lifetime value for each party: the bank and the retailer. Our

[3] *Retailers, Brands Tap Into the Power of Retail Media Networks*, Deloitte WSJ CMO Today. `https://deloitte.wsj.com/cmo/retailers-brands-tap-into-the-power-of-retail-media-networks-b3bd5078`

[4] *Chief Data Officers: Invest in Your Data Sharing Programs Now*, Forrester Research. `https://www.forrester.com/report/chief-data-officers-invest-in-your-data-sharing-programs-now/RES164496`

[5] See, for example, Loyalty Programs—Once a Perk—Now Help Airlines Survive, Wall Street Journal, 2024; Why Co-Branded Cards Can Be High Risk, High Reward for Banks, Wall Street Journal, 2023.

analysis is uniquely positioned in two ways. First, we observe retail marginal costs and profits, providing a complete measurement of how the program impacts retailer profitability — a significant improvement over previous loyalty program studies that rely primarily on behavioral proxies or revenue-based metrics. Second, despite common ownership, the retailer and bank operate independently with separate management teams and incentive structures, mirroring typical business-to-business relationships. This organizational structure allows our findings to generalize to broader contexts where firms are separately owned and independently optimize their own payoffs.

We begin by developing a theoretical framework that formalizes data valuation in collaborative marketing settings. Our model captures how a retailer and a bank evaluate credit card applicants based on their expected contribution to each business. When the retailer shares data with the bank, the bank can better estimate which customers will be more profitable to it. However, this improved estimation doesn't necessarily make both parties better off. Because the bank and the retailer may have different objectives, finer predictions may actually reveal more instances where an applicant benefits one party (the bank) but harms the other (the retailer). As a result, while data sharing typically benefits the decision-maker (who can better target profitable customers), its impact on the non-decision maker depends on whether the new targeting rule is more or less correlated with the non-decision-maker's optimal targeting strategy, namely, how closely the updated targeting rule aligns with that firm's own objective.

Building on our framework, we proceed to our empirical analysis. We first recover causal estimates of the co-branded credit card approval decision. We exploit quasi-randomness in the approval process to identify causal effects. Applications are initially scored by a proprietary algorithm based on financial data to assess default risk. After scoring, they are assigned to human reviewers who make the final decision. This assignment is primarily based upon reviewer availability, not strategic considerations. Although all reviewers follow the same guidelines, they vary in approval leniency — meaning the same applicant could receive different (counterfactual) approval outcomes depending on the assigned reviewer. This quasi-random assignment introduces exogenous variation to the approval process. We then use a doubly robust difference-in-differences design that accommodates machine learning (ML) models while ensuring valid statistical inference. ML is particularly important here given the high-dimensionality nature of the data. This approach yields estimates of the average treatment effect on the treated (ATT), which we use to assess program effectiveness, and conditional average treatment effects (CATEs), which form the foundation of our counterfactual analysis. We use the CATEs of the relevant outcomes (e.g., retail profit lift, credit card spending, default loss) to construct each party's payoff under different decision scenarios—specifically, with and without data sharing. This allows us to quantify the value of data sharing. We then examine how this value varies across alternative contract structures.

Our empirical results reveal that the co-branded credit card program significantly increases retail profit by 33% on average, demonstrating its effectiveness. However, this average masks substantial heterogeneity: 40.2% of approved customers actually have negative retail profit lift from card approval. These customers primarily use the card outside the partner retailer, concentrating only 14% of their purchases at the retailer compared to 24% for positive-effect customers. This is an incentive misalignment problem: customers who are most valuable to the retailer (generating 473.28 local dollars[6] in incremental profit per applicant) often produce losses for the bank (-$60.49), while those who reduce retailer profit (-$316.83) can still be profitable for the bank ($2.62). This opposing pattern in customer value creates a tension: applicants profitable to one party can harm the other.

Building on our estimates of conditional average treatment effects (CATEs) which we then treat as primitives, we conduct a counterfactual analysis to assess how data sharing and contractual design interact. Under the status quo — where the bank controls approval and retains all interchange revenue — access to retailer data reduces alignment between parties, from 53.8% to 50.8%. The bank gains (+$0.72 per applicant), while the retailer is worse off (-$0.88 per applicant), due to divergent targeting priorities. Strikingly, when we simulate a decision policy that prioritizes the retailer's payoff (retailer-centric), the same data become 50 times more valuable, from

[6] Throughout the remainder of this paper, we will use dollar or $ to refer to this currency.

0.72 to 36.59 dollars per applicant. This dramatic difference arises because, under the retailer-centric scenario, the additional retail data identify higher-lift applicants at the margin compared to the bank-centric scenario. Furthermore, as profit-sharing increases or participation constraints are imposed, data sharing becomes less asymmetric and can even benefit both parties. In addition, when examining data sharing under a linear contract, the value of data is over 40 times greater (+$32.77). These findings demonstrate that data's value is not intrinsic but instead shaped by how decisions are made and how gains are allocated between partners.

These findings have important managerial and policy implications. For managers, they highlight the need to consider not just data access but also how contracts mediate incentives when forming data partnerships. For policymakers, our results suggest that the welfare effects of data sharing depend critically on institutional structures. Policies focused only on access or interoperability — without regard to how decisions and gains are distributed — risk overlooking key determinants of value creation and distribution. Put simply, data regulation should not only consider access but also account for the incentives embedded in decision and payoff structures.

This paper contributes to the growing literature on data economics and the value of data. Previous research has established that data provide substantial value to firms through improved targeting, higher productivity, and smarter decision-making across pricing, product development, advertising campaigns, and customer acquisition efforts (Bergemann et al., 2018; Bergemann and Morris, 2019; Rossi et al., 1996; Brynjolfsson et al., 2011; Lambrecht and Tucker, 2013; Einav and Levin, 2014; Goldfarb and Tucker, 2019; Wu et al., 2023). In related work, Lee et al. (2024) show that grocery purchase data can provide significant incremental predictive power to access creditworthiness, highlighting how behavioral signals outside traditional financial records can improve lending decisions. Similarly, Wernerfelt et al. (2024) find that off-site data can enhance the targeting precision of advertisers, demonstrating the value of cross-contextual data flows. While these studies offer valuable insights, they primarily examine data's benefits in single-firm contexts where a company leverages its own information to enhance performance. Our study extends this line of research by exploring how data creates value when shared between partnering entities in a collaborative marketing arrangement. We quantify not only the predictive value of data, but also how that value is shaped by the strategic and contractual environment in which the data are used when one party shares information with another. We show that contractual arrangements not only influence decision making by different stakeholders, but can also determine the value of data. In doing so, we provide empirical support for the recent theoretical work by Galperti et al. (2024), who argue that the value of data depends critically on its intended use and institutional context, not just on predictive value.

This paper also contributes to the incentive aspects of marketing activities. The marketing literature has documented how multiple stakeholders in marketing activities — advertisers, platforms, agencies, and intermediaries — often operate under misaligned incentives that limit marketing effectiveness (Johnson and Lewis, 2015, Xu et al., 2016, Lewis and Wong, 2022, Frick et al., 2023). Researchers have identified attribution biases, targeting inefficiencies, and measurement distortions as consequences of these misalignments, proposing remedies including incrementality-based measurement, revised contracting structures, and firm-side adjustments. Our paper contributes to this growing literature by providing new empirical simulation evidence on contractual structures directly shaping incentive (mis)alignments in multi-party marketing coordination, as well as an economic framework to analyze how alternative contractual arrangements shift the incentive (mis)alignments.

This paper also relates to the literature on loyalty programs. The literature on loyalty programs has extensively examined their impact on customer behavior, with previous research investigating repeat purchase patterns, customer retention, relational investments, purchase frequency effects, and revenue outcomes (Sharp and Sharp, 1997, De Wulf et al., 2001, Lal and Bell, 2003, Verhoef, 2003, Lewis, 2004, Leenheer et al., 2007, Hartmann and Viard, 2008, Taylor and Hollenbeck, 2021, Gopalakrishnan et al., 2021, Iyengar et al., 2022). While these studies have advanced our understanding of the effectiveness of loyalty program through various behavioral and revenue metrics, the extant literature has primarily focused on these areas due to data constraints regarding profitability metrics at the retail level. Our paper contributes to this rich body of research in two key ways. First, we provide the first direct estimate of co-branded credit card programs' effectiveness on retail profit, leveraging a unique dataset that reveals retail margins, costs, and net profitability at product level. This novel contribution advances the literature by bridging the revenue-profit gap, offering direct retail profit estimates rather than revenue-based proxies, and enhancing managerial decision-making through a profit-centric evaluation framework that acknowledges the reality that revenue increases may not always translate to equivalent profitability gains — a crucial distinction for retailers making investment decisions in loyalty initiatives. Second, we offer new empirical evidence on a rapidly emerging format in loyalty program design — bundled program that combine traditional loyalty mechanisms with co-branded credit cards. These hybrid loyalty-credit card models are increasing common but understudied. Our findings shed light on their effectiveness, providing empirical implications for how such programs may influence profitability beyond traditional engagement or spending metrics.

The rest of the paper is organized as follows: In Section 2, we introduce our theoretical framework that formalizes data valuation in collaboration marketing settings. In Section 3, we detail the institutional background and data. Section 4 outlines our causal identification strategy, estimation approach, and empirical results. Section 5 provides counterfactual simulations that empirically examine the value of data sharing under different contractual arrangements. Section 6 concludes.

2 Theoretical Framework

To formalize and explore the economic tension introduced above, we develop a stylized framework that captures how data sharing affects decisions and payoffs in collaborative marketing scenarios. Our goal is not to model every institutional detail, but to isolate the key forces that govern when and for whom data sharing is valuable. Our framework builds on the mechanism design approach developed by Bergemann et al. (2018); below, we adapt their model to our coordinated marketing setting.

Consider a setting in which the decision to approve or decline a co-branded credit card application, made by customer $i \in \{1, \cdots, N\}$, involves two business units—a bank and a retailer. The approval decision is made unilaterally by the bank, but both parties are affected by the outcome. The bank faces this decision problem under uncertainty, as the profitability of each applicant is not fully known at the time of approval. The bank's action set is $\{\text{decline}(D_i^B = 0), \text{approve}(D_i^B = 1)\}$. For a given applicant i, let $\theta_i = (\theta_i^R, \theta_i^B)$ denote the causal impact (or lift) of approving the application on the retailer's and bank's respective payoffs, relative to the case where the application is declined. By focusing on the incremental effect of approval, we normalize the payoff from declining any applicant to zero.

Critically, the payoff vector θ_i is not intrinsic to the applicant but is shaped by the contractual allocation plan in place between the two parties—such as how interchange fees, rewards costs, or customer value are shared. Different contractual arrangements can lead to different θ_i for the same applicant i, depending on how the costs and benefits of approval are divided. In our theoretical framework, we treat the contract as exogenously given and assume both parties commit to the predetermined allocation plan. This allows us to isolate the role of data sharing. Here, we study how access to additional data affects each party's lift and approval incentive misalignment under a fixed contract. In the empirical analysis, we complement this approach by varying the contract exogenously to examine how different allocation rules shape the value and consequences of data sharing in practice.

Given a fixed contract that determines the payoff vector θ_i, the next question is how the bank makes approval decisions based on its own incentives. Under full knowledge of the causal impact, θ_i, the bank approves applicant i if and only if the approval is expected to increase its own payoff, i.e., $\theta_i^B > 0$.[7] The retailer, by contrast, prefers to approve applicants with $\theta_i^R > 0$. Misalignment arises whenever θ_i^R and θ_i^B differ in sign, as shown in Figure 1.

In practice, the bank does not observe θ_i^B directly. Instead, using its own data X_i^B, it forms an estimate $\hat{\theta}_i^B(X_i^B)$ upon which to base approval. Applicants with $\hat{\theta}_i^B(X_i^B) > 0$ are approved;

[7] We treat $\theta_i = 0$ as not warranting approval.

Figure 1: Alignment Scenarios Diagram

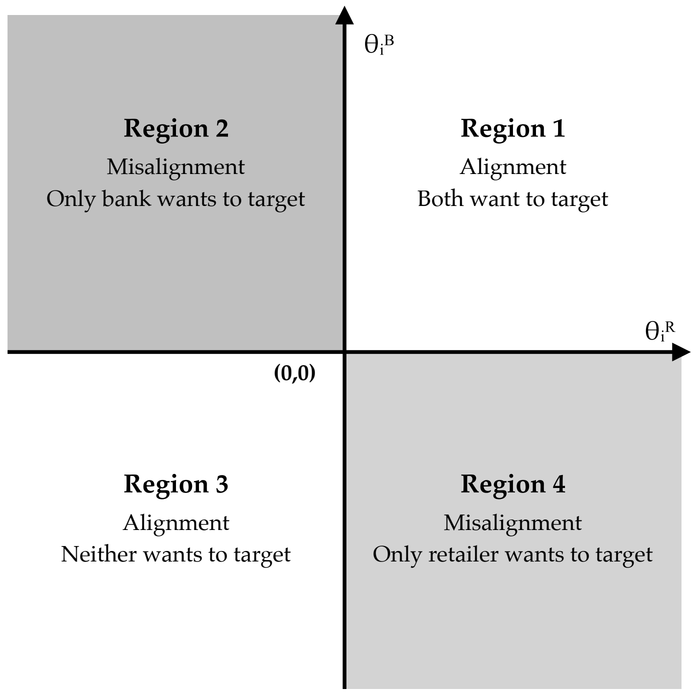

otherwise are declined. The bank's payoff from approval is therefore given by $V^B(X^B) = \sum_{i=1}^N V^B(X_i^B) = \sum_{i=1}^N \theta_i^B \cdot \mathbb{1}[\hat{\theta}_i^B(X_i^B) > 0]$, where $X^B = (X_1^B, \cdots, X_i^B, \cdots, X_N^B)$. Since the approval decision is based solely on $\hat{\theta}_i^B(X_i^B)$, the retailer's payoff is $V^R(X^B) = \sum_{i=1}^N V^R(X_i^B) = \sum_{i=1}^N \theta_i^R \cdot \mathbb{1}[\hat{\theta}_i^B(X_i^B) > 0]$.

Now suppose the retailer shares its data $X^R = (X_1^R, \cdots, X_i^R, \cdots, X_N^R)$ with the bank. The bank can now use both its own data X^B and the retailer's data X^R to estimate applicant i's payoff and make the approval decision. Let the updated estimate be $\hat{\theta}_i^B(X_i^B, X_i^R)$, then the bank approves applicant i if and only if $\hat{\theta}_i^B(X_i^B, X_i^R)) > 0$, and the bank's payoff is $V^B(X^B, X^R) = \sum_{i=1}^N V^B(X_i^B, X_i^R) = \sum_{i=1}^N \theta_i^B \cdot \mathbb{1}[\hat{\theta}_i^B(X_i^B, X_i^R) > 0]$. Similarly, the retailer's payoff is $V^R(X^B, X^R) = \sum_{i=1}^N V^R(X_i^B, X_i^R) = \sum_{i=1}^N \theta_i^R \cdot \mathbb{1}[\hat{\theta}_i^B(X_i^B, X_i^R) > 0]$.

