Ke Cheng HSU (Clio 許可證)
Department of Economics, 5067 Lee Shau Kee Business Building
The Hong Kong University of Science and Technology
Clear Water Bay, Kowloon, Hong Kong
khsu@connect.ust.hk cliohsu1997.github.io

2 November 2026

Professor Rui Shen (申睿)
School of Management and Economics
The Chinese University of Hong Kong, Shenzhen
ruishen@cuhk.edu.cn

Dear Professor Shen,

I am writing to apply for the Post Doctoral Fellow position (PR2026/035/01) in your group at the School of Management and Economics, CUHK-Shenzhen. The lab's emphasis on *data cleaning* and on using cleaned data for research in finance, accounting, economics, and management matches how I already work, and I would like to contribute to that pipeline.

Most of my research rests on large, messy raw data that I turn into analysis-ready panels. For my job market paper on Taobao apparel return policy, I clean and link order, price, insurance, and return records so that any number traces back to the code that produced it, then estimate a structural model of purchase, insurance, and return decisions. With Professors Kohei Kawaguchi and Kei Kawai, I help build and maintain analysis-ready panels from Japanese bookstore point-of-sale records—tens of billions of records across more than 1,000 stores—merging them with product attributes scraped from the web. To turn unstructured product images and listings into analyzable variables, I use Python pipelines that combine web scraping, APIs, and AI-based OCR, with tests at each step so the cleaned data stay auditable. I organize the cleaned records in a relational database so that orders, products, prices, and returns can be linked and queried consistently across projects. I use R for structural estimation, simulation, and counterfactuals; and C++ for performance-critical tasks such as Monte Carlo simulation and iterative equilibrium solving, on SLURM and AWS.

I am comfortable both building the data infrastructure and using it for empirical research, and I would be glad to bring that combination to your group's projects. I am available from February 2027. My curriculum vitae, application form, and a representative paper are attached; two referees' contact details are listed on the CV.

Thank you for considering my application.

Sincerely,
Ke Cheng HSU

Ke Cheng HSU (Clio) 許可證

khsu@connect.ust.hk

https://cliohsu1997.github.io/

BUSINESS ADDRESS

Department of Economics, 5067 Lee Shau Kee Business Building
The Hong Kong University of Science and Technology, Clear Water Bay, Kowloon, Hong Kong

RESEARCH FIELDS

Industrial Organization, Digital Economics, Quantitative Marketing

EDUCATION

Hong Kong University of Science and Technology

Ph.D. in Economics	2022 – 2027 (expected)
M.Phil. in Economics	2020 – 2022
B.Sc. in Mathematics and Economics (Minor in Actuarial Mathematics)	2015 – 2019

WORKING PAPERS

- "Designing Return Policy in E-Commerce: A Trade-off between Purchase and Return Efficiency" *(Job Market Paper)*

  Abstract: We study the effect of return policies by asking whether they encourage efficient returns and purchases. A stricter return policy curbs inefficient ex-post returns, but the resulting lower return rates make higher product prices more attractive for sellers, leading to possible exclusion of ex-ante efficient purchases. We compare automatic coverage that reimburses return-shipping costs (RI) with optional coverage (CI) by estimating a structural model of purchase, insurance, and return decisions under monopoly pricing across six apparel categories. A lenient return policy like RI performs better in high-margin, high-salvage categories by expanding valuable purchases, whereas a strict return policy like CI performs better in thinner-margin categories by curbing excessive returns. Partial RPD, a menu pairing a discounted price with a medium restocking fee as penalty, performs better than RI and CI for most categories by better balancing purchase and return efficiency.

  Presented at (for scheduled):* APIOC 2026*; 2026 Econometric Society Australasia Meeting (ESAM 2026)*; 中國經濟學年會夏季論壇 2026*; HKUST Business School PhD Conference; CEP Brownbag in HKUST.

- "Product Design with Heterogeneous Search Costs", with Pak Hung AU

  Abstract: Building on Wolinsky's (1986) sequential search model with differentiated products, this paper

introduces search cost heterogeneity to analyze how firms leverage observable product designs to influence consumer search behavior. In equilibrium, consumers' search patterns form a partitioned structure: those with low search costs tend to seek niche products, while consumers facing higher search costs gravitate toward standardized products. In the case of two product designs, a reduction in search costs may lead to higher prices for both if the endogenous shift in consumer search modes—the composition effect—dominates the search-encouraging effect. Additionally, we demonstrate that prices respond asymmetrically to decreases in search costs based on the relative position within the distribution where the shock occurs.

Presented at: APIOC 2023; AMES 2023 NTU Singapore.

WORK IN PROGRESS

- "Pass-Through and Incidence with Product Returns", with Ce Bian and Liang Zhao

  Abstract: We study pass-through and incidence with product returns. In contrast to the case without returns, a tax such as a platform commission fee or a sales tax can create negative pass-through and improve social welfare. When a high price significantly raises consumers' return probability, the seller reduces the quantity supplied because of a lower effective margin. This downward-sloping supply curve can lower the equilibrium price in response to a tax increase, raising social surplus by reducing the return rate.

- "Too Much Unsold Stock?: Environmental Consequences of Price Commitment", with Kohei Kawaguchi and Kei Kawai

  Abstract: Resale price maintenance is procompetitive because it encourages sellers to carry inventory in the low-demand state. When realized demand falls short of that inventory, however, the seller cannot lower the price to clear the market, and the unsold copies are returned at an environmental cost. We build a dynamic oligopoly model with demand uncertainty, estimate it on monthly bookstore data, and ask how much unsold stock is due to price commitment and how a deposit-refund tax can reduce waste.

