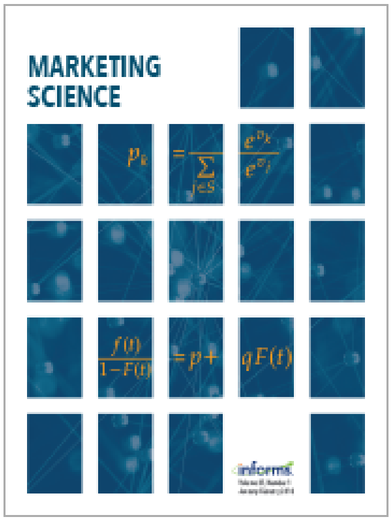

Marketing Science

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Returnless Partial Refunds: A Large-Scale Field Experiment in E-Commerce

Lu Fang, Chunhua Wu, Zhe Yuan, Kaifu Zhang

To cite this article:
Lu Fang, Chunhua Wu, Zhe Yuan, Kaifu Zhang (2026) Returnless Partial Refunds: A Large-Scale Field Experiment in E-Commerce. Marketing Science

Published online in Articles in Advance 23 Sep 2026

. https://doi.org/10.1287/mksc.2025.0618

informs
https://pubsonline.informs.org/journal/mksc

MARKETING SCIENCE
Articles in Advance, pp. 1–16
ISSN 0732-2399 (print), ISSN 1526-548X (online)

Returnless Partial Refunds: A Large-Scale Field Experiment in E-Commerce

Lu Fang,[a,b] **Chunhua Wu,**[c] **Zhe Yuan,**[d,*] **Kaifu Zhang**[e]

[a] School of Management, Zhejiang University of Finance and Economics, Hangzhou 310018, China; [b] Laboratory for Human-AI Collaboration and Digital Intelligence Innovation, Zhejiang University, Hangzhou 310058, China; [c] UBC Sauder School of Business, The University of British Columbia, Vancouver, British Columbia V6T 1Z2, Canada; [d] Center for Research of Private Economy, School of Economics, Future Regional Development Laboratory, and Academy of Financial Research, Zhejiang University, Hangzhou 310058, China; [e] Partner Company

*Corresponding author

fl.fanglu.fl@gmail.com (L. Fang), chunhua.wu@sauder.ubc.ca (C. Wu), yyyuanzhe@gmail.com (Z. Yuan), zhangkaifu314@gmail.com (K. Zhang)

https://orcid.org/0009-0000-4184-4929 (L. Fang), https://orcid.org/0000-0001-6445-8000 (C. Wu), https://orcid.org/0000-0002-7502-9102 (Z. Yuan)

Received: November 18, 2025
Revised: December 12, 2025; May 1, 2026; July 20, 2026
Accepted: July 27, 2026
Published Online in Articles in Advance: September 23, 2026

https://doi.org/10.1287/mksc.2025.0618

Abstract. Product returns create substantial financial burdens for e-commerce firms and hassle for customers. We study a *returnless partial refund policy*, which lets customers keep products while receiving claim-dependent partial refunds. In collaboration with a leading global cross-border e-commerce platform, we conducted a randomized field experiment involving 853,521 customers assigned to either returnless partial refund offers or the traditional return process. The platform implemented the policy through a multiagent artificial intelligence system that evaluated multimodal evidence and generated real-time offers at scale. Overall, the policy generated net transaction value gains of $1.44 per customer over 30 days and $4.41 over 90 days, driven by reduced return processing costs and increased subsequent customer spending, despite increasing refund payouts. With the experimental sample, the policy reduced return shipments, yielding estimated carbon savings equivalent to the sequestration of roughly 2,000 mature trees. These benefits were accompanied by more subsequent return requests with patterns consistent with lower-quality, opportunistic claims, revealing a tradeoff between service recovery and strategic behavior. Leveraging variation in refund generosity across two phases, we calibrate a revenue-maximizing average refund ratio of approximately 41% that balances customer value, operating costs, and opportunism.

History: Catherine Tucker served as the senior editor.
Funding: L. Fang acknowledges financial support from the National Natural Science Foundation of China [Grant 72501258], Major Special Project in Philosophy and Social Sciences of the Ministry of Education of China [Grant 2026JZDZ020], Major Program of the National Natural Science Foundation of China [Grant 72192803], and Major Project of the National Social Science Fund of China [Grant 22&ZD081]. Z. Yuan acknowledges financial support from the National Natural Science Foundation of China [Grants 72673168, 72141305, and 72192803], the National Key Research and Development Project [Grant 2024YFB3312900], Major Project of the National Social Science Fund of China [Grants 26ZDA030 and 25&ZD149], and the Fundamental Research Funds for the Central Universities.
Supplemental Material: The online appendices and data files are available at https://doi.org/10.1287/mksc.2025.0618.

Keywords: returnless refunds • e-commerce • service recovery • strategic customer behavior • field experiment

1. Introduction

Managing product returns requires e-commerce firms to balance operational efficiency against customer experience. The scale of this challenge is substantial: In 2024, U.S. retailers faced an estimated $890 billion in returns, equivalent to 16.9% of total sales, whereas the online return rate was 21% higher than the overall rate (NRF and Happy Returns 2024). The burden falls on both sides of the market: Firms incur processing and depreciation costs that erode margins (Ofek et al. 2011, Nageswaran et al. 2020), whereas customers bear time, effort, and psychological hassle that can deter return completion, reduce satisfaction, and weaken lifetime value (Ertekin 2018, Abdulla et al. 2022).

To address these return-related frictions, retailers including Amazon, Walmart, and Shein have introduced *returnless refund policies*, which provide full refunds without requiring customers to return products (Hadero 2024, Holman 2024). These policies typically apply to low-value items for which return-processing costs exceed product value. A *returnless partial refund policy* offers a middle ground between a full returnless refund and a traditional return: Customers keep products while receiving partial refunds based on their claims and supporting evidence, thereby reducing return hassle while limiting refund payouts.

This paper studies a large-scale implementation of returnless partial refunds. We assess whether the policy creates net economic value for the platform relative to traditional returns and examine how refund generosity shapes the tradeoff among refund payouts, processing cost savings, future customer spending, and strategic behavior.

We evaluate the policy using a randomized field experiment involving more than 850,000 customers on a leading global cross-border e-commerce platform. Eligible customers initiating a return were randomly assigned to either a treatment group, under which they could receive a returnless partial refund offer, or to a control group, under which only the traditional return process applied. The experiment comprised two phases: a high-refund phase with an average refund ratio of 50%, followed by a low-refund phase averaging 30%. This design identifies the policy's overall effect through customer-level random assignment and uses cross-phase variation in refund generosity to inform how its performance varies with refund levels.

Implementing returnless partial refunds at scale requires rapid and consistent evaluation of multimodal claim evidence across millions of returns, which could be prohibitively costly and time-consuming for human agents. The platform addressed this bottleneck through a multiagent artificial intelligence (AI) system that evaluated evidence and generated offers in real time, making large-scale implementation feasible.

The returnless partial refund policy creates competing economic and behavioral effects. By reducing shipback hassle, the policy may improve service recovery and stimulate future purchases while lowering the platform's logistics, inspection, and handling costs. Yet this same convenience also carries countervailing risks: Easier refunds may increase return requests and refund payouts and, more concerningly, invite strategic behavior as customers test the policy's limits at low cost. To organize these tradeoffs, we develop an economic framework that decomposes the policy's impact into three channels: (i) immediate net transaction value, reflecting refund payouts and return processing costs; (ii) future customer value, reflecting subsequent spending; and (iii) potential strategic behavior as customers adapt to the policy. This decomposition guides our empirical analysis.

We find that the policy reduces immediate net transaction value but generates a positive overall gain once subsequent customer spending is included. At the initial return, higher refund payouts exceed operational cost savings, reducing immediate net transaction value by 3.2%. This average masks important heterogeneity driven by refund generosity: Immediate net transaction value declines in the high-refund phase but increases by 1.3% in the low-refund phase. The long-term effects are strongly positive. Over the following 30 days, treated customers spend 2.2% ($1.92 per customer) more, with gains appearing in both phases and persisting through 90 days. An instrumental variables analysis confirms a positive effect of more generous refund offers on future spending.

These benefits are accompanied by patterns consistent with strategic behavior. Treatment increases subsequent return requests by 15.5%. Within the treatment group, approval rates and physical return follow-through decline most for claims most susceptible to opportunistic behavior, including those involving subjective reasons, high-priced items, and rapid repeat requests. Aggregating the immediate and long-term effects, the policy generates net transaction value gains of $1.44 per customer over 30 days and $4.41 over 90 days. It also reduces return shipments, yielding estimated carbon savings equivalent to the sequestration of roughly 2,000 mature trees in our experimental sample. Finally, calibration using cross-phase variation in refund generosity identifies an average refund ratio of approximately 41% as revenue-maximizing when long-term customer value is included.

This paper contributes to marketing and service operations research by providing the first large-scale causal evidence on a returnless partial refund policy. We quantify its effects on immediate costs, subsequent customer value, strategic behavior, and environmental outcomes. We further document how these effects vary with refund generosity and provide a framework for calibrating refund levels. The setting also illustrates how a multiagent AI system can make evidence-based partial refunds operational at scale.