The value of data sharing to each party then is defined as the incremental payoff when the bank has access to the retailer's data. That is, $\Delta V^B = V^B(X^B, X^R) - V^B(X^B)$ and $\Delta V^R = V^R(X^B, X^R) - V^R(X^B)$. Importantly, the value of data sharing can be interpreted as the (lift-weighted) incremental probability of taking the correct action. The notion of a "correct" action is party-specific: from the bank's perspective, the correct action is to approve applicants with $\theta_i^B > 0$, while from the retailer's perspective, it is to approve those with $\theta_i^R > 0$. Therefore, whether data sharing benefits both parties or favors one at the expense of the other depends on how it affects the alignment between their correct actions. If access to additional data increases the overlap in applicants that both parties view as desirable, then data sharing enhances mutual value. Conversely, if the new data shifts the approval policy toward applicants preferred by one party but not the other, it can exacerbate misalignment and reduce the other party's payoff.

While the above analysis assumes the bank controls the approval decision and maximize its own payoff, the core intuition applies equally when the retailer's payoff is maximized (an extreme case where the bank prioritize the retailer's payoff for the sake of long-term partnership). In that case, the bank estimates θ_i^R using X^B and chooses whom to approve based on $\hat{\theta}_i^R$. The value of data sharing and the potential for misalignment remain governed by the same principle: whether additional data improves or distorts alignment between the parties' approval incentives.

3 Empirical Application: Co-Branded Credit Cards

We now turn to our empirical application, which focuses on a large-scale co-branded credit card program between a national retailer and a partnering bank. Our empirical analysis serves two key purposes.

First, we evaluate the effectiveness of this co-branded credit card program under the observed approval policy, specifically the impact on the profits of both the retailer and the bank. Co-branded credit cards are widely used as a marketing tool to foster customer loyalty and boost financial performance. While prior research (Gopalakrishnan et al., 2021; Iyengar et al., 2022; Taylor and Hollenbeck, 2021; Zhao et al., 2022) has shown that offering rewards and discounts can increase customer spending and revenue, these studies often lack a direct analysis of profitability, largely due to limited access to cost data. Moreover, the impact of these programs on the partnering bank has remained largely unexplored for similar reasons. Our empirical analysis addresses these gaps by providing causal estimates of the effect of credit card approval on both parties' profits, leveraging comprehensive data on both revenue and cost.

Second, we investigate incentive misalignment and the value of data sharing under alternative contract structures that are commonly observed in practice. In this counterfactual analysis, we use conditional average treatment effects (CATEs) as the primitive objects with which to uncover how data sharing affects the divergence between retailer and bank. Using our causal estimates based on these detailed, real-world credit card approval and transaction data, we empirically quantify how misaligned incentives affect business outcomes and assess the strategic and economic value of data sharing in a collaboration setting.

To address our first goal — assessing the effectiveness of co-branded credit cards program — we employ a difference-in-differences (DiD) design that leverages variation in credit card approval decisions. Specifically, we compare the before-and-after changes in outcomes for approved customers (treated group) versus declined customers (control group), centered around the time of the credit card approval decision. This approach allows us to estimate the causal effect of receiving a co-branded credit card on both retail and financial outcomes.

The main outcomes we analyze include: (1) retail profit, (2) retail spending, (3) credit card discount usage, (4) credit card spending at the retailer, (5) total credit card spending, and (6) default losses incurred by the bank. To estimate the effectiveness of the observed approval policy, the target estimand is average treatment effect on the treated (ATT). But to conduct the counterfactual analysis of other contract structures, we estimate conditional average treatment effects (CATEs) for each of these outcome components. With estimated CATEs as the primitive objects, we can recover expected individual-level payoff for both the retailer and the bank, under various contractual arrangements. These constructs are crucial for understanding the heterogeneous impact of the program and for evaluating how data sharing affects business outcomes under different incentive structures.

3.1 Institutional Background

We now turn to the institutional context and data that enable our analysis. The co-branded credit card program we study involves a partnership between a major national retailer and a large financial institution, offering a rich setting with detailed customer-level information on applications, transactions, and outcomes. In what follows, we describe the key features of the program and the specific datasets used in our analysis.

Our study draws on proprietary data from a large conglomerate in South American that operates in both the retail and financial sectors. The co-branded credit card program is jointly operated by the retail and banking divisions of the conglomerate, with credit approval decisions made solely by the banking division. Once approved, cardholders can use the proprietary credit card both at the conglomerate's retail stores and at external merchants. When customers purchase with the card at the retailer, they enjoy payment method-exclusive discounts on selected products that are not available with cash, debit cards, or other (third-party) credit cards. Each month, the retailer designates a rotating set of these cardholder-exclusive promotions across multiple categories such as groceries and apparel. The discount is not personalized to individual customers but applies uniformly to all cardholders.

Importantly, the full cost of these discounts is covered by the bank, effectively making them a cross-division transfer from financial to retail units. In return, the bank also benefits financially from credit card usage through interchange fees. For each transaction made with the co-branded card — whether at the retailer or at an outside merchant — a 3% interchange fee is charged to the merchant processing the transaction. Of this amount, 1% is paid to the transaction network (Visa), while the remaining 2% accrues to the bank. Thus, the bank earns the same interchange revenue regardless of where the card is used, while it incurs additional costs only when the card is used at the retailer (via the subsidized discounts). This asymmetry introduces potential incentive misalignment. The bank prefers customers who spend outside the retailer's stores, where transactions generate interchange revenue without triggering discount costs. Conversely, the retailer prefers customers who concentrate their spending within its stores, leveraging the discount-driven customer loyalty despite incurring costs to the bank. This difference in cost-benefit exposure across divisions plays a central role in our analysis of incentive alignment and data-sharing strategy. Given that cardholder benefits — and the associated costs and revenues for each division — are only realized after approval, the credit approval decision plays a central role in shaping the effectiveness of the program and the distribution of value between the bank and the retailer. We next describe in detail how these approval decisions are made in practice.

The approval process begins with a proprietary risk-scoring algorithm developed by the bank to assess applicants' likelihood of default. This algorithm relies exclusively on financial information — such as outstanding debt, and credit history — and produces a risk score used to guide the decision-making process. After scoring, applications assigned to human reviewers, who verify income and employment documentation and make the final approval decision. This assignment is based on reviewer availability rather than any strategic considerations. While all reviewers follow the same general approval guidelines, they differ in their leniency — some are more likely to approve marginal applicants than others. Importantly, this variation in reviewer behavior is plausibly unrelated to applicants' future outcome trends — such as retail spending or credit default — conditional on their financial information. As a result, this quasi-random assignment generates exogenous variation in approval decisions, which we will leverage in our identification strategy, described in detail when we introduce our identification strategy.

3.2 Data and Descriptive Evidence

We now turn to a detailed description of the data used in our analysis. We begin with all customers who applied for the proprietary credit card during 11 distinct application months between July 2021 to May 2022, resulting in 271,260 unique applicants. For each applicant, we construct a panel that tracks their behavior from six months before the application decision (month $= -6$) to thirteen months after (month $= +12$), where month 0 denotes the first post-approval decision month. We merge retail transaction records from the retailer and financial information, such as card usage from the bank, to build a unified dataset capturing both pre- and post-application behavior across the two business units.

The transaction data is sourced from the retail unit, which spans multiple product categories: groceries, electronics, clothing, and home goods. The dataset contains 3,531,323 transactions made by the sampled applicants. The data captures item-level transaction details, and each transaction record includes the timestamp, quantity purchased, unit price, any credit card-specific discounts applied, the marginal cost (procurement) to the retailer, sales tax on each item, and the payment method used (e.g., cash, debit card, or credit card). This rich information allows us to compute key outcome variables such as retail spending, retail profit, and discounts received at the transaction level. We then aggregate these outcomes to the applicant-month level, which forms the basis of our panel dataset used for empirical analysis. As illustrated in Table 1, customers in our sample make, on average, 0.74 shopping trips per month, spend approximately, \$152.60 per month, and generate an average monthly profit of \$16.74 for the retail unit.

Table 1: Consumer-Level Summary Statistics

	Mean	Std. Dev.
Trips/Month	0.74	0.80
Retail Spend/Month	152.60	214.87
Retail Profit/Month	16.74	29.32

The bank's database integrates two primary sources of information related to the proprietary credit card program. The first is the credit card application data, which includes detailed records on applicants. This dataset contains information such as the application date, the approval decision (approved or declined), the date of decision, and applicant-level demographic and financial attributes, including self-reported income and occupation. The second source is the credit registry data from a national banking supervisory authority in its region. This monthly data provides a comprehensive view of each applicant's financial engagements across the broader financial system, not limited to the focal conglomerate. It includes records on total outstanding debt, instances of delinquency, types of credit instruments held (e.g., credit cards, loans), credit card utilization rates, and current balances. Using these two sources, we define a default event as the failure to repay credit card debt for more than 60 days past the due date. This definition aligns with both the internal risk management standards of the bank and industry norms. Additionally, the bank uses this data to generate a risk score, representing the estimated probability of default for each applicant, which plays a central role in guiding approval decisions.

By combining the credit application records with the external credit registry data, we construct a detailed profile of each applicant's financial health and creditworthiness. This integrated view enables us to assess the financial implications of credit card issuance and usage for the bank and, when linked with the retail data, to examine how approval decisions influence business outcomes across both the bank and the retailer.

With the data sources and variable construction in place, we now present preliminary descriptive evidence to illustrate key patterns in approval decisions, customer characteristics, and outcomes. These patterns provide initial insights into how approved and declined applicants differ, how customers use the co-branded credit card post-approval, and how such usage connects to the business performance of both the bank and the retailer.

We next present descriptive evidence of how purchase behavior differs before and after customers apply for the credit card. We use declined applicants as the control group and approved applicants as the treated group, allowing us to draw preliminary comparisons in spending, shopping frequency, and overall profitability between these two sets of customers in the pre- and post-application periods. The results are displayed in Table 2: On average, declined applicants show only modest increases in monthly trips, spending, and profit. By contrast, approved applicants exhibit a substantially larger uptick in all three measures — trips per month, spending per month, and profit per month — following their credit card approval, foreshadowing the potential impact of credit card adoption on customer engagement and profitability. This simple comparison implies that credit card adoption may be associated with increased trips, spending, and profit. However, to causally attribute these differences to the credit card itself — rather than underlying differences between approved and declined applicants — we introduce a more formal causal analysis in Section 4.1.

Table 2: Consumer-Level Summary Statistics

	Control(After)	Control(Before)	Treated(After)	Treated(Before)
Trips/Month	0.42	0.39	0.88	0.55
	(0.58)	(0.69)	(0.91)	(0.86)
Retail Spend/Month	63.86	56.75	193.02	93.71
	(124.19)	(167.08)	(258.85)	(225.22)
Retail Profit/Month	8.38	8.18	19.94	12.68
	(18.24)	(26.64)	(36.00)	(35.60)
Default Loss	0.00	0.00	70.76	0.00
	(0.00)	(0.00)	(687.16)	(0.00)
Observations	40455	40455	230805	230805

4 Causal Effect Estimation

4.1 Doubly Robust Machine Learning Diff-in-Diff Estimation

We now outline our identification strategy. The approval decision for credit card applications is primarily based on a risk score trained on the bank's data. When an applicant submits an application, it is assigned to a reviewer who makes the final decision, guided by this risk score. Recall that the assignment of applications to reviewers is random. This randomness induces variation in approval decisions among customers with the same risk score, providing a source of exogenous variation.

To estimate causal effects, we employ a doubly robust difference-in-differences (DiD) framework (Callaway and Sant'Anna, 2021; Sant'Anna and Zhao, 2020), which integrates two key components: an outcome model and a propensity score model. For the outcome model, we model outcomes as a flexible function of pre-treatment (pre-application) customer characteristics. When evaluating the impact of the co-branded credit card loyalty program on both the retailer and the bank, we incorporate customer characteristics from both the bank's and retailer's databases, allowing us to capture heterogeneous effects of pre-application characteristics on post-approval outcomes. Similarly, in the propensity score model, we include customer characteristics in the same manner as in the outcome model. Most importantly, we include the risk score, as it is the primary observable used by the bank in approval decisions. We also include other pre-treatment variables to improve the precision of the causal estimates.

The doubly robust approach not only enhances the validity of the causal estimates but also provides greater flexibility in estimation. As long as either the outcome model or the propensity score model is correctly specified, the causal estimates remain consistent. Moreover, this approach allows for the incorporation of machine learning (ML) algorithms to estimate these models while still enabling valid statistical inference. The use of ML estimators is particularly important in this setting, as the databases contain more than a hundred variables. ML methods excel at modeling high-dimensional functions and capturing complex, flexible functional forms, which enables a more precise characterization of individual causal effects.

We now introduce the DiD approach in our context. We define group $g \in \{1, \cdots, 11\}$ based on the month when the applicants receive credit card approval. For each approval-month group g, we define their event time relative to the approval decision as $t \in \{-6, \cdots, 0, \cdots, T\}$, where $t = 0$ represents the month in which the approval decision is made. For example, if group g consists of applicants approved in January 2022, then for this group, $t = 1$ corresponds to February 2022 (one month after approval), and $t = -1$ corresponds to December 2021 (one month before approval). We set $T = 12$ to evaluate the impact of approval on key outcomes over a 13-month period, capturing its contribution to customer lifetime value.

To conduct valid causal inference and enable counterfactual analysis with a DiD design, we lay out the key assumptions:

Assumption 1 (Irreversibility of Treatment). *For any given applicant i, treatment assignment is given by $D_{i,t} \in \{0, 1\}$ for all t, where $D_{i,t} = 0$ represents a declined applicant (control), and $D_{i,t} = 1$ represents an approved applicant (treated). In addition, treatment status is non-reversing over time, such that for all t, $D_{i,t-1} \leq D_{i,t}$. Furthermore, we assume that no applicant is treated before the approval decision month, meaning that $D_{i,-1} = 0$.*

Assumption 1 states that once an application is approved and the credit card is issued, the consumer does not cancel the card during the sample period. Empirically, credit card cancellations are extremely rare within our sample window, making this assumption reasonable in our context.