EXPERIENCE

Research

Research Assistant for Prof. Benjamin Bernard, National Taiwan University 2020

Teaching

Lecturer, Math Camp for MSc in Economics, HKUST (4.53/5) 2026

Teaching Assistant Coordinator, Business School, HKUST 2025 – 2026

Teaching Assistant, Department of Economics, HKUST (average 4.72/5) 2020 – 2025

ECON2123 – Macroeconomics (4.53/5)
ECON4134 – Economics of Uncertainty and Information (4.83/5)
ECON4999T – Economics of Uncertainty and Information (4.56/5)
ECON5210 – Microeconomic Theory I (4.85/5; 4.82/5)

FELLOWSHIPS AND AWARDS

Research Travel Grant	2023; 2026
Red Bird PhD Award	2022
Postgraduate Studentship from the HKUST	2020 – present
Academic Achievement Medal	2019
University Scholarship (730K HKD)	2015 – 2019

SKILLS

Technical Skills: R, Python, C++, LaTeX, MATLAB, AWS

Languages: English (fluent), Mandarin (native), Cantonese (basic)

REFERENCES (THESIS SUPERVISION COMMITTEE)

Pak Hung AU
Associate Professor
Department of Economics
Hong Kong University of Science and Technology
aupakhung@ust.hk

Kohei Kawaguchi
Associate Professor
Department of Economics
Hong Kong University of Science and Technology
kkawaguchi@ust.hk

Kei Kawai
Professor
Department of Economics
University of Tokyo
kawaixkei@gmail.com

許可證（Ke Cheng HSU, Clio）

khsu@connect.ust.hk
https://cliohsu1997.github.io/

通讯地址
香港九龙清水湾香港科技大学李兆基商学大楼 5067 室
经济系

研究方向
产业组织、数字经济、量化营销

教育背景

香港科技大学

经济学博士	2022 – 2027（预计）
经济学哲学硕士（M.Phil.）	2020 – 2022
数学与经济学理学学士（副修精算数学）	2015 – 2019

工作论文

- 《电商退货政策设计：购买效率与退货效率的权衡》（Designing Return Policy in E-Commerce: A Trade-off between Purchase and Return Efficiency）（求职论文）

  摘要：本文研究退货政策是否同时促进有效退货与有效购买。更严格的退货政策可以抑制事后无效退货，但退货率下降会使卖家更倾向于提高价格，从而可能把事前有效的购买排除在外。我们在六个服装品类、垄断定价下估计购买、运费险与退货决策的结构模型，比较自动报销退货运费的保障（RI）与可选保障（CI）。宽松政策（如 RI）在高毛利、高残值品类中表现更好，因为它扩大了有价值的购买；严格政策（如 CI）在毛利较薄的品类中表现更好，因为它抑制过多退货。部分退货罚金（Partial RPD）——将折扣价与中等程度补货费配对的菜单——在多数品类中比 RI 和 CI 更好地平衡购买效率与退货效率。

  报告（* 为已安排）： APIOC 2026*；2026 Econometric Society Australasia Meeting（ESAM 2026）*；中国经济学年会夏季论坛 2026*；香港科技大学商学院博士生会议；香港科技大学 CEP Brownbag。

- 《异质搜寻成本下的产品设计》（Product Design with Heterogeneous Search Costs），与 Pak Hung Au

  摘要：在 Wolinsky（1986）差异化产品序贯搜寻模型的基础上，本文引入搜寻成本异质性，分析企业如何利用可观察的产品设计影响消费者的搜寻行为。均衡中，消费者的搜寻呈现分割结构：搜寻成本较低者倾向于寻找利基产品，搜寻成本较高者倾向于标准化产品。在两种产品设计下，若搜寻模式的内生转变（构成效应）主导了鼓励搜寻的效应，搜寻成本下降可能使两种产品的价格同时上升。此外，价格对搜寻成本下降的反应取决于冲击落在成本分布中的位置，因而并不对称。

  报告： APIOC 2023；AMES 2023，新加坡南洋理工大学。

进行中的研究

- 《含退货的价格转嫁与税负归宿》(Pass-Through and Incidence with Product Returns)，与 Ce Bian、Liang Zhao

  摘要：本文研究存在产品退货时的价格转嫁与税负归宿。与无退货情形不同，平台佣金或销售税可以产生负向转嫁并提高社会福利。当高价格显著提高消费者的退货概率时，有效利润率下降，卖家减少供给。这条向下倾斜的供给曲线可以使均衡价格随税收上升而下降，并通过降低退货率提高社会剩余。

- 《滞销过多？价格承诺的环境后果》(Too Much Unsold Stock?: Environmental Consequences of Price Commitment)，与 Kohei Kawaguchi、Kei Kawai

  摘要：转售价格维持具有促进竞争的作用，因为它鼓励卖家在低需求状态下持有库存。但当实现需求低于库存时，卖家不能降价出清，未售出的商品被退回并带来环境成本。我们构建含需求不确定性的动态寡头模型，用月度书店数据估计，并考察有多少滞销来自价格承诺，以及押金—退款税如何减少浪费。

经历

研究

研究助理，Benjamin Bernard 教授，台湾大学 2020

教学

讲师，经济学硕士数学营，香港科技大学（4.53/5） 2026

教学助理协调员，商学院，香港科技大学 2025 – 2026

教学助理，经济系，香港科技大学（平均 4.72/5） 2020 – 2025

ECON2123 宏观经济学（4.53/5）
ECON4134 不确定性与信息经济学（4.83/5）
ECON4999T 不确定性与信息经济学（4.56/5）
ECON5210 微观经济理论 I（4.85/5；4.82/5）

奖助

研究差旅资助 2023；2026

Red Bird 博士奖 2022

香港科技大学研究生助学金 2020 – 至今

学术成就奖章 2019

大学奖学金（73 万港元） 2015 – 2019

技能

技术：R、Python、C++、LaTeX、MATLAB、AWS
语言：英语（流利）、普通话（母语）、粤语（基础）

推荐人（论文指导委员会）

Pak Hung Au
副教授
经济系
香港科技大学
aupakhung@ust.hk

Kohei Kawaguchi
副教授
经济系
香港科技大学
kkawaguchi@ust.hk

Kei Kawai
教授
经济系
东京大学
kawaixkei@gmail.com

☐ ☐ *(02/2020)*

香 港 中 文 大 学（深圳）
THE CHINESE UNIVERSITY OF HONG KONG, SHENZHEN
职位申请表
APPLICATION FORM

注：Note:

申请人所提供的资料将用作招聘或其他与雇佣有关事宜，并会供大学有关部门、委员会或其他处理招聘事宜的人士查阅。招聘程序完成后，未获取录的应征者資料，当无需保留时将全部销毁。

The information provided will be used for appointment and other employment-related purposes in the University. It may be accessible to offices, committees or persons who will process appointment matters. Information on unsuccessful candidates will be destroyed after the recruitment exercise when no longer required.