Our findings offer practical managerial guidance: Firms should evaluate returnless partial refunds using both immediate costs and subsequent customer value and calibrate refund level to balance service recovery benefits against refund payouts and strategic behavior. The framework also applies to related service contracts, such as warranties and claims management, involving similar economic and behavioral tradeoffs.

Our study relates to three streams of research: return policy design, service recovery and strategic behavior, and AI-enabled service operations.

1.1. Return Policy Design

Return policies create option value for consumers but impose costs on firms, requiring retailers to balance customer value against operational costs (Anderson et al. 2009). Their effects depend on several dimensions of leniency, with monetary compensation and return effort particularly important for consumer demand (Janakiraman et al. 2016; Abdulla et al. 2019, 2022), helping explain the growing adoption of returnless refund policies. Lenient return policies can signal quality (Moorthy and Srinivasan 1995), reduce prepurchase risk (Che 1996), stimulate purchases (Wood 2001), and foster long-term loyalty (Bower and Maxham 2012), but they can also increase return rates and erode margins (Davis et al. 1998). Returnless partial refunds are distinctive because they combine high effort leniency—customers avoid shipping products back—with intermediate monetary leniency—a partial rather than full refund. Recent evidence documents brand-building effects of returnless returns (Costello and Bechler 2026), whereas analytical models suggest that partial refunds can manage dissatisfaction, limit opportunism, and improve retailer profits relative to full refunds (Chu et al. 1998, Su 2009). We extend this literature by providing the first large-scale causal evidence on a returnless partial refund policy, quantifying its economic, behavioral, and environmental effects and examining how performance varies with refund generosity.

1.2. Service Recovery and Strategic Behavior

Service recovery research shows that recovery quality can matter more than the initial failure (Smith et al. 1999): Effective recovery can strengthen customer satisfaction, trust, commitment, and loyalty (Hart and Heskett 1990, Tax et al. 1998, De Matos et al. 2007). Consistent with this view, generous return policies can increase long-term customer spending despite higher immediate costs (Petersen and Kumar 2009, Bower and Maxham 2012). Yet greater generosity and reduced hassle can also create moral hazard, encouraging strategic behavior ranging from claim exaggeration to overt abuse such as "wardrobing" (Shang et al. 2017, Zhang et al. 2020, Altug et al. 2021). Existing research leaves unclear how these opposing forces jointly shape the value of a return policy. We address this gap by showing that returnless partial refunds increase both subsequent spending and return requests, with the latter accompanied by patterns consistent with strategic adaptation. We further examine how subsequent spending and strategic behavior vary with refund generosity.

1.3. AI-Enabled Service Operations

Digital technologies have long reshaped customer-service delivery (Ray et al. 2005, Goldfarb and Tucker 2019). Recent field studies show that generative AI assistance can increase service productivity and quality, particularly among less-experienced workers (Brynjolfsson et al. 2025, Ni et al. 2026), although its effectiveness can decline on complex, context-dependent tasks (Dell'Acqua et al. 2026, Zhang and Narayandas 2026). Most existing research examines AI as a tool that augments human agents rather than as infrastructure for implementing evidence-intensive service policies (Autor 2024, Krakowski et al. 2026). Emerging agentic AI systems extend beyond assistance by integrating reasoning and action across complex tasks (Yao et al. 2023, Durante et al. 2024). Our setting illustrates how a multiagent AI system can evaluate multimodal claims and generate real-time offers, making a returnless partial refund policy operational at scale.

The rest of the paper is organized as follows. Section 2 details the context and experimental design. Section 3 presents the economic framework and empirical results. Section 4 concludes.

2. Context, Experimental Design, and Data

2.1. Research Context

Our research setting is a leading global cross-border e-commerce platform that serves hundreds of millions of shoppers across more than 100 countries and operates in approximately 20 languages. The platform's scale and geographic and linguistic diversity make consistent, 24/7 and real-time customer service operationally challenging, particularly for return claims that require nuanced, evidence-based evaluation.

Prior to our intervention, the platform employed a basic rule-based, evidence-free returnless refund policy for its first-party products.[1] Under this policy, customers returning low-value items (typically under \$10) automatically received full refunds without undergoing claim review and shipping the items back. Although simple to administer, this policy relied on a rigid, binary cutoff: Qualifying low-value items received full returnless refunds, whereas other items were routed to the traditional return process. It therefore provided no flexible resolution for the vast majority of moderate-value products requiring more nuanced evaluation.

Our experiment introduced the returnless partial refund policy that expanded the existing approach in two ways. First, it extended the returnless option to moderate-value products (above \$10). Second, and more critically, it replaced the all-or-nothing compensation structure with flexible claim-dependent partial refunds, allowing customers to keep the products while receiving partial refunds.

The platform implemented the policy through a three-agent AI pipeline that sequentially performed customer intent recognition, multimodal evidence analysis (text, image, and video), and refund resolution. The pipeline processed each case in approximately five seconds, making an otherwise prohibitively costly policy—estimated to require over half a million dollars in labor for our experimental sample alone and roughly five minutes per manual decision—operationally scalable across millions of claims. The system's prompts, parameters, and evaluation logic remained fixed throughout the experiment; it evaluated submitted evidence but did not model customer responses, predict acceptance or ship-back decisions, or use reinforcement learning to optimize refund ratios. Online Appendix A provides the workflow diagrams and technical details.

2.2. Experimental Design, Data, and Key Statistics

We conducted a large-scale randomized controlled trial (RCT) from August 2024 to January 2025. Before the experiment began, the platform randomly assigned a subset of customers to treatment or control, with assignment remaining fixed at the customer level throughout the study period. Customers entered the experiment upon initiating a return for a first-party product exceeding the platform's low-value threshold, with claimed reasons deemed appropriate for AI-based verification.[2] Because customers had no prior knowledge of the returnless partial refund policy before initiating the first return, their entry into the experiment, triggered by this return event, is plausibly random with respect to treatment assignment (see Online Appendix D). The final sample comprises 853,521 eligible customers.

In the treatment group (425,578 customers), the AI workflow evaluated eligible return requests and generated real-time partial refund offers for claims it approved, presented through the return interface in Figure 1. Treated customers learned of the policy only if they received an offer.[3] They could accept the offer and keep the item or decline it and proceed with the standard return process, but they could not negotiate the refund level. In the control group (427,943 customers), only the traditional return process applied: Customers could ship the item back for a full refund or abandon the request without compensation. They were not informed about the new policy. Because implementing returnless partial refunds through human review would not have been economically or operationally feasible, our estimates capture the combined effect of the AI-enabled policy rather than the incremental effect of AI alone.

A key design feature was the platform-initiated variation in refund generosity across two experimental phases, independent of the AI evaluation process. The platform adjusted the fixed refund ratios in a predefined mapping table linking AI-assessed product category, claim reason, and issue severity to offer amounts, while leaving the evaluation logic and structure of the mapping unchanged throughout.[4] The initial high-refund phase (August–October 2024, Phase H) offered average refunds of 50% of the transaction value, providing an exploratory benchmark. The subsequent low-refund phase (October 2024–January 2025, Phase L) reduced the average refund ratio to 30%. The reduction was not uniform across all scenarios, with some category-reason pairs adjusted more than others, whereas the average reduction was similar across the broader objective and subjective reason groups. This exogenous variation provides the identification source

Figure 1. (Color online) Example of the Returnless Partial Refund Offer Interface

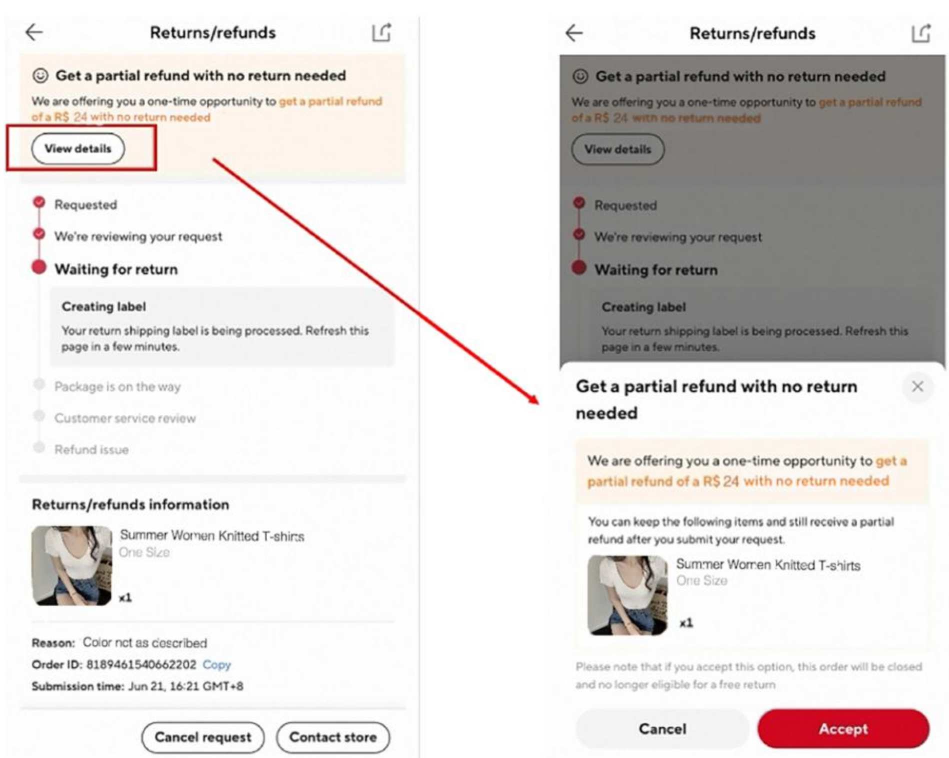

to estimate the causal effects of refund generosity and to calibrate the optimal refund level within our economic framework.