Assumption 2 (Conditional Parallel Trends; Unconfoundedness With Respect To Trend)**.** *For any given group* g, $E[Y_{i,t}(0) - Y_{i,t-1}(0)|X_i, g_i = g, D_{i,t} = 1] = E[Y_{i,t}(0) - Y_{i,t-1}(0)|X_i, g_i = g, D_{i,t} = 0]$ *and* $E[Y_{i,t}(1) - Y_{i,t-1}(1)|X_i, g_i = g, D_{i,t} = 1] = E[Y_{i,t}(1) - Y_{i,t-1}(1)|X_i, g_i = g, D_{i,t} = 0]$, *where* $Y(0)$ *and* $Y(1)$ *denote the potential outcomes under the control and the treated conditions, respectively.*

Assumption 2 enables us to impute the treated group's counterfactual outcome under the control condition using the control group's observed outcome and, conversely, to impute the control group's counterfactual outcome under treatment using the treated group's observed outcome. Unlike the unconfoundedness condition in a static setting, this is an unconfoundedness condition with respect to trends rather than levels. By imposing this assumption, we can identify not only the average treatment effect on the treated (ATT), which reflects the effect of the observed policy, but also the average treatment effect (ATE) and, in particular, the conditional average treatment effect (CATE). We require CATEs as our primitive objects so we can study the effect under a counterfactual policy. This assumption is stronger than the usual one researchers have made to identify ATT, but it can be justified based on our understanding that review agents are effectively randomly assigned, ensuring that there is no systematic difference in the trends after conditioning on pre-application financial characteristics. Additionally, by conditioning on g_i, we allow for group-specific trends, which provides further flexibility/robustness in capturing heterogeneous dynamics across applicant groups.

Assumption 3 (Sufficient Overlap)**.** *The probability (propensity score) of being approved is strictly bounded between 0 and 1:* $0 < P(D_{i,t} = 1|g_i, X_i) < 1$.

Assumption 3 ensures that there is sufficient overlap in the covariate distribution for both treated and control groups. This assumption can be justified based on the random assignment of review agents, and we also show evidence that it is satisfied in Appendix A.1. Assumptions 2 and 3 together allow for meaningful counterfactual comparisons.

After having stated our assumptions, we now move to our target causal parameters and the estimation methods. We first define $\Delta Y_{it} = Y_{it} - Y_{i,-1}$ as our observed outcome (we choose event month $t = -1$ as our baseline period, as is common practice), as well as $\Delta Y_{it}(0) = Y_{it}(0) - Y_{i,-1}(0)$ and $\Delta Y_{it}(1) = Y_{it}(1) - Y_{i,-1}(1)$ as the potential outcomes under the control and the treated conditions, respectively. Because the unconfoundedness holds with respect to the trend, we then have the following doubly robust formulations of average treatment effect on the treated (ATT) and average treatment effect (ATE) for group g in month t, as derived by Sant'Anna and Zhao (2020); Callaway and Sant'Anna (2021):

$$\begin{aligned}\theta_{ATT}(g,t) &= E[\Delta Y_{it}(1) - \Delta Y_{it}(0)|g_i = g, D_i = 1] \\ &= E\Big\{[\frac{D_i}{E[D_i]} - \frac{e_g(X_i)(1-D_i)}{1-e_g(X_i)}/E[\frac{e_g(X_i)(1-D_i)}{1-e_g(X_i)}]] \cdot [\Delta Y_{it} - \mu_{0t,g(X_i)}]\Big\}\end{aligned} \tag{1}$$

$$\begin{aligned}\theta_{ATE}(g,t) &= E[\Delta Y_{it}(1) - \Delta Y_{it}(0)|g_i = g] \\ &= E\Big\{[\mu_{1t,g}(X_i) - \mu_{0t,g}(X_i)] \\ &+ [\frac{D_i \cdot (\Delta Y_{it} - \mu_{1t,g}(X_i))}{e_g(X_i)} - \frac{(1-D_i) \cdot (\Delta Y_{it} - \mu_{0t,g}(X_i))}{1-e_g(X_i)}]\Big\}\end{aligned} \tag{2}$$

For notational simplicity, when presenting the doubly robust estimators, we treat g as defining the population of interest and omit indicator functions that index whether an observation belongs to group g. This allows us to streamline the notation without loss of generality, as all quantities are interpreted conditional on membership in group g. In these equations, $e_g(\cdot)$ represents the propensity score model for group g, which characterizes the probability of treatment assignment given observed characteristics. The functions $\mu_{0t,g}(\cdot)$ and $\mu_{1t,g}(\cdot)$ denote the outcome models for the potential outcomes under the control and the treated conditions, respectively, in month t for group g. These models capture how the observed characteristics influence the expected outcomes for each treatment condition. For our empirical analysis, we aggregate the group-time ATT,

$\theta_{ATT}(g,t)$, to conduct an event study analysis, which tracks the dynamic effects of the approval decision over time. This aggregation facilitates the examination of both short-term and long-term impacts of the co-branded credit card approval on key outcomes. Following Callaway and Sant'Anna (2021), we compute this aggregated effect by taking a weighted average of $\theta_{ATT}(g,t)$, using the relative frequency of each group g as weights for each event time period t.

From the above discussion, it is clear that conducting the event study requires us to first estimate both $\theta_{ATT}(g,t)$ and $\theta_{ATE}(g,t)$. Estimating these causal parameters, in turn, depends on obtaining reliable estimates of the propensity score model, $e_g(\cdot)$, and the outcome models, $\mu_{0t,g}(\cdot)$ and $\mu_{1t,g}(\cdot)$. To estimate these functions, we employ random forest models. The random forest is a nonparametric machine learning method well-suited for this setting, as it can flexibly approximate complex, nonlinear relationships without requiring strong assumptions about the underlying functional forms — an important consideration given the high dimensionality and complexity of our covariates. The doubly robust formulations of these causal parameters ensure that the use of machine learning algorithms does not introduce regularization bias and does not compromise valid statistical inference. As shown by Chernozhukov et al. (2018), even when flexible ML methods such as random forest are used to estimate the nuisance components (e.g., propensity scores and outcome models), the resulting doubly robust treatment effect estimators remain consistent and asymptotically normal, provided standard sample-splitting or cross-fitting procedures are applied.

To obtain the estimated doubly robust scores (i.e., the terms in $E\{\cdot\}$ of Equation 1 for ATT and Equation 2 for ATE), we follow a two-step procedure. First, for each group-time pair (g,t), we randomly partition the data into approximately equal-sized folds, indexed by $f \in \{1, \cdots, F\}$. For each fold f, we estimate the nuisance components — namely, the propensity score model $\hat{e}_{[f]}$ and the outcome models $\hat{\mu}_{[f]}$ — using only the data outside of fold f. In the second step, we apply the corresponding out-of-fold estimates $\hat{\mu}_{[f]}$ and $\hat{e}_{[f]}$ to each observation and compute its doubly robust score using the relevant formula. These scores are then aggregated to construct the estimator for the target causal parameter.

4.2 Causal Effects of Credit Card Adoption

We now examine the causal impact of credit card approval. In what follows, we present estimates of the treatment effects on key retail and financial outcomes, offering direct insight into how credit card adoption affects customer behavior and drives value for both the retailer and the bank.

We begin by presenting ATTs for key outcome components. Figure 2 plots the event-time dynamics of treatment effects for six key constructs — changes in shopping frequency, spending behavior, discount utilization, and overall retailer profitability, highlighting both short-term adjustments and long-term trends.

A clear pattern emerges across these key metrics: an initial jump immediately after approval, followed by a decline that eventually stabilizes at a substantial and persistent increase. For store visits, the ATT rises sharply at month 0 and eventually stabilizes at approximately 0.2 additional trips per month, representing a 29% increase.[8] This sustained increase indicates that the card effectively drives greater in-store engagement over time. Similarly, the effect on spending at the retailer peaks at approximately 200 dollars per month immediately after approval, before stabilizing at around 50 dollars monthly, translating to a 35% increase. This pattern suggests that the promotional discounts tied to the card effectively influence consumer behavior, with an impact that persists well beyond the initial adoption period. Retail profit follows a comparable trajectory, stabilizing at around 5 dollars per month, a 33% increase over pre-approval levels. This substantial and persistent lift in profitability demonstrates the card's effectiveness.

Card usage sustains over time, with total spending stabilizing around $400 per month. Spending at the partner retailer spikes to $250 immediately after approval, then stabilizes around $50. This gap indicates that a substantial share of card spending occurs outside the partner retailer.

While the average effect on retail profits is positive, further analysis reveals substantial heterogeneity in individual-level responses to credit card approval. As shown in Figure 3, 40.2%

[8] This is based on a back-of-the-envelope calculation using the $\frac{\text{Effect}}{\text{Observed post-approval outcome} - \text{Effect}}$. Observed post-approval values are from the *Treated (After)* column in Table 2. The percentage increase of spending and profit below are calculated similarly.

Figure 2: ATTs of Credit Card Adoption

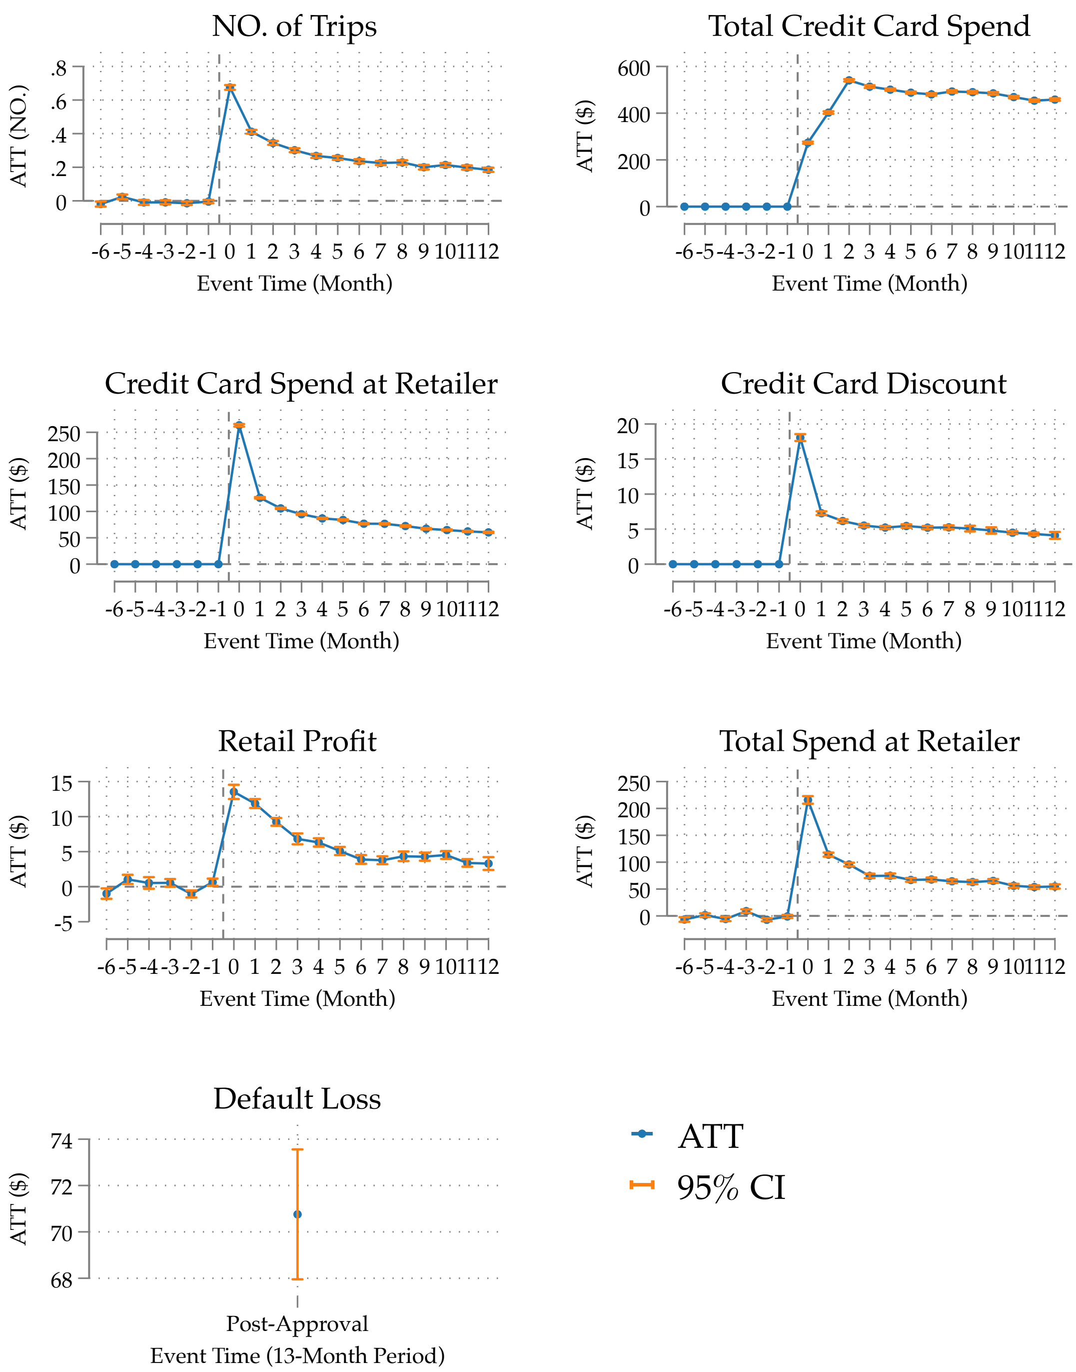

Notes: This figure displays the dynamic average treatment effects on the treated (ATTs) under the observed credit approval policy. The outcomes correspond to key dimensions of the co-branded credit card program. The subplots for Total Credit Card Spend, Credit Card Spend at Retailer, and Credit Card Discount pertain specifically to the focal co-branded card, with pre-approval values essentially at zero. Event time 0 marks the approval month. Default loss reflects the cumulative amount over the 13-month post-approval period.

of approved applicants exhibit negative retail profit effects — meaning the retailer would have been better off not approving these customers. This finding adds a new perspective to the loyalty program literature, which has largely examined whether such programs justify the firm's own spending on rewards (Liu and Yang, 2009; McCall and Voorhees, 2010; Bijmolt et al., 2011). Our setting differs in that the retailer does not bear the cost of incentives, yet approving certain customers still results in worse outcomes than not approving them.