申请职位 Post Applied for Post Doctoral Fellow - Shen Rui Project, School of Management and Economics (博士后 - 经管学院申睿项目; Ref. PR2026/035/01)

I. 个人资料 Personal Particulars

英文姓名（姓氏先行） Name in full (English) (Surname first) Hsu, Ke Cheng		中文姓名 Name (Chinese) 許可證	* 称谓 / Title Mr. / 先生
国籍 Nationality China (Taiwan)	户口所在地 City of Residency Taipei, Taiwan, China	出生日期 Date of Birth 12 May 1997	# 护照号码 / 身份证号码 Passport No. / I.D. Card No. 353355701
通讯地址 Correspondence Address Department of Economics, 5067 Lee Shau Kee Business Building, The Hong Kong University of Science and Technology, Clear Water Bay, Kowloon, Hong Kong			
住址（如与上址不同） Residential Address (if different from above)			
电话号码 Telephone No. Mobile: (852) 6840 9753			电子邮址 E-mail Address khsu@connect.ust.hk
必要时通知 姓名 及 关系 Emergency Contact (Name & Relationship) Xu Jiayi (friend)	必要时通知 电话 及地址 Emergency Contact number & address (852) 9295 0662 香港 上環畢街 15 號 6 樓		

II. 学历 / 学术荣誉 / 奖状 (请顺序列出) Academic Qualifications / Awards / Distinctions (in chronological order)

由 From 月 / 年 Month / Year	至 To 月 / 年 Month / Year	教育机构 / 院校名称 (请列明所在国家) Institution of Learning / College / University Attended (Please specify country)	所获学历 Qualifications / Awards / Distinctions Obtained	主修科目 / 范畴 Major Subject / Field of Study
09/2015	12/2019	The Hong Kong University of Science and Technology (Hong Kong)	B.Sc.	Mathematics and Economics (Minor: Actuarial Mathematics)
08/2020	08/2022	The Hong Kong University of Science and Technology (Hong Kong)	M.Phil.	Economics
08/2022	06/2027 (expected)	The Hong Kong University of Science and Technology (Hong Kong)	Ph.D. (in progress)	Economics (IO, Digital Economics, Quantitative Marketing)

III. 工作经验（请顺序列出）Working Experience (in chronological order)

由 From 月 / 年 Month / Year	至 To 月 / 年 Month / Year	工作机构及地址 Name and Address of Employer	职位 Appointment Held (如属兼职，请予注明) (if part-time, please specify)	工作性质 Nature of Work
01/2020	07/2020	National Taiwan University Taipei, Taiwan	Research Assistant (for Prof. Benjamin Bernard)	Conducted simulations and wrote proofs for a financial-network model of a bank crisis on an interbank lending network.
08/2020	Present	Hong Kong University of Science and Technology Clear Water Bay, Hong Kong	Teaching Assistant / TA Coordinator / Lecturer (Math Camp)	TA for undergrad/grad economics (macroeconomics, information economics, microeconomic theory); TA Coordinator, Business School; Math Camp lecturer

* 请于适当空格加✓号 Please tick as appropriate # 请删去不适用者 Please delete as appropriate

IV. 专业资格 (请顺序列出) Professional Qualifications / Memberships (in chronological order)

专业团体 (全名) Professional Body (Full Name)	所获专业资格 Qualification / Membership Obtained	获取途径 (如考试、推选等) Channel of Award (e.g. exam., election)	颁发日期 Date of Award
N/A			

V. 曾担任之研究工作/出版著作 (教学/研究员申请人适用) Major Research/ Published Work (For Teaching & Research Posts only) 请连同著作撮要一并交回 Please attach abstracts of publications if available.

期间 Period	名称及性质 Name / Nature of The Work	备注 Remarks
Working paper / JMP	Designing Return Policy in E-Commerce: A Trade-off between Purchase and Return Efficiency	Structural IO / e-commerce / platform design
Working paper	Product Design with Heterogeneous Search Costs (with Pak Hung Au)	Search / product design theory
Work in progress	Too Much Unsold Stock?: Environmental Consequences of Price Commitment (with Kohei Kawaguchi and Kei Kawai)	Dynamic oligopoly
Work in progress	Pass-Through and Incidence with Product Returns (with Ce Bian and Liang Zhao)	Tax incidence under returns

VI. 咨询人 (现时 / 前雇主) Referees (Present / Previous Employers)

姓名 Name	公司名称 / 职衔 Company Name / Position	通讯地址 Correspondence Address	传真号码 / 电子邮址 Fax no. / E-mail Address
1. Pak Hung AU 2. Kohei Kawaguchi 3. Kei Kawai	Associate Professor, Dept. of Economics, HKUST Associate Professor, Dept. of Economics, HKUST Professor, Dept. of Economics, University of Tokyo	HKUST, Clear Water Bay, Kowloon, Hong Kong HKUST, Clear Water Bay, Kowloon, Hong Kong University of Tokyo, Tokyo, Japan	aupakhung@ust.hk kkawaguchi@ust.hk kawaixkei@gmail.com

除非特别注明，大学将毋须再行通知应征者并可径向上述咨询人查询。请先行知会上述咨询人。

Unless otherwise specified, consent is deemed given by the applicant to the University to approach the above referees whenever appropriate without prior notification. Please also inform your referees that such consent has been given by you.