Our data set combines information from three primary sources. First, for each initial return, we obtained detailed transaction-level records, including the return reason, product category, return date, AI-generated refund offer, consumer decision (acceptance, ship-back, or abandonment), transaction value, final refund amount, and processing cost (e.g., shipping, inspection and handling). Processing costs vary substantially across return orders depending on product characteristics and logistics routes. Importantly, we directly observe these order-level processing costs and incorporate them into our net transaction value measure to account for cost savings from avoided ship-backs. Second, we track customer behavior for up to 90 days after the initial return, capturing both purchases (transaction value, purchase probability, order count, category variety, price per order) and returns (return frequency, refund amounts, processing costs). Third, we observe customer characteristics (gender, age category, years since registration, country and language index) and purchasing and return activity during the 30 days before enrollment and use these to account for individual heterogeneity.[5]

We assess the validity of randomization using balance tests presented in Table D1 in the Online Appendix. A comparison of extensive pre-experiment demographic and behavioral characteristics revealed no statistically significant differences between the treatment and control groups, confirming that randomization successfully created two comparable cohorts. Moreover, customers were not informed about the experiment, minimizing concerns about potential demand effects and ensuring that observed behavior reflects natural interactions with the platform.

Figure 2 visualizes three key metrics for the treatment group: the offer exposure rate (share of cases receiving a returnless partial refund offer); the average offer refund ratio (offered refund as a percentage of transaction value) among issued offers; and the offer acceptance rate (share of issued offers accepted by consumers). Figure 2(a) shows that the offer exposure rate remained stable between 60% and 70% throughout the experiment, consistent with the system's fixed evaluation logic.[6] This partial exposure attenuates the estimated

Figure 2. (Color online) Offer Exposure Rate, Refund Ratio, and Acceptance Rate over Time

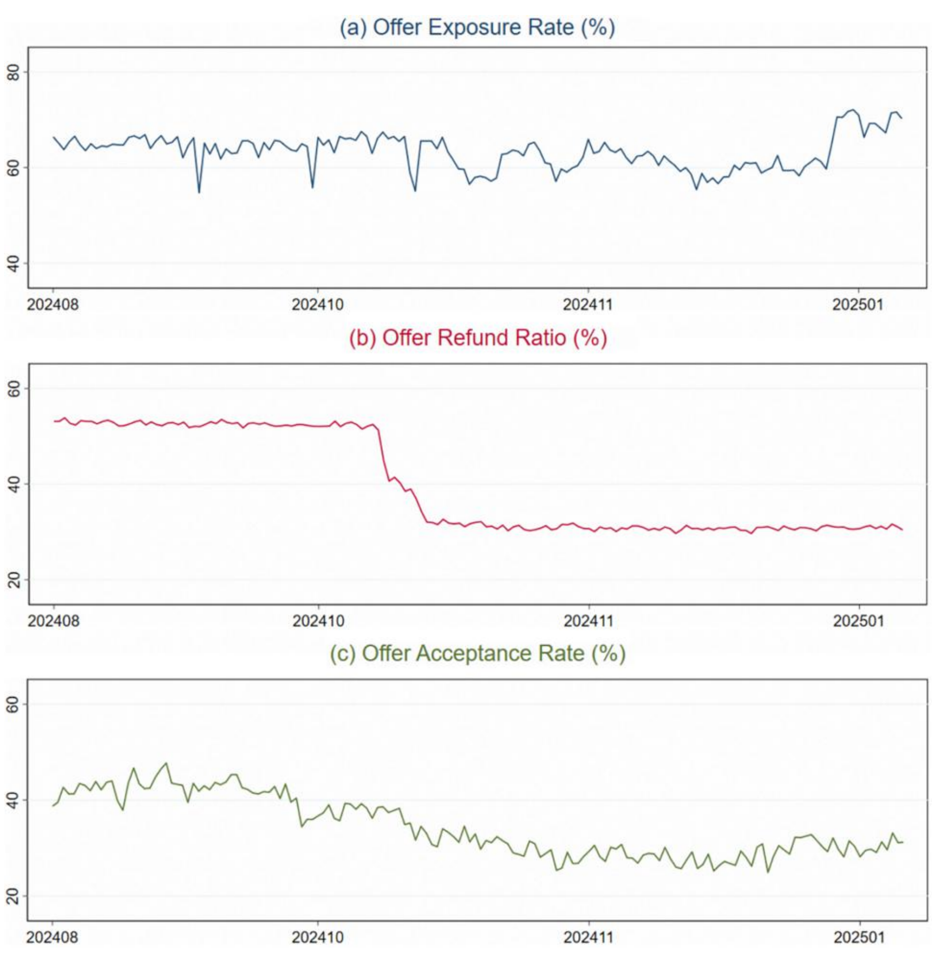

intent-to-treat (ITT) effect, making it a conservative estimate as outcomes are averaged across treatment group customers who do and do not receive an offer. Consistent with the cross-phase refund variation, Figure 2(b) shows the average offer refund ratio declined from approximately 50% to 30% following the phase transition. Correspondingly, Figure 2(c) reveals that the offer acceptance rate decreased from about 40% to 30%, suggesting clear behavioral sensitivity to refund generosity.

A defining feature of the multiagent AI system is its ability to provide nuanced, context-sensitive decisions that vary by product category, return reason, and issue severity, resulting in substantial heterogeneity within the treatment group. Table 1 summarizes these variations across the three metrics. For objective and verifiable claims (e.g., damaged packaging), the system generated offers in 79.9% of cases. These offers averaged 36.3% of transaction value and were accepted 40.2% of the time. For more subjective claims, such as dissatisfaction with product quality or appearance, the offer exposure rate was lower at 53.1%. Although these offers averaged a slightly higher 39.4% of transaction value, customers accepted only 28.0%. These differences demonstrate that the policy did not apply a uniform resolution across claims and motivate our subsequent analysis of heterogeneous customer responses.

The combination of large-scale experimentation, validated randomization, and exogenous refund variation provides a robust foundation for causal inference. In the following section, we develop an economic framework and use our experimental data to quantify the policy's behavioral and financial outcomes.

3. Analysis and Results

This section presents our analysis and results. We begin with a parsimonious economic framework and then examine the policy's immediate and subsequent economic impacts and evidence of customer strategic behavior. We then integrate these findings to assess the policy's overall economic impact and calibrate the revenue-maximizing refund level.

3.1. Economic Framework

3.1.1. Traditional Return Process. Consider a transaction with value V. Under the traditional return process (the control condition), the platform's expected payoff R decomposes into two components:

$$R = NTV + LCV = V[1 - r_s(1 + c - s)] + LCV. \quad (1)$$

Here, NTV (net transaction value) is the current-period payoff accounting for returns; LCV (long-term customer value) is the customer's subsequent-period net transaction value following the initial return request; r_s is the probability of a physical return shipment, and c and s denote the processing cost and salvage value, respectively, each expressed as a fraction of V.

3.1.2. Returnless Partial Refund Process. Under the returnless partial refund policy (the treatment condition), the refund level α influences whether a customer accepts the returnless partial refund, ships the item back through the traditional process, or abandons the return request. With P denoting the "policy" condition, the platform's expected payoff has the same two-component structure:

$$R^P(\alpha) = NTV^P(\alpha) + LCV^P(\alpha) = V[1 - r_a^P(\alpha)\alpha - r_s^P(\alpha)(1 + c - s)] + LCV^P(\alpha). \quad (2)$$

Here, $r_a^P(\alpha)$ is the probability that the customer accepts the offer, $r_s^P(\alpha)$ is the probability that the customer declines and ships the item back, and the residual $r_o^P(\alpha) = 1 - r_a^P(\alpha) - r_s^P(\alpha)$ captures the probability that

Table 1. Returnless Partial Refund Resolution Statistics by Return Reason

Return reason	Offer exposure rate (%)	Offer refund ratio (%)	Offer acceptance rate (%)
Objective reasons (43%)			
Damaged packaging	86.4	25.3	43.4
Product arrived damaged	82.0	37.5	39.8
Scratched	82.4	39.3	41.2
Product flaw	68.7	42.9	36.6
Objective average	79.9	36.3	40.2
Subjective reasons (57%)			
Does not work properly	57.3	37.8	29.2
Poor quality	64.5	40.3	32.0
Appearance not as described	69.0	36.2	23.4
Color not as described	67.7	42.1	29.8
Product not as described	5.9	36.2	31.2
Style not as described	54.2	44.1	22.3
Subjective average	53.1	39.4	28.0

Notes. Offer exposure rate is the share of cases receiving a returnless partial refund offer; average offer refund ratio is the offered refund as a percentage of transaction value, averaged across issued offers; offer acceptance rate is the share of issued offers accepted by consumers. Objective reasons primarily refer to verifiable physical defects, whereas subjective reasons are more related to customer perception.

the customer abandons the return without taking either action. We retain salvage value s for completeness and generality; in our setting, it is effectively zero, because cross-border reverse-logistics and resale-related costs leave negligible product recovery value.[7]

The policy also impacts future customer spending, which we capture through the long-term customer value component $LCV^P(\alpha)$. We express this value relative to the traditional return baseline as

$$LCV^P(\alpha) = LCV[1 + \delta(\alpha)], \tag{3}$$

where $\delta(\alpha)$ denotes the proportional change in subsequent net transaction value induced by the policy at refund level α.