Figure 3: Histogram of Retail Profit Effects

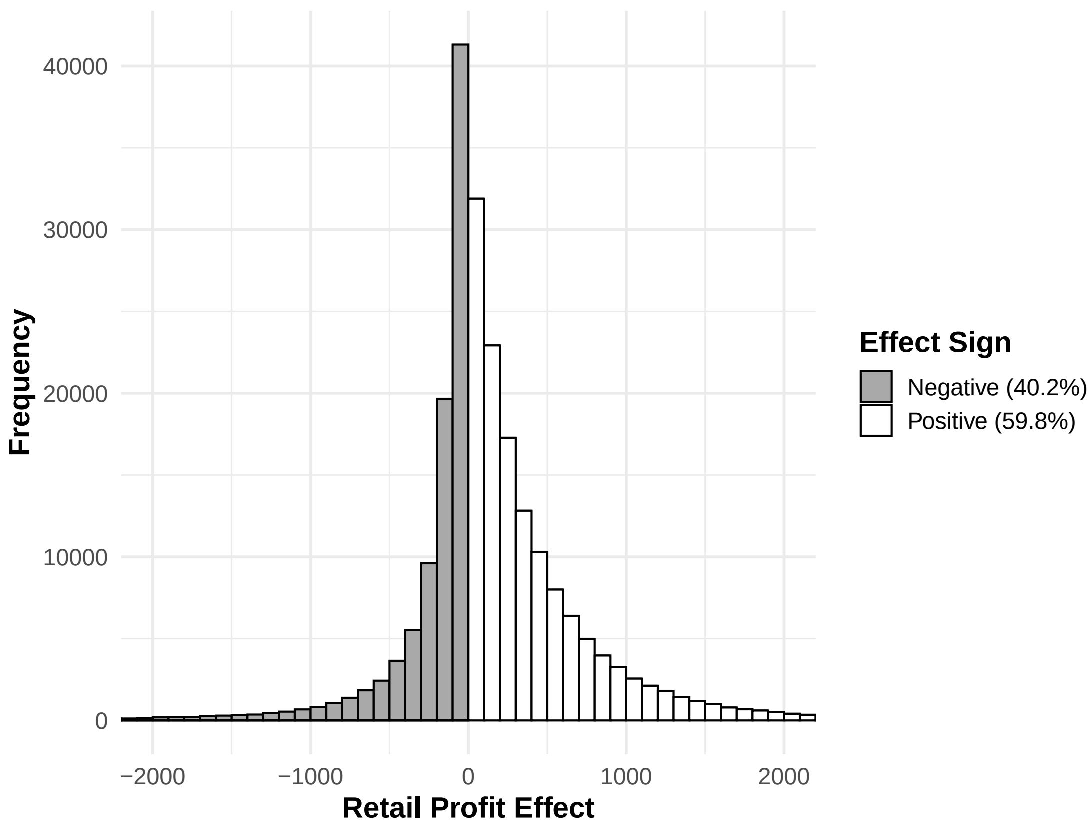

Notes: This histogram shows the distribution of cumulative retail profit effects across approved applicants, measured as the total over the 13-month post-adoption period. The range is truncated at $[-2{,}000, 2{,}000]$ to improve readability, removing 2.36% of the approved applicants. Retail profit reflects revenue net of credit card discounts, which are financed by the bank.

To understand why some customers generate negative profit effects upon approval, Table 3 compares key behavioral metrics across groups. Suprisingly, the decline in retail profit isn't explained by excessive discount usage. On the contrary, the negative effect group redeems fewer total discounts and obtains lower discount amounts per trip compared to the positive-effect group, indicating they're less effective at leveraging the card's promotional benefits. The key difference emerges in their card usage patterns. Compared to their counterfactual without the card, negatively impacted customers visit the retailer less frequently and reduce their overall spending. More revealing is that they concentrate only 14% of their credit card purchases at the partner retailer, compared to 24% for the positive effect group. This pattern suggests that for these customers, the value of general credit access outweighs the retailer-specific promotional benefits, leading them to primarily use the card elsewhere which limits the retailer's return on these approvals.

Table 3: Changes in Shopping and Spending Behavior by Δ Retail Profit Group

	Δ Retail Profit Group	
	Positive	Negative
Δ Retail Profit	473.28	-316.83
#Trip	14.36	7.24
Δ Trip	8.64	-4.17
Retail Spending	3227.17	1442.18
Δ Retail Spending	2285.22	-999.79
Credit Card Discount Used	124.36	19.36
Credit Card Spending (Partner)	1668.34	622.11
Credit Card Spending (All)	7090.77	4513.17

These distinct usage patterns across profit-effect groups have important implications for incentive misalignment between the bank and the retailer in co-branded credit card programs. To investigate this misalignment, we examine how each customer's lifetime value (LTV) is calculated under the current contractual arrangement. For spending with the co-branded credit card, a 3% interchange fee is charged: 1% goes to the network (Visa) and 2% to the bank. The bank also fully covers the cost of card discounts. Based on this allocation, the payoffs to each party are calculated as follows:

$$V_{\text{retailer}} = \Delta\text{retail profit} - (\text{Co-branded Credit Card Spending}_{\text{partner}} \times 3\%)$$

$$V_{\text{bank}} = (\text{Co-branded Credit Card Spending}_{\text{all}} \times 2\%) - C_{\text{discount}} - \text{Default Loss} - \text{Card Production Cost}$$

We exclude interest revenue from the card profit calculation, as it is very rare for cardholders not to pay their balance in full—fewer than 0.01% of cases in our data. We assume a constant card production cost of $2.5, based on industry estimates.[9] The retailer's payoff reflects the incremental retail profit net of interchange fees paid on in-store card usage. The bank's payoff includes interchange revenue from all card transactions (both at the retailer and at outside merchants), minus the costs of discounts, default loss, and card issuance.

When we compare the two parties' payoffs across the positive and negative retail profit effect groups, a clear misalignment emerges. As shown in Table 4, customers who reduce the retailer's profit (−$316.83 per person) still generate a small but positive return for the bank ($2.62). This is driven by their distinctive spending behavior: they use the card primarily outside the partner retailer, resulting in $90.26 in interchange revenue for the bank. At the same time, their limited in-store purchasing leads to relatively low discount costs ($19.36). Although their total spending is lower than the positive group, the reduced promotional burden more than offsets the decline in interchange revenue. They also exhibit slightly lower default losses ($65.78 vs. $75.44), further contributing to their net positive value to the bank.

In contrast, customers who are most valuable to the retailer ($423.23 per person) impose losses on the bank (−$60.49). Their spending is concentrated in-store, which increases retailer profit but also triggers high discount costs ($124.36) that exceed the bank's interchange revenue. This inverse relationship in customer value highlights a fundamental challenge in co-branded credit card programs: approval decisions that benefit one party may simultaneously erode value for the other.

Overall, these results provide strong evidence that the co-branded credit card program in-

[9] https://jleconsultants.com/how-much-does-a-credit-card-cost-to-make

Table 4: Payoff Components by Δ Retail Profit Group

	Δ Retail Profit Group	
	Positive	Negative
Interchange Fee	141.82	90.26
Credit Card Discount Cost	124.36	19.36
Default Loss	75.44	65.78
Δ Retail Profit	473.28	-316.83
Bank Payoff	-60.49	2.62
Retailer Payoff	423.23	-335.49

creases joint lifetime value for the bank and the retailer.[10] Yet this overall gain conceals important differences across customers. Over 40% of approved applicants reduce the retailer's profit relative to their counterfactual without the card. Compared to the profitable group, these customers are more likely to use the card outside the partner retailer, indicating that the credit effect outweighs the discount effect. Looking more closely at the same comparison from the perspective of both parties' payoffs, we find an asymmetry: the customers who are most valuable to the retailer often generate losses for the bank, and vice versa. This divergence in lifetime value reveals a structural tension in the program — approval decisions that benefit one party can simultaneously harm the other.

These patterns motivate our counterfactual analysis in the next section, which focuses on the challenge of valuing data when the data user and the data owner are not the same entity. So far, we have examined average treatment effects under the observed policy. We now move beyond that to conduct a counterfactual analysis using conditional average treatment effects (CATEs) as primitives. This will allow us to quantify how data sharing can intensify or mitigate incentive misalignment, and how it changes the payoff of each party under alternative contract structures.

[10] See Section A.3 for estimates of lifetime value for the retailer, the bank, and the combined total.

5 Incentive Misalignment, Data Sharing, and Value of Data

Motivated by the findings above, we now turn to a set of counterfactual simulations to explore the dynamics of incentive alignment, the value of data sharing, and the potential role of contractual coordination mechanisms. Before presenting counterfactual analyses, we first provide evidence that the retail shopping attributes of applicants prior approval (X^R) can increase the precision of the predictions of treatment effects (thereby enabling clearer segmentation). We focus on an outcome variable for which the treatment effect is directly observable — specifically, where the non-treated counterfactual is zero. Table 5 reports the improvement in prediction accuracy when incorporating X^R, in addition to the bank's information (X^B). The results indicate a modest improvement for default loss prediction, but a substantial gain in accuracy for outcomes such as credit card discount usage and credit card spending behavior.

Table 5: Root Mean Squared Error by Outcome and Data Sources

Outcome	X^R & X^B	X^B	Ratio
Default Loss	82.663	84.491	0.978
Credit Card Discount Used	56.522	109.017	0.518
Credit Card Spending (Partner)	311.455	896.445	0.347
Credit Card Spending (All)	855.149	1891.444	0.452

Notes: The RMSE is calculated among the approved (treated) individuals. The ratio is the RMSE with both X^R and X^B divided by the RMSE with X^B only. A lower ratio indicates greater predictive gain from retail data.

To isolate the effect of data access from other institutional factors, we simulate counterfactual scenarios in which the bank adopts an optimal policy that maximizes its own expected payoff—rather than relying on the observed (factual) approval decisions. We avoid using the factual policy as a baseline for two reasons. First, the actual approval decisions were made by senior management within the conglomerate, who may have considered a broader set of strategic or organizational priorities beyond pure financial optimization. These considerations may reflect context-specific practices that are not generalizable to typical co-branded credit card partnerships. Second, by comparing two counterfactual policies — both optimized under clearly defined objective functions — we can isolate the effect of data access. Holding constant the decision-maker's objective and the underlying targeting algorithm ensures that differences in outcomes reflect only the difference in information set (i.e., access to X^R), rather than institutional discretion or unobserved decision criteria.

We consider a class of profit-sharing contracts as our main payoff distribution mechanism, where the parties split the profit generated by the card — interchange revenue minus defaults and card discounts as well as card production cost — to examine the valuation of data for several reasons. Such arrangements are widely prevalent in co-branded credit card partnerships across various industries. As noted in industry analyses, profit sharing is recognized as one of the fundamental frameworks for these agreements.[11] Evidence from the airline industry shows that payments from credit card partners generate approximately 15% of airlines' revenue,[12] which implies the presence of profit-sharing mechanisms. In addition to its prevalence, profit sharing creates a direct alignment of financial interests between the retailer and the issuing bank, as the retailer's earnings become intrinsically linked to the overall success and profitability of the co-branded credit card program.

A defining feature of these programs—especially in retail settings—is that the issuing bank covers the cost of cardholder rewards and benefits. Banks are willing to take on this cost burden as part of the broader negotiation to secure and maintain partnerships with high-profile or high-volume retail brands. In competitive co-brand markets, banks frequently agree to such terms in exchange for the growth potential or strategic value of the partnership. [13]

Under these contractual arrangements, the retailer receives a share of the profits generated by the co-branded credit card, in addition to its direct retail profit. The payoffs to the retailer and

[11]See, for example, discussions in *CRM Trends: Credit Card Programs* (`http://www.crmtrends.com/CreditCardPrograms.html`) and the CFPB's report *Issue Spotlight: The High Cost of Retail Credit Cards* (`https://www.consumerfinance.gov/data-research/research-reports/issue-spotlight-the-high-cost-of-retail-credit-cards/`).

[12]*Co-Branded Credit Cards: The Allure and The Reality* (`https://bankingandpaymentsgroup.com/co-branded-credit-cards-the-allure-and-the-reality/`)

[13]See, for example, Why Co-Branded Cards Can Be High Risk, High Reward for Banks, Wall Street Journal, 2023.

the bank are defined as follows:

$$V_{\text{retailer}} = \Delta\text{retail profit} + (\pi_{\text{card}} \times s)$$

$$V_{\text{bank}} = \pi_{\text{card}} \times (1 - s)$$

where $s \in [0, 1]$ denotes the share of card-generated profits allocated to the retailer, and π_{card} represents the net profit from the credit card program, calculated as:

$$\pi_{\text{card}} = (\text{Co-branded Credit Card Spending}_{\text{all}} \times 2\%) - C_{\text{discount}} - \text{Default Loss} - \text{Card Production Cost}$$

As previously noted, we exclude interest revenue from card profit, since fewer than 0.01% of cardholders do not pay in full each month.

This class of contracts nests the observed contract as a special case where $s = 0$, meaning the retailer receives no share of the card's financial returns beyond retail profit, and the bank retains full exposure to both the upside and downside of card usage.

This profit sharing arrangement is well suited to our research question because it allows us to quantify the degree of incentive alignment using the scalar s, which governs how card-generated profit (π_{card}) is shared between the retailer and the bank. When s = 0 (observed), the retailer receives no share of card profits and evaluates applicants solely based on their direct contribution to retail profit, minus any interchange fees paid on card transactions. As s increases towards 1, the retailer gradually internalizes card profitability, becoming more willing to approve applicants who generate substantial card revenue but modest retail margin. This shift in incentives creates a natural alignment spectrum: at lower values of s, the retailer and bank often disagree on which applicants to target, while at higher values, their approval preferences increasingly converge. By systematically varying this parameter, we can empirically measure how different profit sharing structures affect targeting decisions, data sharing outcomes, and the overall value distribution between partners.

5.1 Data Sharing and Incentive Misalignment

We begin by examining how data sharing affects the degree of incentive (mis)alignment between the retailer and the bank. As outlined in our theoretical framework (Section 2), we define alignment as the share of applicants for whom both parties independently reach the same approval decision — either both approve or both decline — based on their respective payoffs.

To isolate the effect of data sharing on alignment, we compare two distinct counterfactual information environments: one where the decision maker makes decisions using only bank data (X^B), and another where it also has access to the retailer's data (X^B and X^R). This analysis does not assume that either party actually has access to the other's data in practice, nor does it imply that they should share it. Instead, it examines how access to additional data (X^R) would affect the degree of agreement in their approval preferences – holding fixed the contract and each party's objective. A decline in alignment rate — i.e., more frequent disagreement on which applicants should be approved — indicates that while data sharing enables each party to more precisely identify the applicants who are most profitable according to their own objective, it leads them to favor different subsets of applicants. This introduces a potential tension: although sharing data may improve the decision-maker's ability to target applicants and improve its own payoff, it may simultaneously reduce the other party's payoff. We return to this incentive compatibility issue in later analysis.

The left panel of Figure 4 shows the alignment rate across a range of profit-sharing parameter s. Across all values of s, data sharing leads to a lower alignment rate. For example, under the observed contract ($s = 0$), alignment falls from approximately 53.8% without data sharing to 50.8% with data sharing. The gap narrows as s increases and incentives become more naturally aligned through profit-sharing, since a higher s gives the retailer a larger stake in the card's profitability and makes its approval preferences more similar to the bank's. This pattern suggests that richer data allows each party to optimize more precisely according to its own objective—but this sharper targeting may lead to more divergence, not less, in who they consider profitable.