VII. 过往薪酬纪录 Past Salary Record

期间 Period	公司名称 / 职衔 Company Name / Position	每月基本薪酬 Monthly Basic Salary	每月保险及公积金供款 Monthly Insurance and Housing Fund Contribution	其他津贴 Other Allowances
N/A	N/A	N/A	N/A	N/A

VIII. 其他 Others

1. 现职离职通知期
   Notice Period Required by Present Employer: Not Applicable
2. 如获聘任可到任日期
   Earliest Date Available if Appointed: February 2027
3. 请问从何处得悉本职位空缺？
   How did you learn about this vacant position? CUHK website
   (例：期刊 / 报章名称及日期) (e.g. Name and date of journal / newspaper, etc.)
4. 申请人如有近亲现职于香港中文大学 (深圳)，请列明其姓名、职位及与申请人的关系
   If you have close relatives currently employed by The Chinese University of Hong Kong, Shenzhen, please state their names, posts and relationship to you.

   None
5. 其他与申请此职位有关的资料
   Applicants may use the following space to provide any other relevant information in support of your application.

   Job market Paper Attached

IX. 声明 Declaration

本人谨此声明以上所提供的资料均属真实及所附证明文件皆为真确副本。如获聘任，本人将提供有关身份及资历文件的正本 / 核实真确本予大学作查核。本人明白倘若故意虚报资料或隐瞒重要事实，香港中文大学(深圳) 可取消已发出的口头或书面聘约或纵使已获聘任亦可遭解雇。

I declare that the information given above is correct and complete to the best of my knowledge and that the documents provided by me in connection with this application are true copies. I will produce the original / certified true copies of all identification and qualification documents as required by the University upon assumption of duty at the University if appointed. I understand that if I knowingly supply false information or withhold any material information, The Chinese University of Hong Kong, Shenzhen shall have the right to rescind any verbal / written offer of appointment and I shall render myself liable to dismissal if I am eventually appointed by the University.

日期
Date 2 November 2026

签署
Signature HSU Ke Cheng

Designing Return Policy in E-Commerce: A Trade-off between Purchase and Return Efficiency

Ke Cheng HSU*

November 1, 2026

Abstract

We study the effect of return policies by asking whether they encourage efficient returns and purchases. A stricter return policy curbs inefficient ex-post returns, but the resulting lower return rates make higher product prices more attractive for sellers, leading to possible exclusion of ex-ante efficient purchases. We compare automatic coverage that reimburses return-shipping costs (RI) with optional coverage (CI) by estimating a structural model of purchase, insurance, and return decisions under monopoly pricing across six apparel categories. A lenient return policy like RI performs better in a category with a high margin, high net scrap value, and high uncertainty by expanding efficient purchases, whereas a stricter return policy like CI performs better in a category with the opposite attributes by curbing excessive returns. Partial RPD, a menu pairing a discounted price with a medium restocking fee as a penalty, performs better than RI and CI for most categories by better balancing the ex-ante efficiency of purchases and the ex-post efficiency of returns.

Keywords: structural modeling, product returns, return policy, welfare analysis, online retailing.

1 Introduction

Product returns are a central challenge in online retail. Return rates have reached 19.3% in the United States and exceed 35% in China.[1] In monetary terms, total retail returns in the United States are projected to reach $849.9 billion in 2025. The problem is especially acute for apparel, whose fit and quality are difficult

*Department of Economics, School of Business and Management, Hong Kong University of Science and Technology. Email: khsu@connect.ust.hk

[1]National Retail Federation (2025), 2025 Retail Returns Landscape; Wu Dan Li (2024), Zhejiang University School of Management.

to assess online. Industry benchmarks place online apparel return rates in the United States between 20% and 40%.[2] These high rates arise in part because consumers must purchase before inspecting the product. They learn fit only after delivery and then decide whether to keep or return it. Under a typical return policy, consumers are refunded the product price but bear the cost of return shipping, while the seller incurs handling costs and recovers only part of the product's value. Because the consumer does not bear the seller's handling cost, returns can exceed the efficient level. When the expected value from keeping is not enough to offset that loss, the purchase itself can generate negative expected social surplus and be ex-ante inefficient.

By considering the efficiency of both purchases and returns, the platform can design a return policy to improve the joint welfare of consumers and sellers. A return should occur when the seller's recovery value from the returned product exceeds the consumer's value from keeping it plus return shipping and handling costs. A purchase should occur when the expected value the transaction creates exceeds the expected costs of production and returns. From a welfare perspective, the existing return literature mostly asks whether consumers return efficiently conditional on purchase (Su, 2009), but not whether the purchase itself should occur in the first place.[3] A return policy changes how costly a return is, which shifts return probabilities and the expected social surplus from purchase, so the two margins are linked. Only by evaluating purchases and returns jointly can we determine when purchase and return efficiency trade off, when both can improve together, and when each margin matters more for welfare.

Return Shipping Insurance (RSI), which is commonly used on Chinese e-commerce platforms, changes both purchases and returns by reimbursing the buyer's return-shipping expense. We examine two RSI policies: *retailer insurance* (RI), under which the seller purchases the insurance and provides automatic coverage for all buyers, and *consumer insurance* (CI), under which the consumer self-selects whether to buy optional coverage at checkout. Automatic coverage encourages inefficient returns with positive consumer benefit but negative social cost. However, it can encourage purchases that create positive value but would not occur if consumers feared paying for return shipping. On the other hand, optional coverage disciplines returns by limiting reimbursement to those who buy insurance, but consumers with a high chance of returning may be especially likely to select into it. Since a seller's profit from a sale depends on whether the product

[2] Eightx (2026), Average eCommerce Return Rate by Category; McKinsey & Company (2020), Returning to order.