Equations (1)–(3) identify three empirical targets: (i) the immediate net transaction value NTV^P, capturing how the policy redistributes customer behavior across three outcomes—offer acceptance (at rate r_a^P), ship-back (r_s^P), or abandonment (r_o^P); (ii) the long-term customer value impact LCV^P, reflecting the net effect of enhanced future purchases driven by frictionless service recovery (service recovery effect) and increased opportunistic return behavior induced by easier refunds (moral hazard effect); and (iii) the behavioral parameters $r_a^P(\alpha)$, $r_s^P(\alpha)$, and $\delta(\alpha)$, which govern the policy's response to refund generosity.

3.2. Immediate Policy Outcomes

3.2.1. Estimation. We estimate the immediate ITT effect of the returnless partial refund policy using the following linear regression:

$$y_i = \beta_0 + \beta_1 Treatment_i + \boldsymbol{X}_i'\boldsymbol{B} + \epsilon_i, \tag{4}$$

where y_i denotes an outcome from customer i's initial return, $Treatment_i$ indicates assignment to the policy group, and $\boldsymbol{X}_i$ includes customer demographics, pre-experiment behavior, and initial return characteristics. The coefficient β_1 captures the effect of treatment assignment (ITT) regardless of whether the customer received or accepted an offer. When measuring net transaction value, β_1 corresponds to $NTV^P - NTV$ (Equations (1) and (2)); estimates for refund amounts and processing costs isolate the refund payout channel ($r_a^P\alpha V + (r_s^P - r_s)V$) and ship-back processing cost channel ($(r_s^P - r_s)cV$) underlying this difference.

To examine how the immediate effect differs across experimental phases, we estimate

$$\begin{aligned} y_i = \beta_0 &+ \beta_{11} Treatment_i \times I(Phase_i = H) \\ &+ \beta_{12} Treatment_i \times I(Phase_i = L) + \boldsymbol{X}_i'\boldsymbol{B} + \epsilon_i. \end{aligned} \tag{5}$$

Here, β_{11} and β_{12} capture the ITT effects during the high- and low-refund periods, respectively, whereas the difference $\beta_{12} - \beta_{11}$ measures how the treatment effect changes following the reduction in refund generosity.

3.2.2. Overall Effect and Phase Heterogeneity. Table 2 reports the immediate ITT estimates. Panel A presents the overall policy effect across both phases, and Panel B reports the treatment effects separately for the high- and low-refund phases.

The results reveal the policy's core economic tradeoff. Treatment increased refund payouts by \$0.63 per customer, or 4.2% (column 2) while reducing processing costs by \$0.35, or 11.8% (column 3). The operational cost savings offset 56% of the additional refund

Table 2. Immediate Impact of Returnless Partial Refund Policy: Overall Effect and Phase Heterogeneity

Variables	(1) Transaction value	(2) Refund amount	(3) Processing cost	(4) Net transaction value (NTV)
Panel A: Overall intent-to-treat effect				
Treatment	0.003	0.633***	−0.352***	−0.277***
	(0.053)	(0.048)	(0.008)	(0.046)
Percentage of change	0.0%	4.2%	−11.8%	−3.2%
Observations	853,521	853,521	853,521	853,521
Panel B: Treatment effect by experimental phase				
Coefficients				
Treatment:High Refund	0.130	1.490***	−0.457***	−0.903***
	(0.088)	(0.081)	(0.013)	(0.075)
Treatment:Low Refund	−0.057	0.126**	−0.285***	0.103*
	(0.070)	(0.062)	(0.009)	(0.060)
Effect size (%)				
Treatment:High Refund	0.5%	9.4%	−15.3%	−10.0%
Treatment:Low Refund	−0.2%	0.9%	−9.6%	1.3%

Notes. Standard errors in parentheses. Panel A reports the overall ITT effect across both phases from Equation (4). Panel B decomposes the treatment effect by experimental phase using Equation (5) (High: average 50%; Low: 30%), excluding the one-week transition period in October during which refund levels gradually declined. Net Transaction Value = Transaction Value − Refund Amount − Processing Cost. Percentage changes are relative to control-group means. Controls include demographics, prior 30-day purchase and return history, and initial return characteristics (reason, product category, and date).

***$p < 0.01$; **$p < 0.05$; *$p < 0.1$.

payouts, leading to a \$0.28, or 3.2% decline in immediate net transaction value (column 4).[8]

Critically, the overall estimate masks substantial differences across phases. During the *high-refund phase* (average 50%), the policy increased refund payouts by \$1.49 and reduced processing costs by \$0.46. The significant increase in payouts outweighed the operational savings, resulting in a decline of \$0.90 in net transaction value. In contrast, during the *low-refund phase* (average 30%), the increase in refund payouts was marginal (\$0.13), whereas processing costs fell by \$0.29. Immediate net transaction value consequently increased by \$0.10, or 1.3%. Thus, reducing refund generosity reversed the direction of the policy's immediate economic effect, and at moderate refund levels, the policy delivers immediate net benefits by balancing operational savings and refund payout costs. We next examine how acceptance and ship-back decisions vary with the offered refund level.

3.2.3. Refund Generosity and Customer Decisions. Differences in refund generosity systematically alter how customers respond to the returnless partial refund offers.[9] As shown in Figure 3, across both objective and subjective return reasons, higher refund offers are associated with higher offer acceptance and lower ship-back rates.[10] These patterns illustrate the behavioral channels through which refund generosity can affect immediate policy outcomes.

3.2.4. Reallocation of Return Outcomes. The returnless partial refund policy introduces a new intermediate resolution between the two traditional outcomes of receiving no refund after abandoning the request and receiving a full refund after shipping the product back. Figure 4 visualizes this treatment-control distributional shift in return outcomes. The treatment group exhibits pronounced new peaks at intermediate realized refund ratios (primarily between 20% and 70%), accompanied by declines near 0% (abandonment, zero refund) and 100% (physical returns, full refund).

Figure 4. (Color online) Distributional Shift in Refund Outcomes: Treatment vs. Control

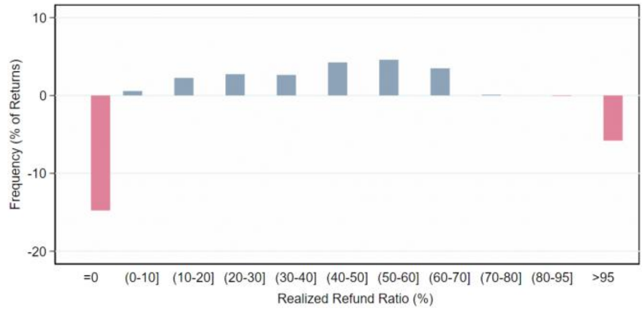

Note. Each bar shows the percentage point (p.p.) difference (Treatment − Control) in the share of returns by realized refund ratio (realized refund payout/transaction value of the return).

This reallocation produces two countervailing effects on immediate platform value. When a customer who would otherwise abandon the return accepts a partial refund, the platform incurs an additional payout. When a customer who would otherwise ship the product back accepts a partial refund, the platform avoids processing costs and pays less than a full refund. The net financial outcome depends on the relative magnitudes of these two flows. Our results indicate that the policy draws more customers from abandonment (a 14.8-percentage-point (pp) shift) than from physical returns (a 6.8-pp shift), explaining the overall net increase in refund payouts in Table 2. At the same time, faster and less burdensome resolutions may generate subsequent customer value, which we examine next.