The right panel of Figure 4 sheds further light on the nature of this divergence by decomposing the increase in misalignment into two components: (1) applicants whom the bank would approve but the retailer would not, and (2) applicants whom the retailer would approve but the bank would not. We find that the share of applicants favored by the retailer but not the bank decreases slightly with data sharing, and this decline remains relatively stable across all values of s. In contrast, the share of applicants favored by the bank but not the retailer increases with data sharing, though the magnitude of this increase diminishes as s rises. This pattern reflects the fact that as the retailer receives a larger share of card profit, it becomes more willing to approve applicants who are primarily valuable to the bank. The gain from card profit offsets more of the retailer's initial reluctance, bringing its approval preferences closer to the bank's.

Figure 4: Change in (Mis-)Alignment Rate due to Data Sharing

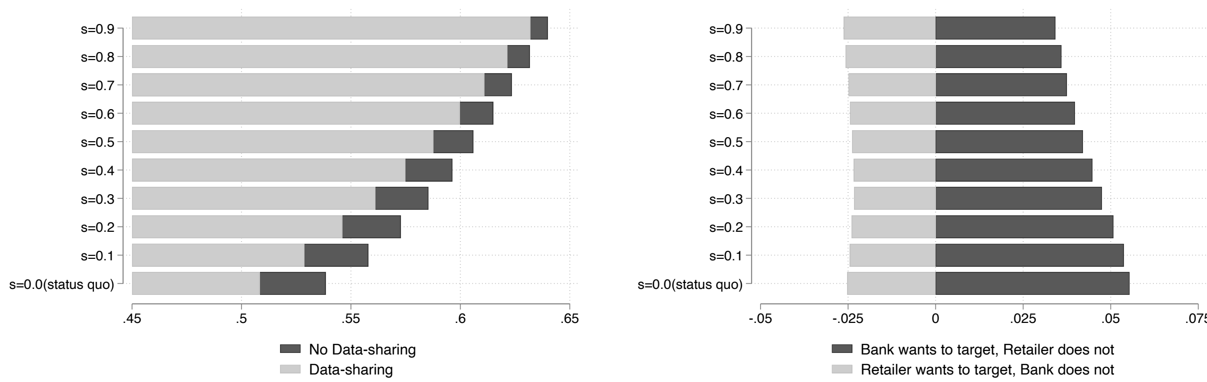

5.2 Welfare Implications and Sources of Misalignment

The previous analysis shows that data sharing can increase incentive misalignment. This section investigates how data sharing affects the realized payoff for each party. As described in Section 2, the value of data for the bank (ΔV^B) and the retailer (ΔV^R) is defined as the difference in realized payoff resulting from their approval decisions under different levels of data availability. For each contract, the benchmark is the realized payoff under a scenario in which decisions are based solely on bank data. To better understand how the decision maker's objective influences these outcomes, we consider two cases: (i) the factual setting in which the bank controls approval decisions to optimize its own objective, and (ii) a theoretical benchmark in which approval decisions are made to maximize the retailer's objective. This represents an extreme case where the retailer's payoff is fully prioritized, providing a useful contrast to the bank-centric scenario.[14] It is relevant for understanding contexts in which the retailer has significant influence over targeting decisions—such as when the retailer leads program design or dictate approval criteria. This setup also allows us to assess the value of the same additional retail data (X^R) by comparing outcomes in which decisions are optimized either using bank data alone (X^B) or the combined data (X^B, X^R), but comparing both under the retailer's objective. In both cases, we compute the value of data by comparing realized payoffs with and without data sharing, and examine how it varies across the full range of the profit-sharing parameter s.

Figure 5: Change in Uplift due to Data Sharing

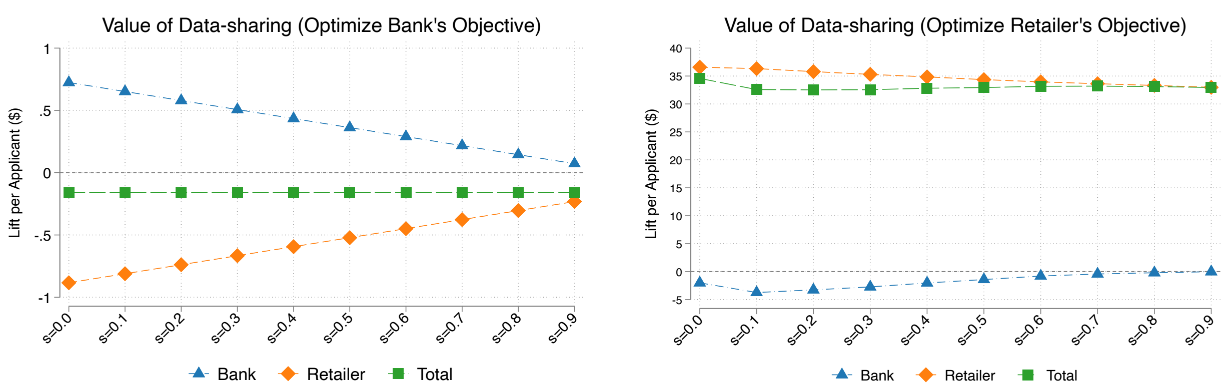

Figure 5 presents the effects of data sharing on realized value — showing the lift in realized value per applicant under different profit-sharing contracts. The left panel shows outcomes when the bank optimizes its own objective, while the right panel shows outcomes when optimizing the retailer's objective. For each scenario, we display the change in realized payoffs for the bank, the retailer, and their combined total value.

[14]Both regulatory and industry sources support our claim that banks consider partner merchants' objectives when making approval decisions. The Federal Deposit Insurance Corporation (FDIC), *Credit Card Lending* Guidance, acknowledges that banks may adjust underwriting standards to accommodate retail partners. The Consumer Financial Protection Bureau (CFPB), *Issue Spotlight: The High Cost of Retail Credit Cards*, highlights that approval decisions can reflect a retailer's interest in long-term customer relationships. Industry evidence also shows these partner-focused priorities can sometimes lead to issuer losses, as documented in: `https://www.nerdwallet.com/article/credit-cards/big-companies-shuttering-branded-credit-cards`, `https://www.americanbanker.com/payments/news/how-banks-co-brand-card-relationships-go-sour`.

In the left panel, when the bank optimizes its own payoff, data sharing yields modest gains for the bank but results in losses for the retailer. This pattern holds across all values of the profit-sharing parameter s. As s increases, the bank's gains diminish, while the retailer's losses also shrink. This is because a larger share of the gain is transferred to the retailer through the profit sharing arrangement. However, the gains from increased card profit do not compensate for the losses on the retail side, resulting in a net decline in total payoff across all s. This pattern reflects that, with access to more data, the bank can better target applicants who maximize its own profit, even if doing so harms the retailer.

A similar tension arises when decisions are made to optimize the retailer's payoff, shown in the right panel. In this case, data sharing consistently improves the retailer's realized value per applicant, while reducing the bank's payoff. As s increases, the bank's loss becomes smaller, because the retailer's and bank's payoffs are more closely aligned under this arrangement. Unlike the bank-optimization case, total realized value increases under all profit-sharing contracts when decisions are made in the retailer's interest.

Two insights emerge from these results. First, data sharing tends to improve outcomes aligned with the decision maker's objective, but often at the expense of the other objective. Data sharing does not necessarily improve both payoffs. Moreover, the distribution of gains and losses varies substantially with the contract structure.

Second, the value of data differs dramatically depending on which objective is being optimized. When retail data is used to optimize the bank's objective, it yields a modest gain of just \$0.72 per applicant. However, when the same data is used to optimize for the retailer, its value jumps to \$36.59 per applicant—a more than 50-fold increase. This stark contrast arises because the retailer has more variation in payoff, as shown in earlier comparisons (Table 4): the mean retail profit lift is −\$316.83 for negatively affected applicants versus −\$473.28 for positively affected ones, highlights a substantial scope for improvement. This discrepancy underscores that the value of data is not intrinsic in collaborative settings. Instead, it is shaped by the contract structure, which defines the decision maker's objective and, in turn, how the data is used. The resulting decisions affect the payoffs of both parties, including those without decision authority. As a result, the same data can yield vastly different economic outcomes depending on how it informs decisions within a given contractual arrangement.

To better understand the underlying drivers of the observed value changes, we decompose the realized value uplift from data sharing into its key components: interchange fees from partner retail, interchange fees from outside merchants, retail profit, discount usage, and default losses. Figure 6 presents the component-level changes when additional data is used to optimize the bank's payoff, and Figure 7 shows the corresponding changes when the data is used to optimize the retailer's payoff.

Figure 6: Breakdown of Payoff Change due to Data Sharing

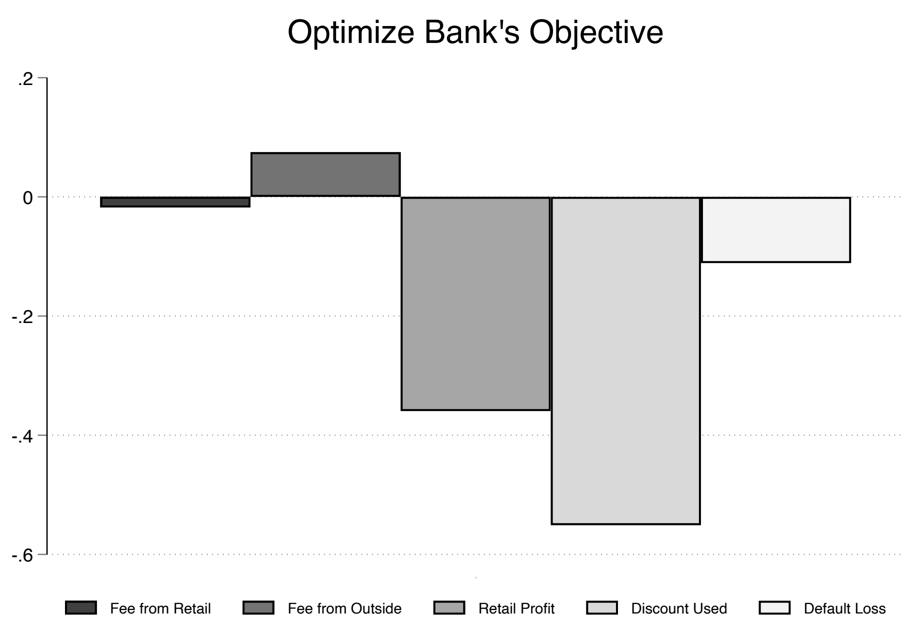

Notes: This bar plot decomposes the changes in the components of the bank's and retailer's payoff functions when additional data is used to optimized the bank's payoff. Since varying the profit-sharing parameter $s \in [0, 1]$ does not affect the bank's approval decisions, we report the component-level changes at $s = 0$ for simplicity. The vertical axis represents the change in value($) per applicant.

Changes in credit card discount usage and interchange fees from outside merchants highlight a key source of misalignment between the bank and the retailer. From the bank's perspective, the ideal applicants are those who adopt the card and use it primarily outside the partner retailer.

Figure 7: Breakdown of Payoff Change due to Data Sharing

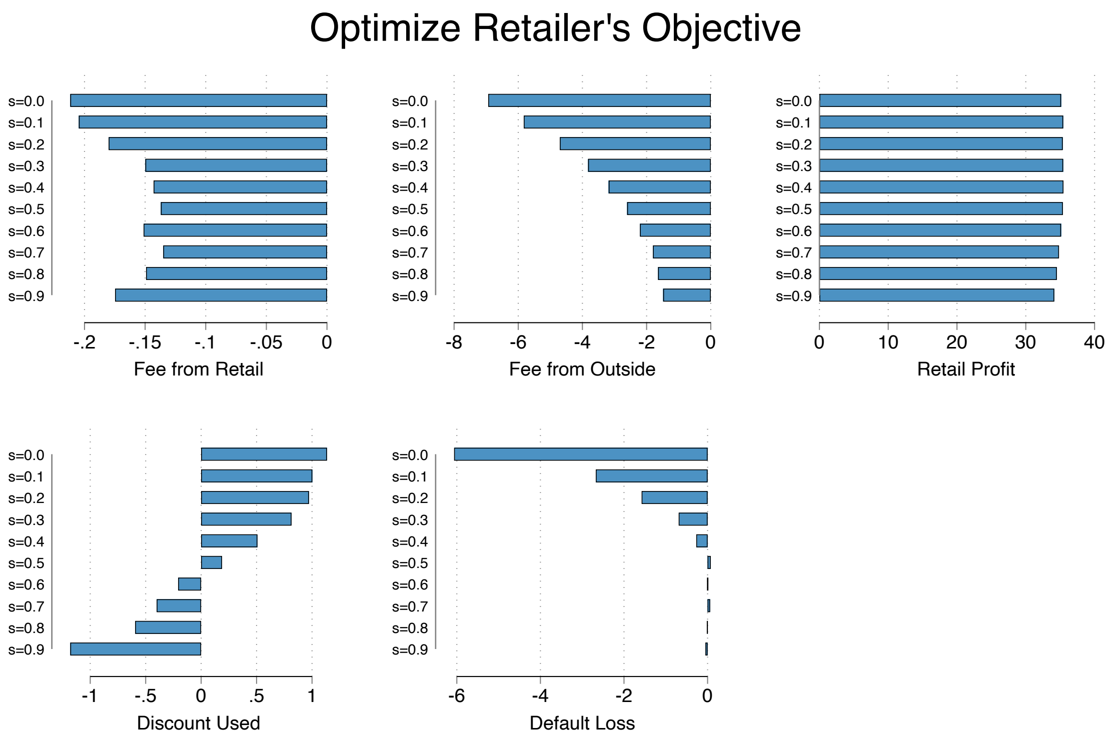

Notes: These bar plots decompose the changes in the components of the bank's and retailer's payoff functions when additional data is used to optimize the bank's payoff. The horizontal axis represents the change in value ($) per applicant.

These customers generate positive interchange revenue without incurring the cost of redeeming credit card discounts at the retailer.

Figure 6 reflects this strategy. When the bank controls approval decisions, discount usage decreases by $0.55, accounting for a substantial 76% of the total increase in the bank's payoff. At the same time, interchange revenue from outside merchants rises, while fees collected from transactions at the retailer decline. This pattern confirms that the bank prefers applicants who allocate their spending away from the partner retailer.

In contrast, Figure 7 shows that when the retailer's objective is being optimized, credit card discount utilization increases by $1.13 under the observed contract ($s = 0$) and interchange fees from outside merchants fall by $6.94. This indicates that optimizing the retailer's objective generates an opposite strategy, prioritizing customers who are likely to shop in-store, even if it means higher discount-related costs to the bank.

However, as the profit-sharing parameter s increases, the retailer's incentives begin to shift. At a high value of $s = 0.9$, discount usage declines by \$1.18 relative to the no data sharing case, indicating that the cost of providing discounts now plays a more significant role in the retailer's decision-making. Nevertheless, even at high levels of profit-sharing, the retailer's objective continues to favor customers who concentrate their spending in their stores: interchange revenue from outside merchants remains negative across all values of s, while changes in retail profit remain consistently positive. This suggests that even when the retailer captures a substantial share of card-based profit, its priorities remain in the retail profits — reflecting the relatively greater value it derives from retail transactions than from credit card activity in this setting. In the Appendix A.4, we further examine which customer segments drive these shifts. Specifically, we analyze how data sharing changes the distribution of applicants across the four targeting regions defined in Figure 1, and how these shifts translate into changes in each party's payoff and its underlying components.