[3] In Che (1996), consumers are homogeneous and all of them buy. Allowing returns can make the seller charge a higher price. The consumer keeps if and only if valuation exceeds that price. When retail cost is small, that price exceeds cost. The consumer then returns even though keeping is efficient. That loss is a suboptimal return. In a cheap-talk setting, a refund above salvage makes seller advice credible (Inderst and Ottaviani, 2013). This produces overpurchase in the seller's recommendation given the signal. It also produces over-return. The question is whether the seller should send that signal. Buyers are not screened by expected social value.

is ultimately kept or returned, each insurance policy changes expected profit at any given price and therefore the price the seller chooses. Once the seller adjusts the price, it is not obvious which arrangement creates more value, or whether that ranking changes with the margin, net scrap value, and match value uncertainty.

These observations lead to three research questions. First, how do automatic and optional return-shipping insurance affect purchases, returns, insurance selection, and seller pricing? Second, is automatic coverage superior to optional insurance, and how does the result vary across product categories? How much of the difference comes from purchases, and how much from returns? Third, can alternative policies that better balance purchases and returns create greater social value than the existing insurance policies?

In our model, a monopolistic seller sets the product price, anticipating consumers' purchase and return decisions. Consumers only know their expected value of a product when deciding whether to buy and learn their final match value after delivery, when they decide whether to keep the item. Under either policy, consumers receive a full refund of the product price upon return, and insurance reimburses return-shipping fees. The seller pays production cost on every sale, earns the price when the product is kept, and recovers scrap value net of return-handling costs on return.

Our model generates direct behavioral effects and an indirect effect through the equilibrium price response. At a given product price, RI raises willingness to purchase more broadly because every buyer receives coverage without paying a separate premium, but it also lowers the return cost of every purchaser. Under CI, the premium limits coverage and participation, but uninsured purchasers bear the return shipping costs. Equilibrium price effects are ambiguous because three forces interact: the seller-paid RI premium creates cost pass-through, the price itself changes the total volume of returns through who purchases and whether those purchases are returned, and the number of purchases ultimately kept affects the profitability of raising the price. The last force is *surplus extraction*: setting a higher price becomes more attractive when fewer purchasers return the product.

We formalize the welfare consequences of these two effects along the ex-ante and ex-post margins. *Ex-ante efficiency* concerns whether purchases create positive expected social surplus. Broader participation under RI can mitigate monopoly underprovision when marginal buyers are socially valuable, but it can also create overprovision by attracting frequent returners whose purchases impose more cost than value. *Ex-post efficiency* concerns whether a purchased product is returned only when doing so creates social value after accounting for shipping, hassle, handling, and recovery value. Return policies should therefore preserve valuable purchases while discouraging inefficient purchases and excessive returns. RI tends to perform

better when the margin, net scrap value, and uncertainty are high. High margin and net scrap value, and the option value arising from high uncertainty (Anderson et al., 2009), raise the expected social surplus from a purchase; weaker surplus extraction and automatic reimbursement of return shipping enable RI to capture more of those purchases. High uncertainty also lowers the ex-post cost of a lenient return policy, because reimbursement shifts return rates little when returns depend mainly on realized match value.

We study these tradeoffs with transaction data from a mid-sized women's clothing retailer on Taobao that repeatedly switched between RI and CI during 2017–2023. Within-retailer switches provide policy variation that is not confounded by retailer identity. For products sold under both policies, RI is associated with lower prices and higher return rates. Under CI, insured transactions return much more often than uninsured ones, consistent with frequent returners selecting into coverage and a behavioral response to reimbursement. Price comparisons vary with product composition and aggregation, which motivates a framework in which prices, purchases, insurance, and returns adjust jointly.

We estimate a structural model of these decisions under monopoly pricing for six apparel categories. The estimates recover ex-ante heterogeneity and ex-post match value uncertainty, along with estimated seller production cost and calibrated net scrap value, to evaluate equilibrium welfare. The model reproduces the main qualitative differences between the policies. It underpredicts purchases by a similar amount under both policies and modestly overstates the return gap that matters for welfare, so we rank policies within each category and discuss how this bias overstates CI's advantage. To measure how much of the welfare loss reflects distortion from market structure that return-policy design cannot correct, we study a benchmark in which returns are socially efficient while the monopoly seller still sets the price. About 24%–29% of first-best welfare remains unrealized because prices exclude efficient purchases.

For dresses, the margin, net scrap value, and uncertainty are high. Both RSI underprovide by leaving purchases that produce positive social surplus. RI performs better by expanding efficient purchases, reaching 76% of first-best welfare compared with 74% under CI. For outerwear, pants, sweaters, and two-piece sets, margins are low to medium, and net scrap value and uncertainty are medium. CI underprovides, whereas RI overprovides for pants and sweaters and slightly underprovides for outerwear and two-piece sets. CI performs better in these categories, with much higher ex-post efficiency and slightly lower ex-ante efficiency. For tops, the margin is intermediate, both net scrap value and uncertainty are low. RI performs worse than CI on both efficiency margins through allowing ex-ante inefficient purchases and ex-post over-return, leading to a loss of about 51% of first-best welfare.

The third research question considers a counterfactual self-selective price-refund menu where the buyer self-selects between a high price with a high refund and a low price with a low refund. In the literature, this menu is shown to increase seller profit (Najafi et al., 2024; Altug et al., 2021). We call this menu *return-penalty discount (RPD)* and study its welfare impact. In contrast to the literature, where the seller sets the full price-refund menu, the platform sets the return penalty. We let the platform choose the return-penalty design because that matches platform practice.[4] Nevertheless, RPD is not necessarily more efficient. Stricter refund terms can strengthen surplus extraction and exclude socially valuable buyers. If the penalty is too severe, it can also deter returns that would be socially efficient. We evaluate both full forfeiture of the refund and partial forfeiture implemented as restocking fees.