3.3. Long-Term Customer Value

3.3.1. Estimation. To evaluate the policy's effect on subsequent behavior and *LCV*, we apply the same ITT

Figure 3. (Color online) Relationship Between Offer Refund Ratio and Customer Return Decisions

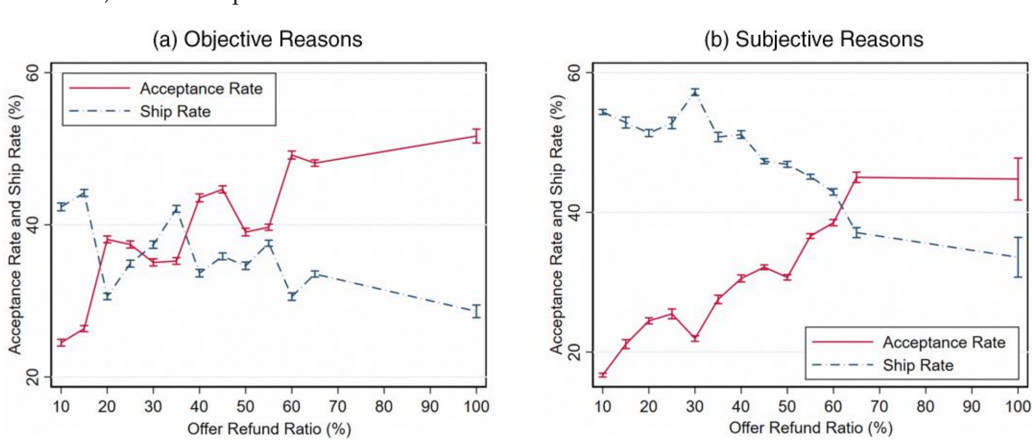

Notes. The solid and dash-dotted lines show the offer acceptance rate and ship-back rate, respectively. Error bars represent standard errors.

specification to outcomes measured during the 30 days following the initial return:

$$z_i = \gamma_0 + \gamma_1 Treatment_i + \mathbf{X}_i'\mathbf{\Gamma} + \varepsilon_i, \tag{6}$$

where z_i includes subsequent purchase and return outcomes. When measuring z_i in net transaction value, γ_1 corresponds to the subsequent payoff difference of $LCV^P - LCV$ (Equations (1) and (2)); the additional spending and return measures trace the behavioral mechanisms—service recovery versus moral hazard—that determine the sign of $\delta(\alpha)$ in Equation (3).

To estimate how the outcomes respond to refund generosity, we employ a two-stage least squares (2SLS) approach, using the phase indicator ($I(\text{Phase}_i = L)$) as an instrument for the initial offer refund ratio α_i. This analysis is restricted to treated customers who received a returnless partial refund offer:

$$z_i = \lambda_0 + \lambda_1 \hat{\alpha}_i + \mathbf{X}_i'\mathbf{\Lambda} + u_i. \tag{7}$$

When λ_1 captures the impact of refund generosity on subsequent transaction value, it traces the service recovery channel within $\delta(\alpha)$ in Equation (3).

3.3.2. Subsequent Spending Patterns. Table 3 reports the 30-day ITT effects on subsequent purchasing. Customers in the treatment group exhibited significantly stronger postreturn engagement. Over the 30-day window, their total transaction value increased significantly by 2.2% ($1.92 per customer, column 1). This aggregate lift was driven by several positive behavioral shifts: They were more likely to make a subsequent purchase (a 1.2% relative increase, column 2), purchased more frequently (1.7% more orders, column 3), and bought from a broader range of product categories (1.6% more, column 4). Average price per order declined slightly by 1.0% (column 5). Because customers were randomly assigned and treated customers learned of the policy only after initiating their initial return, these subsequent spending effects are consistent with service recovery generated by the return experience.

To test whether these effects were transient or persistent, we plot the weekly transaction value for treatment and control customers across the 4 weeks preceding and 12 weeks following the initial return (Figure 5). No significant differences appear in the preperiod, whereas treatment customers consistently spend more during the postreturn period. Extending Table 3 to a 90-day window reveals consistent findings: Treated customers' transaction value increased by 2.1% ($4.89 per customer), driven by a 0.7% higher purchase likelihood, 1.7% more orders, and 1.3% greater category breadth (Table C1 in the Online Appendix). These patterns indicate that the frictionless resolution process fosters a sustained engagement increase rather than a short-lived spending surge.

3.3.3. Effect of Refund Generosity on Spending. Table 4 examines how subsequent spending varies with offer exposure and refund generosity, based on the 2SLS model estimation in Equation (7). The results reveal a positive and statistically significant effect. The coefficient on the instrumented refund level $\hat{\alpha}$ is 11.95 (column 3), indicating that a 10-pp increase in the offered refund ratio raises subsequent 30-day transaction value by approximately $1.20. This finding is consistent with service recovery theory: More generous resolutions enhance return experience and strengthen customer loyalty.

3.3.4. Subsequent Net Economic Impact. Table 5 decomposes the policy's subsequent net economic impact over the 30-day period following the initial return, based on Equation (6). The results demonstrate a significant net gain in long-term customer value. The treatment group's net transaction value was 2.2% ($1.71 per customer) higher than that of the control group (column 4). Specifically, subsequent transaction value increased by $1.92 (column 1), which more than offset a $0.25 rise in refund payouts (column 2). The policy also generates modest processing savings ($0.05, column 3). The positive subsequent net economic effect was

Table 3. Subsequent Impact of Returnless Partial Refund Policy on Customer Spending

Variables	(1) Transaction value	(2) Purchase probability	(3) Purchase frequency	(4) Category expansion	(5) Price
Treatment	1.916***	0.007***	0.130***	0.072***	−0.155*
	(0.617)	(0.001)	(0.050)	(0.016)	(0.087)
Constant	−17.61***	0.469***	1.062***	2.489***	15.480***
	(2.418)	(0.001)	(0.316)	(0.035)	(0.111)
Percentage of change	2.2%	1.2%	1.7%	1.6%	−1.0%
Observations	853,521	853,521	853,521	853,521	521,631
R^2	0.570	0.100	0.373	0.166	0.023

Notes. Standard errors in parentheses. Thirty-day postreturn window. Purchase probability = 1 if any subsequent purchase. Purchase frequency is the number of subsequent orders. Category expansion is the number of product categories purchased. Price per order is the mean transaction value per order among customers with a subsequent purchase. Percentage changes are relative to control-group means. Controls include demographics, prior 30-day purchase and return history, and initial return characteristics (reason, product category and date).
***$p < 0.01$; **$p < 0.05$; *$p < 0.1$.

Figure 5. (Color online) Transaction Value by Weeks Preceding and Following Initial Return Request

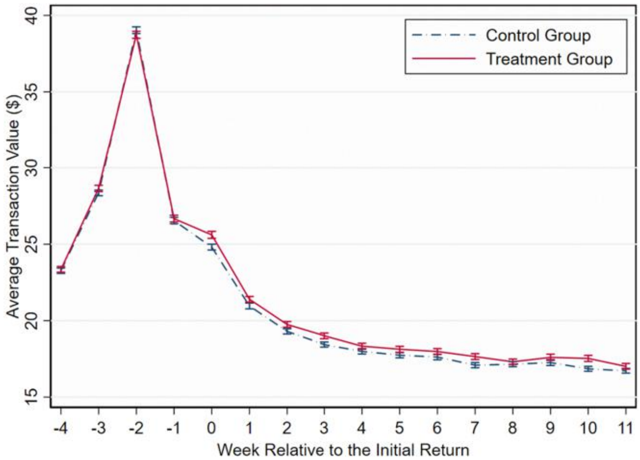

Notes. The solid and dash-dotted lines with standard error bars show the weekly average transaction value for the treatment and control groups, respectively. Weeks are indexed relative to initial return (event week 0). Elevated prereturn transaction value reflects the fact that returns are conditioned on recent purchases, whereas the gradual postreturn decline reflects normal purchase cycles, interpurchase intervals, and natural customer attrition.

consistent across both experimental phases, reinforcing that these benefits are robust to refund generosity.

Importantly, we observe clear signs of moral hazard. The proportion of returned orders for treated customers was 4.3% higher than that of the control group (column 5)—nearly double their spending growth rate of 2.2%, indicating that easier refunds increased opportunistic behavior, which we examine in detail in the next subsection.

Collectively, these results show that the returnless partial refund policy produces positive net value. Although it induces more strategic returns, the accompanying rise in customer spending driven by improved experience is substantially larger, highlighting the economic value of frictionless service recovery. This pattern sustains over a 90-day horizon, with treated customers exhibiting a 2.2% increase in net transaction value ($4.69 per customer; Table C2 in the Online Appendix).

A potential threat to our estimates is cross-phase contamination: if a treated customer initiates returns in both phases within the 30-day window, they may receive different offer levels, affecting their response. We find that only 0.1% of treated customers are potentially affected, and excluding them does not alter our main findings. If anything, these estimates should be interpreted as conservative: Treated customers experiencing both phases may feel disappointment relative to their earlier, higher reference point, attenuating observed gains in customer value, and similarly biasing the 2SLS effect of refund generosity downward.

3.4. Strategic Customer Behavior

3.4.1. Patterns of Strategic Behavior. The overall increase in subsequent customer value reflects two opposing forces: positive service recovery effects and potential moral hazard. To separate these mechanisms, we exploit the feature that the AI system's prompts, parameters, and evidence-evaluation logic remained fixed throughout the experiment and did not learn from customer responses. Changes in approval outcomes therefore cannot be attributed to evolving adjudication standards. Instead, they provide a consistent basis for examining whether the composition or quality of submitted claims changes following policy exposure.