In summary, our analysis reveals that while data sharing can improve the decision maker's payoff, it often does so at the expense of the non-decision-making party, and in some cases, even reduces total payoff. In the co-branded credit card context, this tension arises from a shift in targeting priorities: the bank favors applicants who generate interchange revenue outside the retail channel and avoid discount costs, whereas the retailer prioritizes applicants who drive in-store profit, even if that entails higher discount cost to the bank. More importantly, these findings underscore that the value of data is not intrinsic — the same information can produce vastly different economic outcomes depending on the decision maker's objective, which is determined by the contract structure.

In the next section, we examine contract settings that incorporate incentive compatibility — where the decision maker must ensure that the other party is not made worse off by the use of additional data. We consider two cases: one in which the decision maker must ensure that the other party is no worse off, and another in which a linear contract aligns incentives by maximizing the joint payoff and splitting it between the parties.

5.3 Value of Data Under Constraints

Our previous results show that data sharing does not always lead to mutually beneficial outcomes — particularly when one party optimizes based solely on its own payoff. To ensure that data sharing is not only effective but also incentive compatible, we now impose a participation constraint that guarantees the non-decision-making party is not worse off than under the no-data-sharing scenario. This constraint reflects a realistic requirement for sustainable data-sharing partnerships: unless both parties benefit or are at least no worse off, voluntary cooperation is unlikely. We incorporate this constraint into the decision-maker's policy optimization and evaluate how it alters the realized value from data sharing.

An alternative arrangement that inherently guarantees mutual benefit is a linear contract, in which approval decisions are made to maximize joint payoff, and the resulting gains are split between parties according to a pre-specified share. This setup mirrors a single-firm environment, where more data consistently enhances total value. In this case, both the retailer and the bank benefit from access to additional data. However, the extent to which each party benefits depends on the specified split parameter. Notably, maximizing joint payoff under a linear contract does not necessarily lead to the highest realized value of data, as the baseline payoff (i.e., without data sharing) also varies across contract forms.

Table 6 presents the impact of imposing participation constraints and implementing linear contracts on the value generated through data sharing. To provide a consistent comparison framework, we use the observed contract ($s = 0$) as our baseline reference point. For the linear contract scenario, we present results based on an equal 50-50 profit split between the bank and retailer.[15]

Our findings reveal that imposing a participation constraint (ensuring the objective not being optimized is not worse off) has minimal impact on the decision-maker's potential gains while

[15] For results under alternative values of s or alternative share splits in the linear contract, see Section A.5.

Table 6: Change in Payoff (Δ Value) by Objective and Constraint

Objective		Bank	Retailer	Total
Bank Centric (s = 0)	Δ Value (No constraint)	0.72	-0.88	-0.16
	Δ Value (with constraint)	0.72	0.01	0.73
Retailer Centric (s = 0)	Δ Value (No constraint)	-2.01	36.59	34.58
	Δ Value (with constraint)	0.01	36.59	36.60
Linear Contract (Equal Split)	Δ Value	16.39	16.39	32.77

significantly improving outcomes for the other party. When the bank optimizes with this constraint, its value from data sharing remains almost unchanged at 0.72 dollars per applicant, while the retailer's position improves from a loss of 0.88 to a slight gain of 0.01 dollars. With this constraint, the total payoff increases from $122.72 (no data sharing) to $123.45. Similarly, when the retailer optimizes with the constraint, its substantial gain of 36.59 dollars is preserved while the bank's position improves from a loss of 2.01 to a slight gain of 0.01 dollars, driving the total payoff increases from $189.76 to $226.36. The linear contract, which optimizes joint payoff rather than either party's individual objective, produces a balanced value distribution with both parties gaining 16.39 dollars per applicant, for a total value of 32.77 dollars. This total represents a middle ground between the modest joint value created in the bank-centric case (+0.73) and the substantial value in the retailer-centric case (+36.60). These variations in outcomes strongly reinforce our central finding that data's value is not intrinsic but heavily dependent on contractual structures. The same dataset can generate dramatically different economic returns depending on the decision rights, objective functions, and sharing mechanisms established in the contractual arrangement.

To summarize our findings in this section, Figure 8 provides a direct comparison of data's value across different contract structures, highlighting the dramatic variation in outcomes depending on contractual arrangements. These results demonstrate that the same retail transaction dataset can create significantly different economic value depending on the contractual structure governing its use. Data sharing doesn't automatically benefit all parties; rather, careful contract design is essential to align incentives and unlock the full potential of shared information resources.

Figure 8: Data Value across Contracts

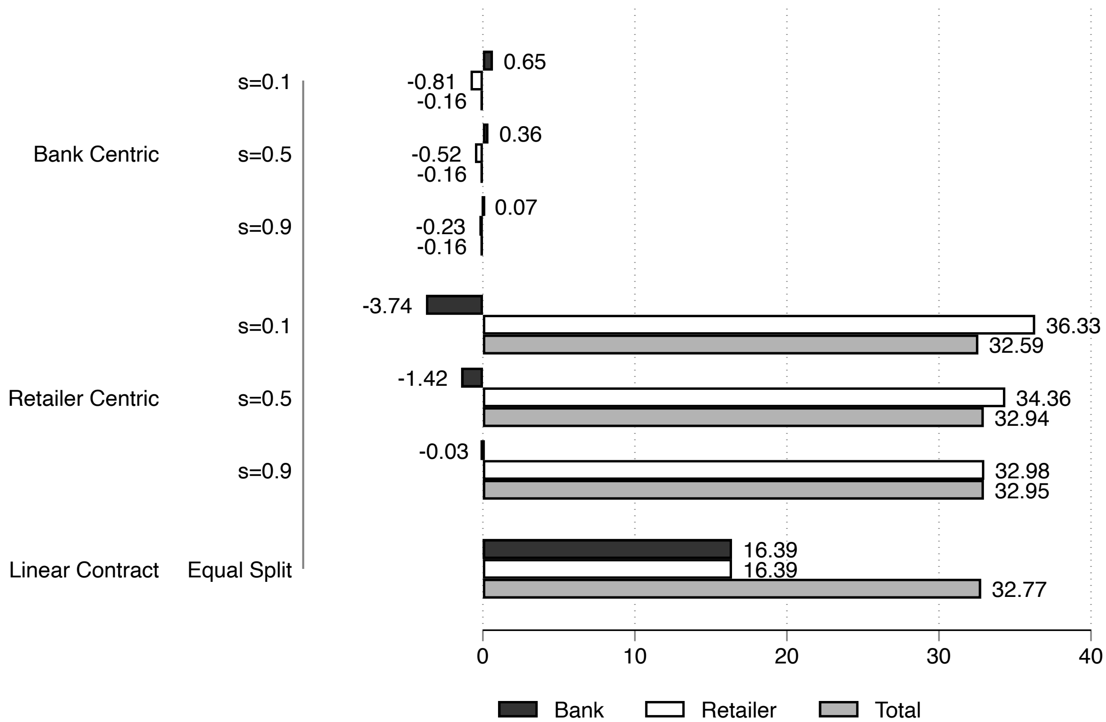

6 Concluding Remarks

This paper investigates how data sharing creates - or fails to create - value in multi-party marketing partnerships. As firms increasingly collaborate by leveraging each other's data, understanding the value of data in these partnerships and how contractual arrangements shape its value becomes essential. We ask: Does shared data lead to mutual gains, and how do contract structures influence the distribution of those gains?

To answer this question, we develop a theoretical framework that formalizes data valuation in collaborative marketing settings. The framework shows that while additional data helps the decision-maker make better decisions, this may leave the non-decision-maker (who is also the data owner) worse off, since their payoffs are not perfectly aligned. We then test these predictions using comprehensive individual-level data, where we observe each applicant's behavior on both the retail and banking sides. Using a doubly robust difference-in-differences approach, we estimate the causal effects of credit card adoption on firm outcomes.

Our empirical results reveals tension between two parties: customers who are profitable for the bank often reduce the retailer's profit, and vice versa, highlighting an incentive misalignment. This misalignment arises because the bank prefers applicants who primarily use the card outside the partner retailer — generating interchange revenue without triggering discount costs — whereas the retailer benefits most from customers who concentrate spending in-store, even though this increases costs for the bank. In our counterfactual analysis, we show that when the bank uses retail data to optimize approvals based on its own objectives, it gains (+0.72 local dollars per applicant), but the retailer loses (−0.88), as the data helps the bank identify customers aligned with its goals but not the retailer's. We further demonstrate that imposing a participation constraint can ensure incentive compatibility: under this arrangement, data sharing yields positive but modest joint gains (+0.73). By contrast, when the same data is used under a linear contract structure that aligns incentives and splits profits, its value rises dramatically — to +32.77 per applicant — underscoring that the realized value of shared data is not intrinsic to the dataset, but shaped by the incentive structure in this collaboration.

These results demonstrate that the economic value of data is not an intrinsic property of the dataset, but a function of how decisions are made and how gains are shared. In collaborative marketing settings, identical data can yield vastly different outcomes depending on the institutional context. This insight has important implications for how firms structure data-sharing agreements, and how policymakers evaluate the welfare implications of interoperability and data portability rules.

Future research could explore how these dynamics unfold in other settings where data is shared across firms with distinct objectives — such as retail media networks, platform ecosystems, or third-party ad attribution. Extensions could also examine how contract terms interact with data quality, enforcement frictions, or market power. More broadly, advancing our understanding of how institutional design shapes the returns to data remains a critical agenda for marketing, strategy, and digital policy.

References

Bergemann, D., A. Bonatti, and A. Smolin (2018). The design and price of information. *American economic review 108*(1), 1–48.

Bergemann, D. and S. Morris (2019). Information design: A unified perspective. *Journal of Economic Literature 57*(1), 44–95.

Bijmolt, T. H., M. Dorotic, P. C. Verhoef, et al. (2011). Loyalty programs: Generalizations on their adoption, effectiveness and design. *Foundations and Trends® in Marketing 5*(4), 197–258.

Brynjolfsson, E., L. M. Hitt, and H. H. Kim (2011). Strength in numbers: How does data-driven decisionmaking affect firm performance? *Available at SSRN 1819486*.

Callaway, B. and P. H. Sant'Anna (2021). Difference-in-differences with multiple time periods. *Journal of econometrics 225*(2), 200–230.

Chernozhukov, V., D. Chetverikov, M. Demirer, E. Duflo, C. Hansen, W. Newey, and J. Robins (2018). Double/debiased machine learning for treatment and structural parameters.

De Wulf, K., G. Odekerken-Schröder, and D. Iacobucci (2001). Investments in consumer relationships: A cross-country and cross-industry exploration. *Journal of marketing 65*(4), 33–50.

Einav, L. and J. Levin (2014). Economics in the age of big data. *Science 346*(6210), 1243089.

Frick, T. W., R. Belo, and R. Telang (2023). Incentive misalignments in programmatic advertising: Evidence from a randomized field experiment. *Management Science 69*(3), 1665–1686.

Galperti, S., A. Levkun, and J. Perego (2024). The value of data records. *Review of Economic Studies 91*(2), 1007–1038.

Goldfarb, A. and C. Tucker (2019). Digital economics. *Journal of economic literature 57*(1), 3–43.

Gopalakrishnan, A., Z. Jiang, Y. Nevskaya, and R. Thomadsen (2021). Can non-tiered customer loyalty programs be profitable? *Marketing Science 40*(3), 508–526.

Hartmann, W. R. and V. B. Viard (2008). Do frequency reward programs create switching costs? a dynamic structural analysis of demand in a reward program. *Quantitative Marketing and Economics 6*, 109–137.

Iyengar, R., Y.-H. Park, and Q. Yu (2022). The impact of subscription programs on customer purchases. *Journal of Marketing Research 59*(6), 1101–1119.

Johnson, G. and R. A. Lewis (2015, Oct). Cost per incremental action: Efficient pricing of advertising. *Simon Business School Working Paper No. FR 15-29*.

Lal, R. and D. E. Bell (2003). The impact of frequent shopper programs in grocery retailing. *Quantitative marketing and economics 1*, 179–202.

Lambrecht, A. and C. Tucker (2013). When does retargeting work? information specificity in online advertising. *Journal of Marketing research 50*(5), 561–576.

Lee, J. Y., J. Yang, and E. T. Anderson (2024). Using grocery data for credit decisions. *Management Science*.

Leenheer, J., H. J. Van Heerde, T. H. Bijmolt, and A. Smidts (2007). Do loyalty programs really enhance behavioral loyalty? an empirical analysis accounting for self-selecting members. *International journal of research in marketing 24*(1), 31–47.

Lewis, M. (2004). The influence of loyalty programs and short-term promotions on customer retention. *Journal of marketing research 41*(3), 281–292.

Lewis, R. and J. Wong (2022). Incrementality bidding and attribution. *arXiv preprint arXiv:2208.12809*.

Liu, Y. and R. Yang (2009). Competing loyalty programs: Impact of market saturation, market share, and category expandability. *Journal of Marketing 73*(1), 93–108.

Manski, C. F. and J. V. Pepper (2018). How do right-to-carry laws affect crime rates? coping with ambiguity using bounded-variation assumptions. *Review of Economics and Statistics 100*(2), 232–244.

McCall, M. and C. Voorhees (2010). The drivers of loyalty program success: An organizing framework and research agenda. *Cornell Hospitality Quarterly 51*(1), 35–52.

Rambachan, A. and J. Roth (2023). A more credible approach to parallel trends. *Review of Economic Studies 90*(5), 2555–2591.

Rossi, P. E., R. E. McCulloch, and G. M. Allenby (1996). The value of purchase history data in target marketing. *Marketing Science 15*(4), 321–340.

Sant'Anna, P. H. and J. Zhao (2020). Doubly robust difference-in-differences estimators. *Journal of econometrics 219*(1), 101–122.

Sharp, B. and A. Sharp (1997). Loyalty programs and their impact on repeat-purchase loyalty patterns. *International journal of Research in Marketing 14*(5), 473–486.

Taylor, W. and B. Hollenbeck (2021). Leveraging loyalty programs using competitor based targeting. *Quantitative Marketing and Economics 19*(3), 417–455.

Verhoef, P. C. (2003). Understanding the effect of customer relationship management efforts on customer retention and customer share development. *Journal of marketing 67*(4), 30–45.

Wernerfelt, N., A. Tuchman, B. T. Shapiro, and R. Moakler (2024). Estimating the value of offsite tracking data to advertisers: Evidence from meta. *Marketing Science*.

Wu, R., Y. Huang, and N. Li (2023). Platform information design and competitive price targeting. *Available at SSRN 4619584*.