Full-forfeiture RPD shows why stronger return discipline is not enough by itself. In most categories, purchases concentrate on the discounted branch, full-price demand is negligible, and return probability on the discount branch is near zero. The policy therefore replaces over-return under RSI with under-return, deterring returns that would be socially efficient. Even so, it substantially reduces ex-post inefficiency. It also strengthens surplus extraction and raises the purchase cutoff, excluding marginal consumers, even though the average price paid is lower. The reduction in return inefficiency dominates this participation loss relative to RI outside dresses. Outside dresses, full-forfeiture RPD outperforms RI by about 8–34 percentage points of first-best welfare and outperforms CI by about 4 percentage points on sweaters and is roughly tied on the rest. For dresses, it falls short of both RSI policies by roughly 18–21 percentage points because many buyers that would have created positive social surplus are excluded.

Partial-forfeiture RPD softens both failures of full forfeiture. Restocking fees of 10%–15% of the discounted price preserve much of the return discipline while reducing under-return and restoring participation. These fees outperform both RSI policies outside the dress category, attaining about 75%–78% of first-best welfare. For dresses, these fees restore much of what full forfeiture loses (raising welfare from about 56% of first best under full forfeiture to about 69%–70%) but remain below RI (about 76%). At low net scrap value, the partial-forfeiture advantage outside dresses survives. At high net scrap value, partial RPD remains ahead for sweaters and tops and below CI by at most 1.5 percentage points for the other categories. This evidence favors a category-specific return policy. A more lenient return policy, RI, should be provided to dresses, the category with high margin, high net scrap value, and high uncertainty, whereas a more strict policy, such as

[4] In practice, refund amounts are often constrained by platform rules or law. Best Buy charges $45 on opened phones and 15% on opened cameras, drones, and projectors. Amazon permits merchant-fulfilled sellers to charge a restocking fee of up to 20% for change-of-mind returns received after the return window (Amazon Seller Central Help).

full forfeiture RPD, should be provided to categories with opposite attributes. When tailoring is infeasible, moderate partial-forfeiture RPD may serve as a practical alternative.

1.1 Literature Review

Many existing empirical studies of return policy examine, from the seller's perspective, a trade-off between stimulating purchase and managing return cost. For purchases, return leniency raises current purchases (Wood, 2001), and a completed return raises later demand and lifetime value (Petersen and Kumar, 2009, 2015; Bower and Maxham, 2012; Griffis et al., 2012). For costs, extra reverse-logistics cost can exceed extra orders (Patel et al., 2021), profit falls when the seller pays return shipping (Hjort and Lantz, 2016), and free shipping promotions raise returns enough to cut profit (Shehu et al., 2020). A meta-analysis finds that different return-policy dimensions, such as time or scope, affect purchases and returns differently (Janakiraman et al., 2016). Although this literature generally treats additional purchases as beneficial to the seller, they can reduce the joint surplus of consumers, the seller, and the insurer when return probabilities are high and margins are thin. Under RI, such overprovision can arise because the seller does not fully internalize insurer claim losses. To our knowledge, no empirical study asks whether return policy causes under- or overprovision or compares this distortion across policies. We use two measures. Ex-ante inefficiency is the surplus lost from distorted participation relative to first best when returns are efficient. The purchase-efficient gap is the actual purchase cutoff minus the zero-surplus cutoff under actual return behavior. Positive values indicate underprovision and negative values overprovision. By this measure, CI and full-forfeiture RPD underprovide in every category. RI underprovides for dresses and slightly underprovides for outerwear and two-piece sets, and it overprovides for pants, sweaters, and tops.

Price determines who participates, and greater return leniency can raise or lower the equilibrium price. A higher refund raises willingness to pay and allows the seller to extract option value (Su, 2009). A money-back guarantee can also support a higher price by screening keepers (Che, 1996). Allowing returns can make demand less elastic and raise prices (Newberry and Zhou, 2024). More generally, sellers may pass return-related costs through to consumers in product prices, including the seller-paid insurance premium under RI (Ren et al., 2021; Li et al., 2023b).[5] Conversely, a full refund may induce the seller to cut price to deter trial and return (Davis et al., 1995). One channel missing from the return literature is the net-demand channel for surplus extraction. In our setting, a more generous policy can thin net demand, weakening surplus extraction and putting downward pressure on price.

[5]These RI–CI models hold each buyer's return probability fixed.

On the ex-post margin, returns need not be socially wasteful. Su (2009) shows that when the refund equals salvage, a return reallocates the good to the party with the higher value. Later work explains why the seller's optimal refund departs from this benchmark. The refund exceeds salvage when the seller compresses information rents (Courty and Hao, 2000; Matthews and Persico, 2007) or makes sales talk credible (Inderst and Ottaviani, 2013). Manufacturer buy-backs can correct a retailer's excessively strict refund (Shulman et al., 2010). The refund falls below salvage to deter clearance deferrals (Altug and Aydinliyim, 2016), to discourage returns or extract surplus from returners (Shulman et al., 2009; Shang et al., 2017a), or to deter switching to a rival (Shulman et al., 2011). With opportunistic renters, the optimal refund can lie on either side of salvage depending on rental value, salvage, and sales volume (Altug et al., 2021). These studies explain departures from the socially efficient refund but do not empirically quantify or compare return inefficiency across policies. We measure that distortion as ex-post inefficiency, the welfare lost when purchasers do not return efficiently, holding fixed who bought. In our counterfactuals, ex-post inefficiency is smallest under RPD, larger under CI, and largest under RI in every category.

Theory on RSI compares RI and CI mainly by seller profit, with some analyses of consumer surplus and social welfare. Li et al. (2023b) distinguish informed buyers, who never return, from uninformed buyers, who return at a fixed rate. RI is more profitable when handling costs are low, whereas CI can benefit the retailer, insurer, and consumers when uninsured buyers anticipate regret. Li et al. (2021) find that RI dominates when the premium is low and the refund equals salvage, while CI may leave more surplus with buyers. Low premiums and high compensation also favor RI in Fan and Chen (2020), although adding RSI to a money-back guarantee need not raise retailer profit (Ren et al., 2021). With quality signaling, RSI can harm both high- and low-quality sellers while benefiting consumers, and higher-rated sellers are more likely to adopt RI (Zhang et al., 2022).[6] We instead evaluate RI and CI empirically in terms of retailer profit, expected consumer surplus, and expected social surplus, showing how their ranking varies with margins and handling costs.