We find three complementary pieces of evidence consistent with strategic behavior, based on customers' return activity within 30 days of their initial request, presented in Table 6. First, treated customers initiated 15.5% more subsequent return requests than the control group. Second, within the treatment group, the system's approval rate for subsequent requests declined by 1–3 pp relative to initial requests. Third, when the system

Table 4. Impact of Offer Refund Ratio on Subsequent Transaction Value

Variables	(1) Transaction value	(2) Transaction value	(3) Transaction value (IV)
If Exposed to Initial Partial Refund Offer	−0.027		
	(1.137)		
Initial Offer Refund Ratio (α)		7.357**	11.950**
		(3.376)	(5.157)
Constant	−105.0***	−109.2***	−146.9***
	(3.371)	(4.214)	(19.80)
Observations	425,578	256,382	256,382
R^2	0.565	0.623	0.623
Kleibergen-Paap Wald rk F-statistic			119,830

Notes. Standard errors in parentheses. Thirty-day postreturn window. Column 1: treatment group. Columns 2 and 3: treatment group receiving a returnless partial refund offer on initial return. Column 3 employed IV estimation using high/low phase as an instrument. Because the estimation is restricted to the treatment group, we demean each customer's transaction value by the corresponding control-group mean within the same return date, reason, and category cell to control for seasonality. Controls include demographics, prior 30-day purchase and return history, and initial return characteristics (reason, product category).

$^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$.

Table 5. Subsequent Impact of Returnless Partial Refund Policy on Key Metrics

Variables	(1) Transaction value	(2) Refund amount	(3) Processing cost	(4) Net transaction value (LCV)	(5) Order return rate
Treatment	1.916***	0.251**	−0.046***	1.711***	0.004***
	(0.617)	(0.112)	(0.018)	(0.586)	(0.0002)
Constant	−17.61***	3.210***	0.526***	−21.34***	0.129***
	(2.418)	(0.263)	(0.062)	(2.357)	(0.0002)
Percentage of change	2.2%	3.1%	−3.5%	2.2%	4.3%
Observations	853,521	853,521	853,521	853,521	6,435,288
R^2	0.570	0.089	0.100	0.575	0.088

Notes. Standard errors in parentheses. Thirty-day postreturn window. Order return rate at order level; other metrics at customer level. Net Transaction Value = Transaction Value − Refund Amount − Processing Cost. Order return rate is the share of returned orders. Percentage changes are relative to control-group means. Controls include demographics, prior 30-day purchase and return history, and initial return characteristics (reason, product category and date).

***$p < 0.01$; **$p < 0.05$; *$p < 0.1$.

did not generate an offer, treated customers' physical return follow-through declined by 2–6 pp for subsequent requests. This reluctance to incur effort after an unsuccessful low-cost attempt suggests opportunistic "probing" rather than genuine complaint resolution.

This behavioral pattern is difficult to reconcile with a pure service recovery explanation, which would predict that a better experience would increase legitimate requests, thereby maintaining approval rates and ship-back follow-through rates. Instead, the combination of higher request volume, lower approval rates, and reduced physical return follow-through after rejection is more consistent with strategic probing than with service recovery alone.

Heterogeneity in these patterns further supports this interpretation and serves as a falsification test. Approval rates declined most precisely where opportunism is most feasible: subjective reasons (–2.18 pp), which are harder to verify; high-priced items (–3.02 pp), which offer stronger incentives to exploit the system; and rapid repeat requests within a week (−3.46 pp), which likely reflect probing behavior. Similar patterns appear for ship-back rates when AI rejects to make an offer. A simple enhancement of customer experience would not generate such concentrated declines in ambiguous, high-incentive contexts; instead, this pattern is most consistent with strategic adaptation to the new refund system.

3.4.2. Refund Generosity Induces Strategic Behavior. To examine whether refund generosity affects subsequent strategic behavior, we apply the same phase-based 2SLS approach to treated customers who received an initial offer. Table 7 presents three findings. First, a more generous initial refund significantly increases subsequent return attempts (column 1) and order return rate (column 2). Second, it reduces the approval rate of subsequent requests, although the estimate is insignificant (column 3). Third, it also reduces the likelihood that customers follow through with a physical return when no offer is made for a subsequent request (column 4), indicating lower willingness to exert effort for potentially strategic claims. The findings remain robust when extended to a 90-day window (Table C3 in the Online Appendix).

Under the instrumental variables identifying assumptions, these results indicate that more generous initial refunds increase subsequent return activity and reduce customers' willingness to complete a physical return when a later request does not receive an offer. Combined with the descriptive evidence of lower approval under fixed adjudication standards, they are consistent with opportunistic probing: Some customers submit additional claims but abandon them when an easy resolution is unavailable.

These findings reveal a central managerial tradeoff. Greater refund generosity strengthens service recovery and subsequent spending, but it also encourages additional return attempts and behavior consistent with strategic adaptation. Although the AI system filters many unsupported claims, evaluating these attempts can still increase operating costs and reduce the policy's net gains. This tradeoff motivates the refund-level calibration that follows.

3.5. Platform Economics Analysis

3.5.1. Cost-Benefit Analysis. Having estimated the policy's impact on immediate net transaction value (*NTV*, Section 3.2) and long-term customer value (*LCV*, Section 3.3), we combine these components to assess the policy's overall economic impact. Table 8 presents the cost-benefit analysis on a per-customer basis.

The results reveal several key effects. First, at the initial return, the policy increased refund payouts by \$0.63, and reduced processing costs by \$0.35, producing a \$0.28 decline in immediate net transaction value. Second, the policy generated substantial long-term gains. Over the subsequent 30 days, treated customers spent an additional \$1.92 accompanied by a modest \$0.25 increase in subsequent refunds and a small reduction in processing costs (\$0.05), yielding a \$1.71 increase in

Table 6. Returns, Approval Rate, and Ship Rate of Subsequent Returns

Metric	Control group	Treatment group	Difference	p-value
Number of subsequent returns	0.10	0.12	+0.02 (15.5%)	<0.01
Return distribution (%)				
By reason				
Objective (e.g., packaging)	40.89	43.85	+2.96 pp	<0.01
Subjective (e.g., quality)	59.11	56.15	−2.96 pp	<0.01
By price tier				
High price	50.33	52.14	+1.81 pp	<0.01
Low price	49.67	47.86	−1.81 pp	<0.01
	Treatment initial return	Treatment subsequent returns	Difference	p-value
Returnless partial refund approval rate (%)				
All	79.83	78.14	−1.69 pp	<0.01
By reason				
Objective (e.g., packaging)	87.10	85.62	−1.49 pp	<0.01
Subjective (e.g., quality)	73.31	71.13	−2.18 pp	<0.01
By price tier				
High price	82.27	79.25	−3.02 pp	<0.01
Low price	77.76	76.91	−0.85 pp	0.06
By request time				
Initial request	79.83	—	—	—
Subsequent requests within seven days	—	76.37	−3.46 pp	<0.01
Subsequent requests beyond seven days	—	79.18	−0.65 pp	0.07
Ship rate conditional on AI not making an offer (%)				
All	61.38	59.00	−2.38 pp	<0.01
By reason				
Objective (e.g., packaging)	47.30	44.90	−2.39 pp	0.11
Subjective (e.g., quality)	67.48	65.59	−1.89 pp	0.05
By price tier				
High price	70.22	63.88	−6.34 pp	<0.01
Low price	55.43	54.16	−1.27 pp	0.26
By request time				
Initial request	61.38	—	—	—
Subsequent requests within seven days	—	58.28	−3.10 pp	<0.01
Subsequent requests beyond seven days	—	59.70	−1.68 pp	0.08

Notes. Thirty-day postreturn window. Approval rate = number of requests for which AI approves a returnless partial refund offer/number of requests for which AI renders an approval or rejection decision. Ship rate conditional on AI not making an offer = number of no-offer requests that result in a physical return shipment/number of requests for which AI does not make an offer. These two panels only include return requests in the treatment group.

Table 7. Offer Refund Ratio and Strategic Behavior Using IV Estimation

Variables	(1) Number of subsequent returns	(2) Order return rate	(3) Approval rate	(4) Ship rate conditional on no offer
Initial Offer Refund Ratio (α)	0.069***	0.014***	−0.022	−0.150**
	(0.011)	(0.002)	(0.021)	(0.069)
Constant	−0.027	0.063***	0.593***	1.091***
	(0.030)	(0.006)	(0.092)	(0.214)
Observations	256,382	2,013,338	28,682	4,708
R^2	0.023	0.051	0.030	0.140
Kleibergen-Paap Wald rk F-statistic	119,830	950,292	12,064	2,243

Notes. Standard errors in parentheses. Thirty-day postreturn window. Column 1 estimated at customer level; column 2 at order level; columns 3 and 4 at return request level. Columns 1–4: restricted to treated customers receiving an initial offer and estimated using high/low phase as an instrument. Columns 1 and 2: outcomes are demeaned by the corresponding control-group mean within the same return date, reason, and category cell to control for seasonality. Column 3: further restricted to subsequent return requests for which AI renders an approval or rejection decision; Column 4: restricted to requests for which AI does not make an offer. Controls include demographics, prior 30-day purchase and refund history, and initial return characteristics (reason, product category).