Xu, J., X. Shao, J. Ma, K.-c. Lee, H. Qi, and Q. Lu (2016, Feb.). Lift-based bidding in ad selection. *Proceedings of the AAAI Conference on Artificial Intelligence 30*(1).

Zhao, N., A. Gopalakrishnan, and C. Narasimhan (2022). The impact of co-branded credit card adoption on customer loyalty. *Available at SSRN 3820361*.

Appendix

A.1 Randomization Check and Balance Check

As discussed earlier, the random assignment of applications to review agents, conditional on the risk score, introduces variation in approval decisions among applicants with similar risk profiles. This feature of the institutional setting makes us to expect overlapping risk score distributions between the treated (approved) and control (rejected) groups. Some applicants with a given risk score are approved, while others are rejected, depending on the reviewer they are assigned to.

This motivates a simple empirical check of the randomization/overlap assumption, which requires that treated and control groups have support over similar regions of the risk score. In this case, since the risk score is a primary driver of approval, we expect its distribution to overlap across groups. Figure A.1 shows the estimated densities of the risk score for approved and rejected applicants. The two distributions display substantial overlap, with no clear separation or mass concentrated at the extremes (0 or 1). This provides reassurance that the overlap assumption holds with respect to the risk score, supporting the plausibility of our assumption.

We now more formally assess the validity of the overlap assumption, which is central to identifying treatment effects. Intuitively, when the overlap assumption is violated, we lack comparable untreated (or treated) observations to reliably predict or impute counterfactual outcomes for some individuals. To evaluate this, we examine the distribution of estimated propensity scores for treated and control groups to ensure that both groups have sufficient representation across the common support. As we can see from Figure A.2, neither plot of each figure exhibits substantial probability mass near 0 or 1. Moreover, for each group, the two estimated densities largely overlap, with most of their respective mass concentrated in shared regions of the propensity score support. These patterns suggest that the overlap assumption holds in our setting, and there is no evidence of meaningful violations that would compromise identification.

Figure A.1: Overlap Check with Risk Scores

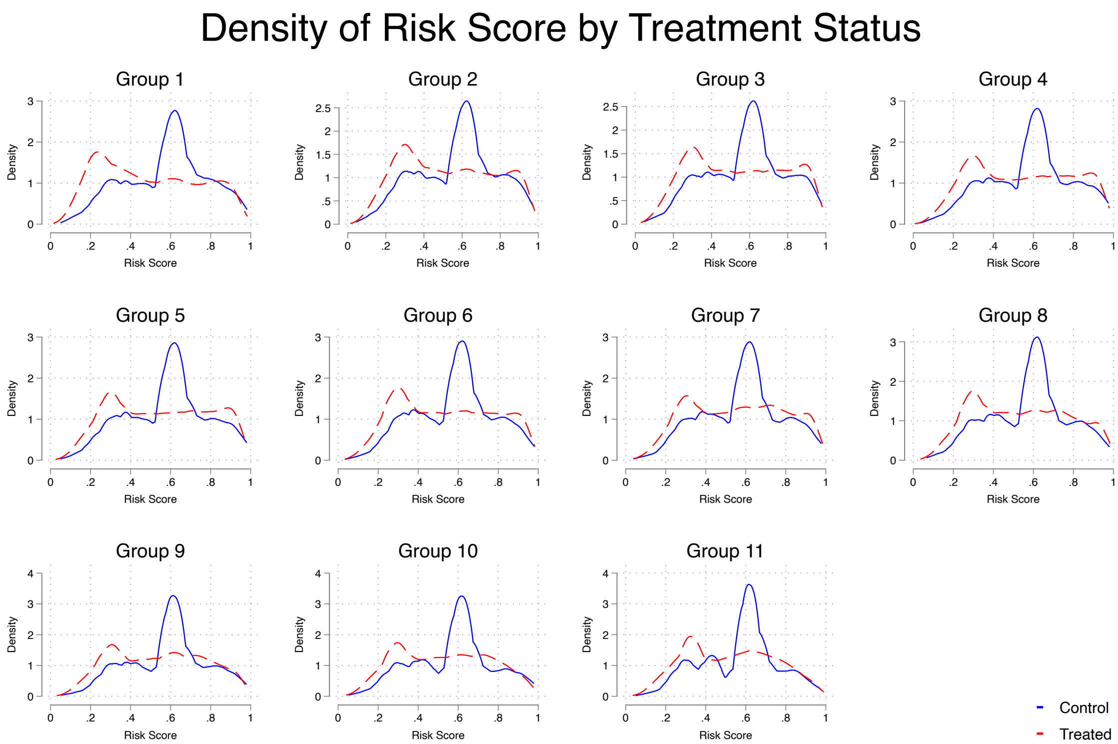

A.2 Conditional Parallel Trend — Sensitivity Analysis

Directly testing for parallel trends using pre-treatment data can introduce pre-testing bias, potentially invalidating subsequent inference. To address this concern, Rambachan and Roth (2023) propose a sensitivity analysis framework grounded in partial identification (Manski and Pepper, 2018), which avoids relying on strong assumptions about pre-trend equivalence. Instead of asserting that the parallel trend assumption holds exactly, their approach quantifies the robustness of treatment effect estimates to deviations from this assumption. We adopt their method to assess the sensitivity of our estimates to violations of the conditional parallel trends assumption, providing a more credible range of treatment effect estimates under a bounded departure from the identifying assumptions.

Specifically, we adopt the sensitivity analysis framework proposed by Rambachan and Roth (2023), which relaxes the conditional parallel trends assumption by imposing smoothness restrictions on potential deviations. The key idea is that any violation of parallel trends in the

Figure A.2: Overlap Check with Propensity Scores

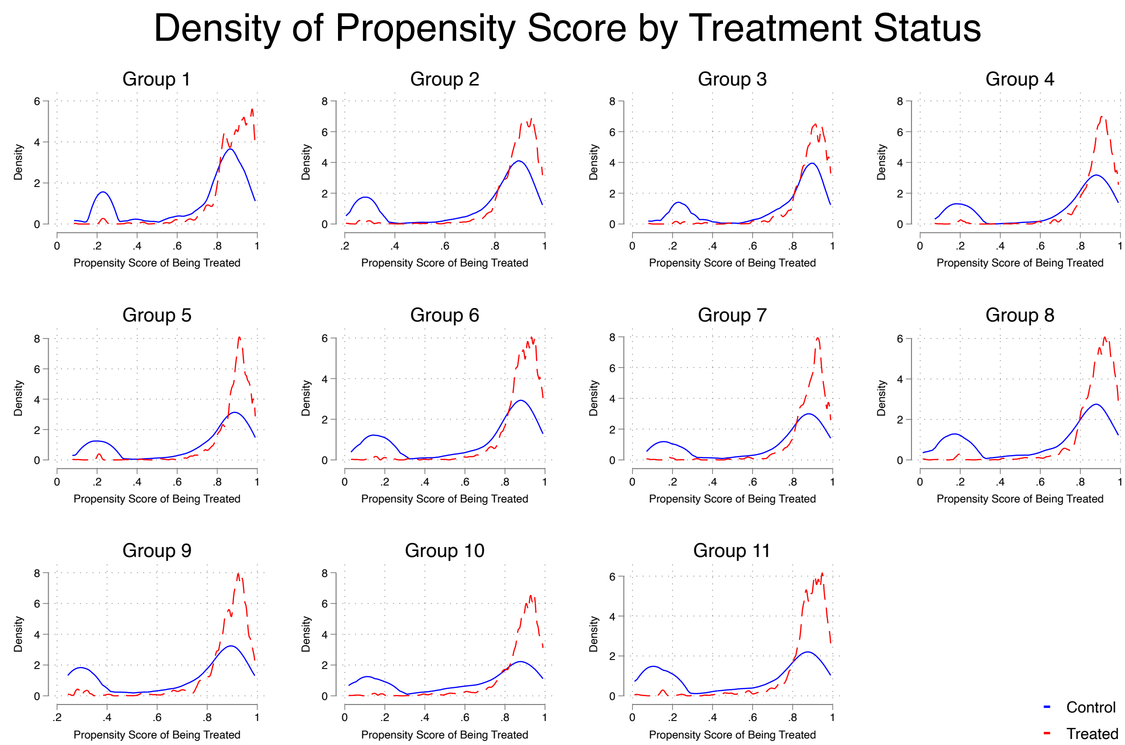

post-treatment period must evolve gradually, rather than arbitrarily. We assume that the slope of the pre-treatment trend can change by no more than a fixed amount M between consecutive post-treatment periods. When $M = 0$, the counterfactual difference in trends is constrained to follow a linear extrapolation of the pre-trend. Larger values of M allow for increasing degrees of non-linearity, thereby weakening the assumption and providing a more conservative estimate of the treatment effect. This approach enables us to quantify the robustness of our results to plausible forms of time-varying confounding.

Figure A.3 presents the results of our sensitivity analysis for three key outcomes: the number of store trips, retail profit, and total spending at the retailer. We focus on these outcomes because, for the remaining variables (such as discount usage, credit card spending, and default loss), the parallel trends assumption is mechanically satisfied by construction, as they are defined to be zero in absence of treatment. We find that the breakdown value — the smallest value of M at which we can no longer reject a null effect — is approximately $M \approx 0.01$. This implies that our estimated effect remains statistically significant unless we allow the slope of the counterfactual trend to deviate by more than 0.01 percentage points across consecutive periods from the linear extrapolation of the pre-trend. In contrast, the results for retail profit and total spending at the retailer are substantially more robust: even when allowing for deviations as large as 0.1 percentage points, the estimated effects remain statistically distinguishable from zero. This provides reassurance that our main findings for these two outcomes are not driven by small or smooth violations of the parallel trends assumption.

Figure A.3: Sensitivity Analysis (Smoothness Restriction)

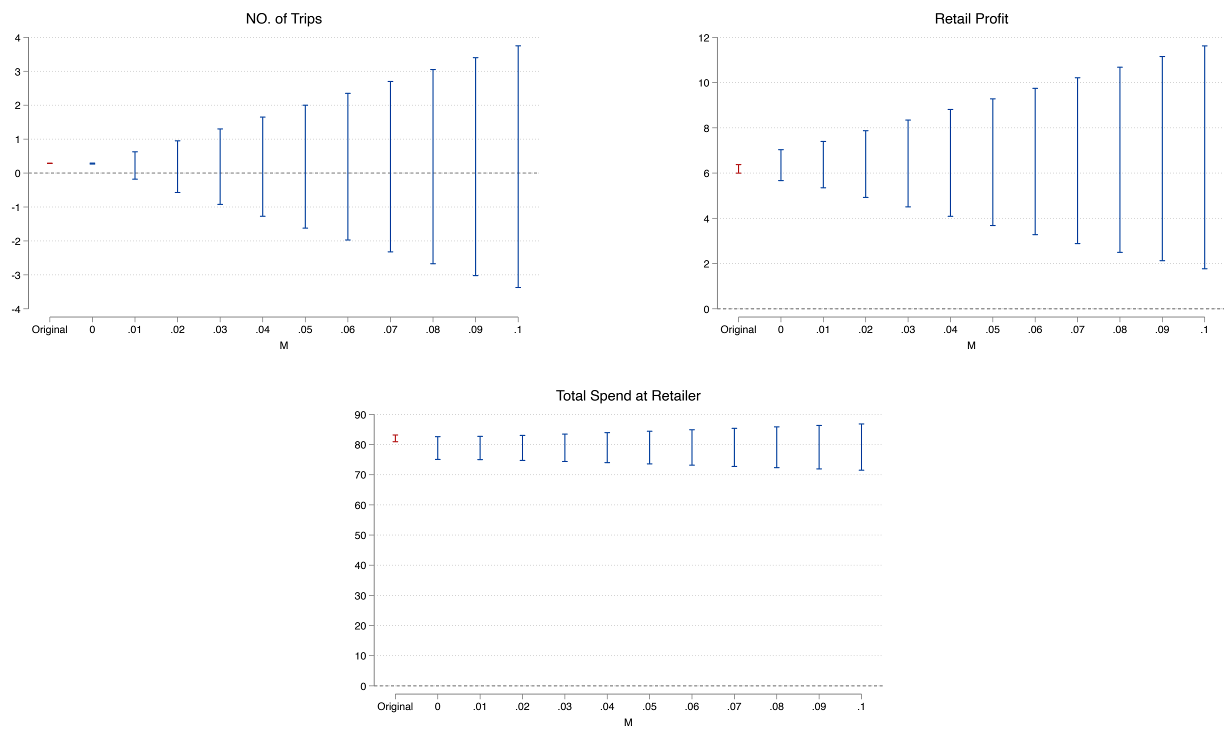

A.3 Lifetime Value across Different Parties

Figure A.4 presents two measures of causal effects on the lifetime value (LTV): ATT and average payoff per applicant (not ATE). The ATT captures the average causal impact of approval among those who were actually approved, providing a clean estimate of the value generated by the program conditional on being treated. In contrast, the average payoff per applicant reflects the realized average payoff across the entire pool of applicants, given the observed approval rule. The left panel shows ATTs for the lifetime value (LTV) generated by the co-branded credit card, disaggregated by the bank and the retailer, along with the total. The largest component, estimated at approximately 118.06 dollars, reflects the increase in cumulative retail profit per approved customer over the 13-month post-approval window. This sizable effect highlights the success of the co-branded credit card in driving higher retail spending and profitability. The corresponding ATT for the bank is negative, estimated at -37.61 dollars. This includes gains from interchange fees, cost of discounts, and losses from defaults. The negative payoff suggests that while the bank benefits from card issuance, the marginal financial return per customer is low, likely due to subsidy costs (e.g., retailer discounts) and default risk. The sum of both effects is approximately 80.45 dollars, with the vast majority accruing to the retailer. The right panel shows a similar pattern for the average payoff per applicant.