Return-policy design uses both nonmonetary and monetary penalties (Abdulla et al., 2019). Early models deter opportunistic returns through hassle under a full money-back guarantee (Davis et al., 1998) or nonrefundable shipping charges (Hess et al., 1996). Later models jointly choose prices and partial refunds (Shulman et al., 2009; Hsiao and Chen, 2012) or offer self-selective screening menus (Shang et al., 2017a; Altug et al., 2021; Najafi et al., 2024; Huang and Zhang, 2020). More recently, Costello and Bechler (2026)

[6]Other empirical work on RSI is scarce. It includes Chen et al. (2018) on how RSI changes purchase and return intentions.

and Fang et al. (2026) examine returnless policies that refund the buyer in full or in part while letting them keep the product. Empirical work on return policy finds a large return option value in catalog apparel (Anderson et al., 2009) but a small willingness-to-pay premium for a full refund in electronics (Shang et al., 2017b). The closest empirical study to ours is Newberry and Zhou (2024), who find that the insurance value of returns accounts for roughly 80% of consumer surplus gains, compared with 20% for quality signaling, which supports abstracting from the signaling channel. They estimate the welfare effects of uniform return policies and show that repricing benefits the seller but harms consumers, because sellers pass return costs through and demand is less elastic. By contrast, we study self-selection under CI and RPD rather than a uniform return policy.

To keep the model parsimonious, we abstract from pre-purchase search (Petrikaitė, 2018; Jerath and Ren, 2024; Janssen and Williams, 2024; Li et al., 2024; Ibragimov et al., 2026), purchase-return interdependence over customer lifetimes (Chen, 2021; Montaguti et al., 2026), quality signaling (Moorthy and Srinivasan, 1995; Zhang et al., 2022; Li et al., 2023a), inter-firm competition over return policy, such as money-back guarantees, restocking fees, or RI (Chesnokova, 2007; McWilliams, 2012; Shulman et al., 2011; Chen and Grewal, 2013; Inderst and Tirosh, 2015), and supply chain buyback contracts (Padmanabhan and Png, 1997; Tran et al., 2018).

Our main contribution is an empirical welfare decomposition that separates losses from inefficient participation, when inefficient purchases occur or efficient purchases are deterred, from losses due to suboptimal return decisions. By comparing these two margins across policies and categories, the framework identifies when they can improve together and when they trade off. In the latter case, we show that whether to prioritize participation or return discipline depends on margins and return-handling costs. Applied to RSI and RPD, the framework provides, to our knowledge, the first empirical evaluation in online retail of both RSI and a self-selective price-refund menu, measuring seller profit, expected consumer surplus, and expected social surplus.[7] These results can guide platforms on when to tighten or relax return policy and support a moderate restocking fee in that menu as a practical alternative across most apparel categories when tailoring requires information the platform lacks.

The remainder of the paper proceeds as follows. Section 2 describes the institutional setting of return shipping insurance, introduces the transaction data, and documents key stylized facts on prices and returns.

[7] Vollmer (2025) structurally compares resale with a single partial refund in college football tickets, reporting primary-seller profit, consumer welfare, and total welfare. The seller posts one refund per game, and a holder returns if realized value falls below it. Buyers do not choose among refund contracts, and the refund is not paired with an upfront discount.

Section 3 sets up the structural model, and Section 4 characterizes equilibrium purchase, insurance, and return decisions, seller pricing incentives, and the welfare framework. Sections 5 and 6 discuss identification, estimation, and calibration. Section 7 reports counterfactual welfare simulations comparing RI, CI, and return-penalty discount (RPD) by category. Section 8 concludes.

2 Institutional Setting and Data

Apparel accounts for the largest share of online retail in China at 22%.[8] Because consumers cannot inspect garments before delivery, return rates are high, historically generating frequent and costly return-shipping disputes.

2.1 Return Shipping Insurance: RI versus CI

To resolve frequent and costly return-shipping disputes, Taobao and Huatai Insurance introduced Return Shipping Insurance (RSI) in 2010 to standardize reimbursement. The RSI market reached approximately 8 billion RMB in 2022.[9] The insurance market on Taobao is centralized. The platform partners with a small set of insurers and assigns providers algorithmically, so coverage terms are uniform across merchants. On approved returns, buyers receive a full refund of the product purchase price from the seller and reimbursement of the 1 kg delivery fee from the insurer (typically 5 to 25 RMB).

RSI operates under two mutually exclusive regimes. Under RI, the seller enrolls at the store level, without a separate consumer fee. The insurer charges the seller a per-transaction premium (typically 1 to 5 RMB) based on the store's three-month historical return rate and category benchmarks. RI dominates the platform, accounting for over 90% of all RSI policies. Under CI, which is offered *if and only if* the seller has not enrolled the store in RI, individual buyers choose whether to purchase coverage as an add-on at checkout. In contrast to RI, the CI premium is consumer-specific, determined mainly by the buyer's province of residence and return history, *independent of the product*. CI premiums typically exceed RI premiums because optional coverage attracts consumers who anticipate returning, and insurers charge higher rates to buyers with higher return propensities.

2.2 Data and Sample Construction

Our dataset comprises transaction records from a single mid-sized women's clothing retailer on Taobao, spanning from 2017 to 2023. Each transaction falls under one RSI policy. Over this period, total revenue is

[8] Ministry of Commerce (2024), *2023 China Online Retail Market Development Report*.
[9] Shanghai Securities News (2022), Sina Finance.

approximately 62 million RMB.