***$p < 0.01$; **$p < 0.05$; *$p < 0.1$.

Table 8. Cost-Benefit Analysis (per Customer Within 30 Days of Initial Return Request)

Metric ($)	Difference	Percentage change	p-value
(1) Costs			
Increased refund amount			
Initial return request	(0.63)	(4.2%)	<0.01
Subsequent return requests	(0.25)	(3.1%)	0.03
Reduced processing cost			
Initial return request	0.35	11.8%	<0.01
Subsequent return requests	0.05	3.5%	<0.01
(2) Benefits			
Increased revenue			
Transaction value of initial return request	0.00	0.0%	0.95
Transaction value of subsequent purchases	1.92	2.2%	<0.01
(3) Platform net transaction value impact	1.44	1.6%	0.01
Back-of-the-envelope calculation			
(4) Platform net profit impact (35% margin)	0.19	1.5%	
(5) Customer retained value	1.49	18.0%	
(6) Total carbon footprint saving	2,000 trees		

Notes. Parentheses indicate costs. Net Transaction Value = Transaction Value − Refund Amount − Processing Cost. Panels (1)−(5) report per-customer values. Profit assumes a 35% margin (within the typical 30%−50% range for cross-border lifestyle, fashion and small consumer electronics e-commerce); retained value assumes a 50% residual product value. Panel (6) reports sample-wide carbon savings from reduced return shipping.

subsequent net transaction value. Aggregating these components, the returnless partial refund policy generated a positive net transaction value gain of $1.44 per customer, confirming that the gains from higher long-term customer value substantially outweigh the modest increase in initial costs. The gain increased to $1.52 when the exploratory high-refund phase was excluded. Because net transaction value does not account for the margin on incremental spending, we also provide an illustrative profit calculation: assuming a 35% gross margin consistent with the platform's primary product categories, the policy yields a positive net profit gain of $0.19 per customer overall and $0.62 during the low-refund phase.

These estimates may represent lower bounds of the policy's value. They exclude several meaningful benefits such as the consumer surplus retained from keeping products, the environmental gains from reduced logistics and product waste, and the intangible psychological value to customers from eliminating hassle costs. As an illustrative calculation, assuming customers retain roughly 50% of the purchase price as the residual value of the products they keep, this would add approximately $1.49 in economic benefit per customer.

Beyond financial performance, the policy also generates meaningful environmental benefits. We conduct a back-of-the-envelope calculation using standard carbon-offset conversion rates and compute the "tree-equivalent" effect. Within our experimental sample alone, the policy reduces the ship-back rate by roughly 12%, translating into carbon savings comparable to the sequestration of about 2,000 mature trees.[11]

The observed effects persist over a longer horizon. Over 90 days, the policy generated $4.41 in net transaction value and an estimated $1.23 in profit per customer, along with $1.89 in retained-product value and sample-wide carbon savings equivalent to the sequestration of about 3,000 mature trees. (Table C4 in the Online Appendix).

3.5.2. Revenue Curve and Refund Level Calibration. Refund generosity is a central strategic managerial decision for the platform. Although the system generates case-specific offers based on claim characteristics, our objective is to calibrate the average refund ratio that maximizes platform value and identify the range over which the policy outperforms traditional returns.[12]

Using the same phase-based 2SLS instrument approach, we estimate linear response functions linking the refund ratio α to offer acceptance ($r_a^P(\alpha)$), physical return ($r_s^P(\alpha)$), and subsequent net transaction value ($\delta(\alpha)$). The estimates show that higher refund ratios increase acceptance of returnless partial refunds, reduce physical returns, and boost long-term customer value (Table B1 in the Online Appendix). We substitute these response functions into the payoff framework to construct immediate and long-term platform-value curves and identify regions in which the policy outperforms traditional returns and calibrate the optimal refund ratio. Online Appendix B provides the estimation and calibration details.

The curves, presented in Figure B1 in the Online Appendix, reveal a stark contrast between short-term and long-term perspectives on the revenue-refund relationship. A myopic evaluation focusing only on the initial *NTV* suggests a highly conservative approach: The revenue-maximizing refund ratio is only 6%, and the policy outperforms traditional returns up to a refund ratio of 36% (left panel). However, incorporating the positive effect on *LCV* substantially alters the calculus. The break-even point expands from 36% to 90%, and the revenue-maximizing refund ratio increases to 41%. Notably, this calibrated long-run revenue-maximizing level lies between the observed phase averages of 50% and 30%, affirming that refund ratios in both periods fall within economically sensible ranges.

Taken together, these findings yield a central managerial insight: a narrow focus on immediate costs leads to an overly conservative refund policy, whereas a dynamic view that incorporates future customer value supports a broader range of generous refund levels that remain more profitable. This exercise also demonstrates a practical approach to fine-tune refund policies while balancing short-term costs, long-term customer value, and the risk of strategic behavior.

4. Conclusion

This study provides the first large-scale field experiment examining the impact of a returnless partial refund policy in customer service. We show that introducing the returnless partial refund option raises immediate refund payouts but substantially reduces operational costs and enhances long-term customer spending, yielding a net gain of about $1.44 per customer over 30 days and $4.41 over 90 days. A simple back-of-the-envelope carbon accounting further suggests that a significant drop in ship-back rates corresponds to emission savings comparable to the sequestration of about 2,000 mature trees in our experimental sample. Together, these results show that flexible, economically rational return contracts can transform returns from a traditional cost center into a strategic source of customer loyalty, growth, and environmental gains.

From a managerial standpoint, the findings underscore the importance of refund generosity and of adopting a dynamic evaluation horizon. Refund levels act as a powerful policy-design lever: At moderate levels, they create measurable long-term value; when overly generous, they invite strategic customer behavior and erode margins. Our calibration shows that a wide range of moderate refund levels can be profitable once long-term customer value is incorporated, with an average of around 41% of transaction value balancing service recovery benefits against the costs of opportunism. More broadly, multiagent AI automation can enable firms to apply novel customer service policies consistently at scale, combining operational efficiency with behavioral insight.

Future research should examine such a policy across platforms and product categories; measure customer and strategic responses beyond 90 days; and study information diffusion and competitive reactions. Experimental designs that separately vary the refund contract and adjudication technology could isolate the contribution of AI relative to human or rule-based implementation. Further work could also examine personalized and adaptive refund policies while addressing transparency, fairness, privacy, and regulatory constraints. More broadly, effective automated service recovery requires firms to design policy incentives and implementation technology jointly, turning routine service resolutions into lasting value.

Acknowledgments

The authors thank the editor and review team for constructive feedback and support; the seminar and conference participants at the 2025 CKGSB Marketing Workshop, the 2025 AI/ML Conference, and the 2026 Empirical and Theoretical Symposium for helpful comments and discussions; and the partner e-commerce company for supporting this project and its employees for patiently answering many questions. The views expressed in this paper are solely those of the authors and do not necessarily reflect those of any organization. All remaining errors are the authors' own.

Endnotes

[1] Both the basic policy and the subsequent returnless partial refund policy applied exclusively to first-party products fully managed by the platform, over which the platform has full discretion regarding return policies.

[2] The experiment included only a subset of claim reasons considered suitable for the AI system to verify through observable evidence, such as "wrong color" or "package damaged." Purely preference-based reasons such as "no longer needed" or "changed my mind" were excluded. Additional claim types can be incorporated in future extensions of the system.

[3] Meaningful word-of-mouth diffusion during the experimental period is practically unlikely. The experiment was conducted as an internal test with no public announcement of the policy. The treatment group represents a very small share (<1%) of the platform's large, geographically dispersed user base. Competitive responses are also unlikely: The policy was highly novel at the time, and to the best of our knowledge, no competing retailer introduced a similar policy during our experimental period.

[4] The multiagent AI system produces a final judgment covering product category, reason validity, and issue severity, and applies a predefined mapping table to link these factors to corresponding refund ratios. To limit financial risk, the system was designed to draw from predefined fixed refund ratios rather than dynamically adjusting them, while the verification process is fully automated by AI.

[5] In accordance with the research agreement, all customer records were de-identified and associated only with pseudonymous hashed identifiers; no personally identifiable information was accessed in this study.

[6] Among initial requests in the treatment group, approximately 80% were eligible for AI adjudication (the remainder were ineligible due to insufficient evidence or ambiguous intent). The AI approved roughly 80% of the eligible requests, resulting in approximately 60%–70% of treatment customers receiving a partial refund offer. In cases where the AI did not generate an offer, treatment group customers were directed to the standard return process without being aware of AI involvement.

[7] In our partner platform focusing on lifestyle, fashion, and small consumer electronics, cross-border reverse-logistics costs and resale-related costs are high relative to the average product value ($27 in our experiment), and not all returned items can be resold as new, implying negligible salvage value. For platforms with substantial salvage value (e.g., furniture and large electronics), forgoing physical returns entails greater opportunity costs from lost recovery, which may lead to lower offered refund levels.

[8] We do not observe item-level product costs or margins and therefore use net transaction value as our primary metric (Net Transaction Value = Transaction Value − Refund Amount − Processing Cost), which is also the platform's key operational indicator. The null treatment difference in initial transaction value (column 1) provides an additional balance check that is consistent with successful randomization, because treated customers were unaware of the policy when placing their initial orders.