Figure A.4: Lifts of Life Time Value

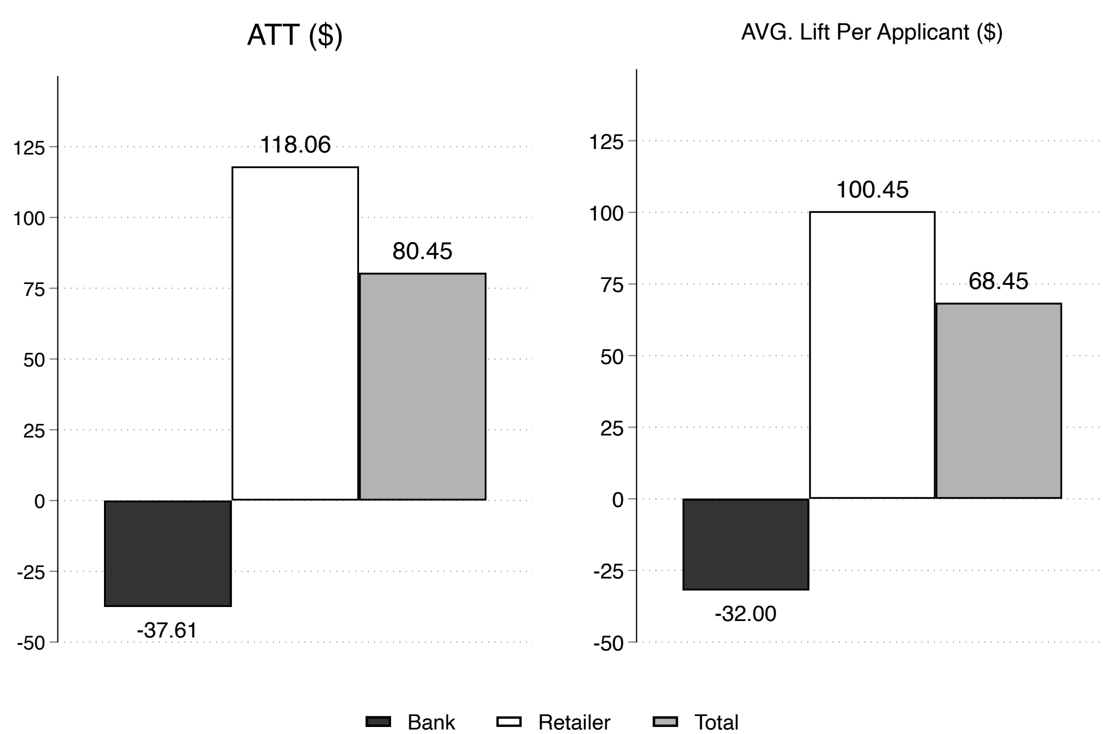

A.4 Alignment Region Changes and Corresponding Component Changes

In this section, we further unpack which segments of customers drive these component shifts. we analyze how data sharing changes the distribution of applicants across the four targeting regions defined in Figure 1. And then we analyze how these region changes affect each parties payoff and the component in the payoff. This approach allows us to connect the observed value changes to specific incentive alignment patterns, revealing how improved predictive signals reshape targeting behavior and business outcomes under different information and contractual environments.

Table A.1: Percentage of Target Region Change

	s=0.0 (%)	s=0.1 (%)	s=0.2 (%)	s=0.3 (%)	s=0.4 (%)	s=0.5 (%)	s=0.6 (%)	s=0.7 (%)	s=0.8 (%)	s=0.9 (%)
1->2	32.84	32.26	31.46	30.54	29.70	28.83	28.05	27.23	26.65	26.00
1->3	0.58	0.60	0.63	0.65	0.67	0.70	0.73	0.75	0.78	0.80
1->4	2.77	2.80	2.84	2.90	2.94	2.98	3.02	3.07	3.10	3.15
2->1	10.11	10.03	10.01	10.02	9.99	9.98	9.93	9.87	9.73	9.59
2->3	2.26	2.29	2.34	2.39	2.44	2.48	2.52	2.58	2.63	2.69
2->4	1.51	1.49	1.50	1.50	1.50	1.49	1.48	1.47	1.45	1.43
3->1	1.19	1.22	1.29	1.34	1.39	1.45	1.49	1.55	1.60	1.68
3->2	1.80	1.81	1.85	1.89	1.92	1.96	1.98	2.01	2.04	2.06
3->4	16.44	16.81	17.10	17.46	17.70	17.86	17.98	18.15	18.15	18.24
4->1	2.63	2.63	2.65	2.68	2.71	2.72	2.73	2.74	2.75	2.77
4->2	0.75	0.74	0.74	0.73	0.72	0.72	0.72	0.72	0.71	0.70
4->3	27.13	27.31	27.61	27.90	28.32	28.84	29.37	29.89	30.41	30.90

Recall that regions 1 through 4 correspond to combinations of approval incentives: region 1 represents mutual approval (both retailer and bank benefit), region 2 is bank-favored (only the bank benefits), region 3 is mutual rejection, and region 4 is retailer-favored (only the retailer benefits). Table A.1 reports the percentage of applicants who transition across these regions as data sharing is introduced under varying profit-sharing contracts. Each row denotes a transition from one region to another (e.g., “1->2” indicates applicants originally in region 1 are reclassified into region 2 after data sharing). These shifts reflect how access to retailer data alters the bank’s and retailer’s estimated treatment effects, and thus affects their joint targeting preferences. To illustrate how data sharing reshapes perceived applicant value and incentive alignment, we focus on three representative transitions: “1->2”, “1->3”, and “1->4”. (1) applicants in “1->2” continue to appear valuable to the bank but are now revealed to be unprofitable for the retailer. Importantly, this transition does not change the targeting decision if the bank remains the decision-maker, since the bank continues to approve these applicants both before and after data sharing. As a result, this reclassification has no impact on anyone's realized value — neither the bank nor the retailer experiences a change in payoff, as the decision and outcomes remain identical. This point is confirmed in Table A.2, which calculates the impact on average payoff per each applicant due to the region change; (2) applicants in the "1->3" category are initially perceived as beneficial to both the bank and the retailer but are reclassified — after data sharing — as unprofitable for either party. Unlike "1->2", this transition leads to a change in the approval decision: the bank, now recognizing negative value, will choose to reject these applicants. Table A.2 sheds light on the underlying behavioral patterns. If approved, these applicants tend to redeem card-linked discounts, but the purchases they make are concentrated in lower-margin products, resulting in minimal retail profit. Additionally, they exhibit limited credit card usage at outside merchants, reducing interchange fee potential for the bank. Importantly, the additional retailer data reveal that — absent the card and discount offers — these customers would likely still shop, and do so with higher-margin purchases, without drawing discount subsidies. Thus, rejecting these applicants after data sharing leads to an increase in retail profit by eliminating inefficient subsidized purchases. Simultaneously, the bank's payoff improves, as the savings from avoided discount costs more than offset the modest loss in interchange fees; (3) for "1->4" applicants, they are still desirable to the retailer but are newly identified as unprofitable to the bank. Since the bank now opts to reject them, retailer value declines sharply (−8.19 dollars) while the bank gains modestly (+0.25 dollars), driven by savings on discount costs, which outweigh the forgone interchange fee revenue from potential card usage. Other region transitions can be interpreted in a similar manner by applying the logic illustrated in our three focal cases: changes in realized value depend on whether the targeting decision changes, which components of value shift, and how these shifts are distributed across the two parties. Table A.3 reports the corresponding results when the retailer's objective is optimized. The interpretation of region transitions follows the same logic as in the bank-optimized case and is omitted here for brevity.

Table A.2: Change in Lift by Target Region Change (Optimize Bank's Objective, $s = 0$)

	Bank's Payoff	Retailer's Payoff	Total Payoff	Retail Profit	Card Spend (Outside)	Card Spend (Retailer)	Card Spend (Total)	Discount Used	Default Loss	Fee from Retail	Fee from Outside
1->2	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
1->3	0.03	0.54	0.57	0.57	-0.70	-0.44	-1.15	-0.04	-0.01	-0.01	-0.01
1->4	0.19	-6.58	-6.41	-6.38	-5.38	-2.56	-7.94	-0.28	-0.05	-0.05	-0.11
2->1	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
2->3	0.15	5.46	5.59	5.59	-2.91	-1.85	-4.77	-0.18	-0.04	-0.04	-0.06
2->4	0.21	-1.23	-1.02	-1.02	-2.12	-1.04	-3.16	-0.24	-0.03	-0.02	-0.04
3->1	0.03	0.82	0.86	0.82	2.79	0.66	3.45	0.03	0.00	0.01	0.06
3->2	0.04	-3.04	-2.98	-3.04	3.88	1.36	5.24	0.05	0.01	0.03	0.08
3->4	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
4->1	0.06	3.50	3.57	3.47	6.75	2.41	9.16	0.10	0.01	0.05	0.13
4->2	0.02	-0.36	-0.34	-0.36	1.45	0.57	2.02	0.02	0.00	0.01	0.03
4->3	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Total	0.72	-0.88	-0.16	-0.36	3.75	-0.89	2.86	-0.55	-0.11	-0.02	0.07

Table A.3: Change in Lift by Target Region Change (Optimize Retailer's Objective, $s = 0$)

	Bank's Payoff	Retailer's Payoff	Total Payoff	Retail Profit	Card Spend (Outside)	Card Spend (Retailer)	Card Spend (Total)	Discount Used	Default Loss	Fee from Retail	Fee from Outside
1->2	-9.95	8.61	-1.55	7.95	-501.32	-50.06	-551.38	-0.84	-0.03	-1.00	-10.03
1->3	0.03	0.54	0.57	0.57	-0.70	-0.44	-1.15	-0.04	-0.01	-0.01	-0.01
1->4	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
2->1	3.24	5.73	9.03	6.17	161.52	36.56	198.07	0.65	0.00	0.73	3.23
2->3	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
2->4	-0.21	1.23	1.02	1.02	2.12	1.04	3.16	0.24	0.03	0.02	0.04
3->1	0.03	0.82	0.86	0.82	2.79	0.66	3.45	0.03	0.00	0.01	0.06
3->2	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
3->4	-6.16	9.65	3.59	7.66	27.57	15.82	43.39	2.46	4.46	0.32	0.55
4->1	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
4->2	-0.02	0.36	0.34	0.36	-1.45	-0.57	-2.02	-0.02	-0.00	-0.01	-0.03
4->3	11.03	9.65	20.50	10.59	-37.32	-13.59	-50.91	-1.35	-10.53	-0.27	-0.75
Total	-2.01	36.59	34.37	35.14	-346.80	-10.59	-357.39	1.13	-6.07	-0.21	-6.94

Table A.4: Change in Lift by Target Region Change (Optimize Bank's Objective, $s = 0.9$)

	Bank's Payoff	Retailer's Payoff	Total Payoff	Retail Profit	Card Spend (Outside)	Card Spend (Retailer)	Card Spend (Total)	Discount Used	Default Loss	Fee from Retail	Fee from Outside
1->2	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
1->3	0.01	0.62	0.62	0.61	-0.99	-0.56	-1.56	-0.06	-0.02	-0.01	-0.02
1->4	0.02	-6.45	-6.46	-6.41	-5.52	-2.60	-8.13	-0.28	-0.04	-0.05	-0.11
2->1	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
2->3	0.02	5.58	5.58	5.57	-3.04	-1.94	-4.97	-0.21	-0.05	-0.04	-0.06
2->4	0.02	-1.02	-1.01	-1.00	-1.57	-0.79	-2.36	-0.19	-0.01	-0.02	-0.03
3->1	0.00	0.92	0.93	0.87	3.85	0.90	4.75	0.04	0.01	0.02	0.08
3->2	0.00	-3.02	-3.01	-3.06	3.66	1.36	5.02	0.05	0.01	0.03	0.07
3->4	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
4->1	0.01	3.48	3.51	3.41	6.24	2.28	8.52	0.09	0.01	0.05	0.12
4->2	0.00	-0.33	-0.32	-0.34	1.12	0.47	1.59	0.01	0.00	0.01	0.02
4->3	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Total	0.07	-0.23	-0.16	-0.36	3.75	-0.89	2.86	-0.55	-0.11	-0.02	0.07

Table A.5: Change in Lift by Target Region Change (Optimize Retailer's Objective, $s = 0.9$)

	Bank's Payoff	Retailer's Payoff	Total Payoff	Retail Profit	Card Spend (Outside)	Card Spend (Retailer)	Card Spend (Total)	Discount Used	Default Loss	Fee from Retail	Fee from Outside
1->2	-0.35	5.99	5.49	8.69	-178.12	-34.68	-212.80	-0.55	-0.02	-0.69	-3.56
1->3	0.01	0.62	0.62	0.61	-0.99	-0.56	-1.56	-0.06	-0.02	-0.01	-0.02
1->4	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
2->1	0.21	5.20	5.46	3.64	106.05	29.91	135.96	0.55	0.00	0.60	2.12
2->3	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
2->4	-0.02	1.02	1.01	1.00	1.57	0.79	2.36	0.19	0.01	0.02	0.03
3->1	0.00	0.92	0.93	0.87	3.85	0.90	4.75	0.04	0.01	0.02	0.08
3->2	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
3->4	-0.33	8.81	8.59	9.26	28.36	16.70	45.05	3.00	1.07	0.33	0.57
4->1	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
4->2	-0.00	0.33	0.32	0.34	-1.12	-0.47	-1.59	-0.01	-0.00	-0.01	-0.02
4->3	0.45	10.10	10.37	9.73	-33.71	-21.32	-55.04	-4.33	-1.11	-0.43	-0.67
Total	-0.03	32.98	32.80	34.14	-74.12	-8.74	-82.86	-1.18	-0.05	-0.17	-1.48

A.5 Value of Data: Participation Constraint and Linear Contract

Figure A.5 illustrates the value of data sharing across different stakeholders. The left panel demonstrates the change in value per applicant under a bank-centric decision-making scenario, while the right panel shows these changes under a retailer-centric decision-making scenario. In both cases, we measure the incremental value to the bank's payoff (blue), the retailer's payoff (orange), and their combined total payoff (green) across varying profit-sharing arrangements. Figure A.6 displays the uplift in payoffs under a linear contract framework where decisions maximize joint total payoff rather than either party's individual objectives. The parameter s represents the proportion of joint payoff allocated to the retailer.

Figure A.5: Change in Uplift due to Data Sharing with Participation Constraints

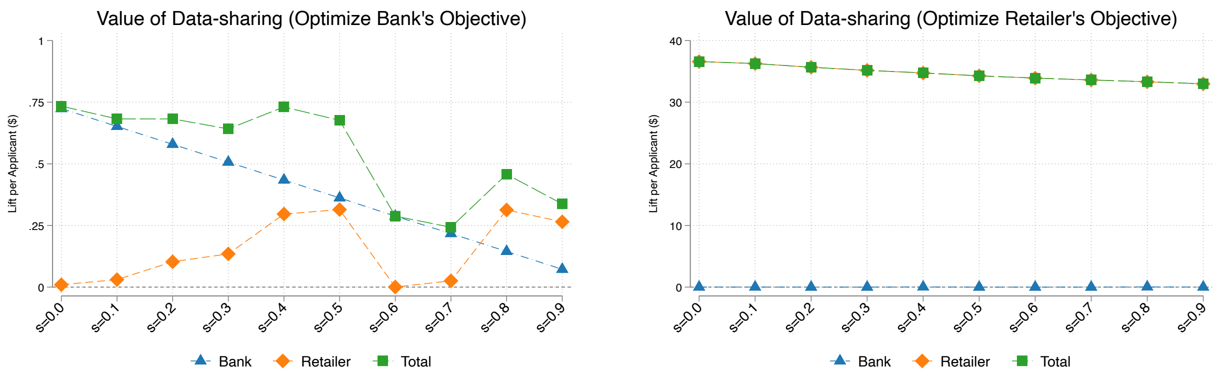

Figure A.6: Change in Uplift under Linear Contract