[9] A likelihood ratio test comparing logistic models of customers' initial acceptance decisions with and without the offered refund ratio yields a chi-square statistic of 1,358 ($p < 0.001$), suggesting that the refund ratio is a key driver of customer response.

[10] These correlations may reflect underlying differences in characteristics such as product categories or return reasons. We formally address this endogeneity using the IV approach detailed in Table B1 (columns 1 and 2) of Online Appendix B.

[11] Scaling this effect to the U.S. retail industry, where return-related activities produced approximately 24 million metric tons of CO_2 emissions in 2022 (Optoro, Inc. 2023), implies an environmental benefit of more than 100 million tree-equivalents annually. Specifically, based on Long and Liu (2025), we assume 0.4 kg of CO_2 emissions per return. A mature tree absorbs approximately 22 kg of CO_2 per year (Source: https://greenerearthalliance.com/how-do-trees-absorb-carbon-how-much/). Thus, a 10% reduction in the 24 million metric tons of CO_2 emissions associated with product returns equals the sequestration of $24{,}000{,}000{,}000 \times 0.1/22 \approx 109{,}090{,}909$ mature trees annually (von Zahn et al. 2025).

[12] Individual offer customization naturally introduces heterogeneity, and further optimization of such personalized offers remains a promising direction for future research.

References

Abdulla H, Abbey JD, Ketzenberg M (2022) How consumers value retailer's return policy leniency levers: An empirical investigation. *Production Oper. Management* 31(4):1719–1733.

Abdulla H, Ketzenberg M, Abbey JD (2019) Taking stock of consumer returns: A review and classification of the literature. *J. Oper. Management* 65(6):560–605.

Altug MS, Aydinliyim T, Jain A (2021) Managing opportunistic consumer returns in retail operations. *Management Sci.* 67(9): 5660–5678.

Anderson ET, Hansen K, Simester D (2009) The option value of returns: Theory and empirical evidence. *Marketing Sci.* 28(3):405–423.

Autor D (2024) Applying ai to rebuild middle class jobs. NBER Working Paper No. 32140, National Bureau of Economic Research, Cambridge, MA.

Bower AB, Maxham JG III (2012) Return shipping policies of online retailers: Normative assumptions and the long-term consequences of fee and free returns. *J. Marketing* 76(5):110–124.

Brynjolfsson E, Li D, Raymond L (2025) Generative AI at work. *Quart. J. Econom.* 140(2):889–942.

Che Y-K (1996) Customer return policies for experience goods. *J. Industrial Econom.* 44(1):17–24.

Chu W, Gerstner E, Hess JD (1998) Managing dissatisfaction: How to decrease customer opportunism by partial refunds. *J. Service Res.* 1(2):140–155.

Costello JP, Bechler CJ (2026) Just keep it: When and why returnless product returns foster brand support. *J. Marketing Res.* 63(1): 191–209.

Davis S, Hagerty MR, Gerstner E (1998) Return policies and the optimal level of "hassle". *J. Econom. Bus.* 50(5):445–460.

De Matos CA, Henrique JL, Rossi CAV (2007) Service recovery paradox: A meta-analysis. *J. Service Res.* 10(1):60–77.

Dell'Acqua F, McFowland E III, Mollick ER, Lifshitz-Assaf H, Kellogg KC, Rajendran S, Krayer L, et al. (2026) Navigating the jagged technological frontier: Field experimental evidence of the effects of artificial intelligence on knowledge worker productivity and quality. *Organ. Sci.* 37(2):403–423.

Durante Z, Huang Q, Wake N, Gong R, Park JS, Sarkar B, Taori R, et al. (2024) Agent AI: Surveying the horizons of multimodal interaction. Preprint, submitted January 7, https://arxiv.org/abs/2401.03568.

Ertekin N (2018) Immediate and long-term benefits of in-store return experience. *Production Oper. Management* 27(1):121–142.

Goldfarb A, Tucker CE (2019) Digital economics. *J. Econom. Literature* 57(1):3–43.

Hadero H (2024) Many retailers offer 'returnless refunds.' Just don't expect them to say for which products. Accessed August 31, 2026, https://apnews.com/article/customer-returns-refunds-keep-items-amazon-54d94df2f40866020300685db6dd32f9.

Hart CW, Heskett JL (1990) The profitable art of service recovery. *Harvard Bus. Rev.* 68(4):148–156.

Holman J (2024) Returns are a headache. More retailers are saying, 'just keep it.' Accessed August 31, 2026, https://www.nytimes.com/2024/11/15/business/returns-online-shopping.html.

Janakiraman N, Syrdal HA, Freling R (2016) The effect of return policy leniency on consumer purchase and return decisions: A meta-analytic review. *J. Retailing* 92(2):226–235.

Krakowski S, Haftor D, Luger J, Pashkevich N, Raisch S (2026) Human-centered artificial intelligence: A field experiment. *Management Sci.* 72(1):57–72.

Long J, Liu J (2025) Hidden footprints in reverse logistics: The environmental impact of apparel returns and carbon emission assessment. *Sustainable Futures* 10:101360.

Moorthy S, Srinivasan K (1995) Signaling quality with a money-back guarantee: The role of transaction costs. *Marketing Sci.* 14(4):442–466.

Nageswaran L, Cho SH, Scheller-Wolf A (2020) Consumer return policies in omnichannel operations. *Management Sci.* 66(12):5558–5575.

Ni X, Wang Y, Feng T, Lu LX, Wang Y, Zhou C (2026) Generative AI in action: Field experimental evidence from Alibaba's customer service operations. Preprint, submitted February 8, https://arxiv.org/abs/2603.29888.

NRF and Happy Returns (2024) 2024 consumer returns in the retail industry. Accessed August 31, 2026, https://nrf.com/research/2024-consumer-returns-retail-industry.

Ofek E, Katona Z, Sarvary M (2011) "bricks and clicks": The impact of product returns on the strategies of multichannel retailers. *Marketing Sci.* 30(1):42–60.

Optoro, Inc. (2023) 2022 impact report: Making returns better for customers, retailers and the planet. Accessed August 31, 2026, https://4771362.fs1.hubspotusercontent-na1.net/hubfs/4771362/2022%20Impact%20Report/Optoro_2022%20Impact%20Report.pdf.

Petersen JA, Kumar V (2009) Are product returns a necessary evil? Antecedents and consequences. *J. Marketing* 73(3):35–51.

Ray G, Muhanna WA, Barney JB (2005) Information technology and the performance of the customer service process: A resource-based analysis. *MIS Quart.* 29(4):625–652.

Shang G, Ghosh BP, Galbreth MR (2017) Optimal retail return policies with wardrobing. *Production Oper. Management* 26(7): 1315–1332.

Smith AK, Bolton RN, Wagner J (1999) A model of customer satisfaction with service encounters involving failure and recovery. *J. Marketing Res.* 36(3):356–372.

Su X (2009) Consumer returns policies and supply chain performance. *Manufacturing Service Oper. Management* 11(4): 595–612.

Tax SS, Brown SW, Chandrashekaran M (1998) Customer evaluations of service complaint experiences: Implications for relationship marketing. *J. Marketing* 62(2):60–76.

von Zahn M, Bauer K, Mihale-Wilson C, Jagow J, Speicher M, Hinz O (2025) Smart green nudging: Reducing product returns through digital footprints and causal machine learning. *Marketing Sci.* 44(4):954–969.

Wood SL (2001) Remote purchase environments: The influence of return policy leniency on two-stage decision processes. *J. Marketing Res.* 38(2):157–169.

Yao S, Zhao J, Yu D, Du N, Shafran I, Narasimhan K, Cao Y (2023) ReAct: Synergizing reasoning and acting in language models. Preprint, submitted October 6, https://arxiv.org/abs/2210.03629.

Zhang S, Narayandas D (2026) Engaging customers with AI in online chats: Evidence from a randomized field experiment. *Management Sci.* 72(1):73–95.

Zhang DJ, Dai H, Dong L, Qi F, Zhang N, Liu X, Liu Z, et al. (2020) The long-term and spillover effects of price promotions on retailing platforms: Evidence from a large randomized experiment on Alibaba. *Management Sci.* 66(6):2589–2609.

Lu Fang is an assistant professor in the School of Management at Zhejiang University of Finance and Economics. Her research examines platform design, artificial intelligence, and platform governance. Her work has appeared in *Management Science* and *Journal of Industrial Economics*, among other journals.

Chunhua Wu is an associate professor with the Finning Junior Professorship in Marketing at UBC Sauder School of Business, University of British Columbia. His research examines the economic and market impacts of technology. His work has appeared in journals such as *Marketing Science*.

Zhe Yuan is an assistant professor in the School of Economics at Zhejiang University. His research focuses on industrial organization, digital platforms, artificial intelligence, and platform governance. His work has appeared in *RAND Journal of Economics, American Economic Journal: Microeconomics*, and *Management Science*, among other journals.

Kaifu Zhang is the founder of Stealth Mode Startup and was the head of the AI Initiative at the partner platform. His research focuses on the economics of artificial intelligence, business model design for online platforms, and big data and machine learning. His work has appeared in *Marketing Science, Management Science*, and *Information Systems Research*, among other journals.