The data include sales, returns, and insurance at the transaction level, with consumer information limited to addresses. Insurance premium and coverage are observed for nearly all RI transactions, excluding international orders and buyers ineligible because of high historical return rates. Under CI, coverage is recorded only when the consumer buys insurance, while for non-purchasers we impute it from same-province addresses. Premiums are not recorded, so we impute the premium from coverage using the 2025 premium-coverage schedule (Online Appendix A).

We code a return as a binary indicator equal to one when the buyer receives a strictly positive refund and zero otherwise, excluding partial refunds for minor product defects. We keep single-item, single-quantity orders with confirmed shipment and no premium shipping, and drop non-product listings and promotion days on which the store temporarily switched RSI policy. Restricting to one item per order rules out complementarities between purchase and return decisions across products, while restricting to quantity one rules out bracketing, in which buyers order multiple sizes or colors of the same listing and return all but one. We keep six apparel categories with at least 50 transactions per month on average: dress, tops, pants, sweater, two-piece, and outerwear. The cleaned sample contains 361,201 transactions (238,612 under RI and 122,589 under CI).

2.3 Stylized Facts

Figure 1 plots 30-day moving averages of the daily return rate from 2017 through 2023 across RI and CI periods. Over these six years, the store repeatedly switched between the two policies. These switches do not seem to be based on the return rate. They occur when returns are low or high and when the rate is rising or falling. Separately, the return series drifts upward over the sample and exceeds 40% toward the end of 2023 during a long period of RI. This late-sample climb plausibly coincides with the rapid expansion of *live-stream* commerce on Taobao, where purchases are typically more impulsive and return-prone.[10]

Table 1 reports prices, premiums, coverage, and return rates under RI and under CI, pooling across categories.

Average list prices are higher under CI than under RI (about 92 versus about 87 RMB). Traditional premium pass-through would predict the opposite. The seller-paid RI premium is a visible cost that could raise list prices under RI. We do not observe higher RI prices, suggesting that other pricing forces operate. One candidate force is surplus extraction. Under CI, uninsured buyers face higher private return costs and

[10] Removing a pooled annual cycle (day-of-year harmonics 1–4) prior to the same 30-day moving average yields a broadly similar pattern, including the pronounced late-sample increase, and still shows no clear ranking of return rates between RI and CI periods.

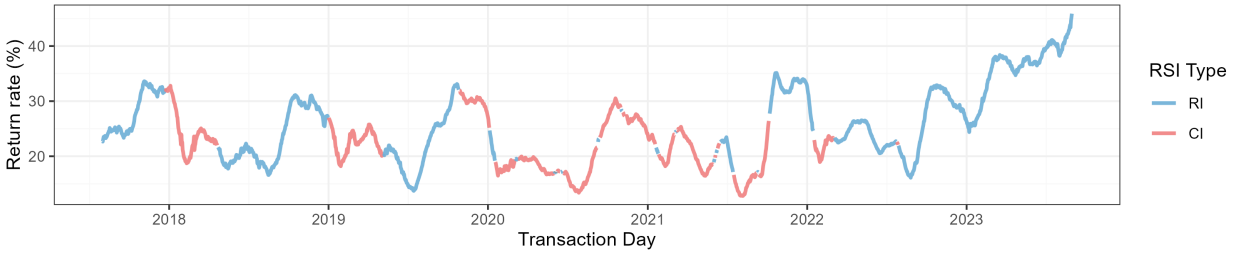


Figure 1: 30-day moving average of the daily return rate (%), 2017–2023, with RI and CI periods shaded.

Table 1: Summary Statistics

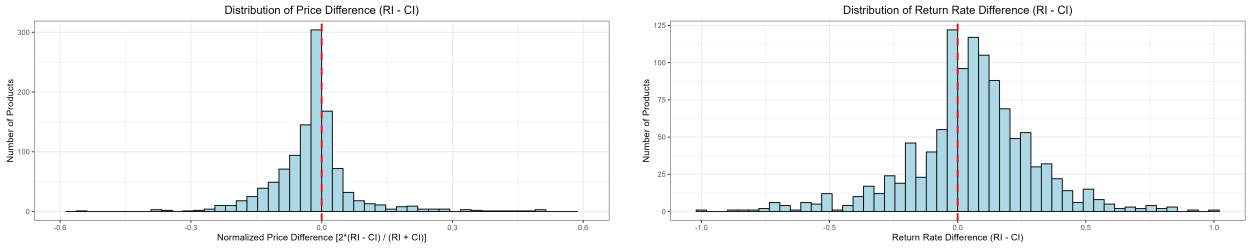
RSI	Variable	Unit	N	Mean	Sd	Min	Median	Max
CI	Price	RMB	122,589	91.50	56.67	18.0	83.65	1038.0
	Premium	RMB	122,589	5.01	0.22	4.6	5.10	7.1
	Coverage	RMB	122,589	10.63	1.46	8.0	11.00	25.0
	Return Rate	%	122,589	21.10	40.80	0.0	0.00	100.0
RI	Price	RMB	238,612	87.36	67.58	15.3	79.00	738.0
	Premium	RMB	238,612	2.67	0.39	1.4	2.50	3.3
	Coverage	RMB	238,612	10.59	1.45	8.0	11.00	25.0
	Return Rate	%	238,612	27.30	44.55	0.0	0.00	100.0

Notes: Return-rate medians are zero because the outcome is binary.

return less often, which can strengthen the seller's incentive to set higher prices.

Average return rates are higher under RI than under CI, with pooled means of about 27% versus about 21%. Under CI, insurance is purchased on 2.91% of transactions. Return rates are much higher among insured than uninsured transactions (46.3% versus 20.4%). The gap is consistent with lower private return costs among insured buyers and adverse selection into coverage.

Pooled comparisons may be confounded by product selection. Figure 2 restricts the sample to 1,140 products that sold under both RI and CI and have more than five transactions per product. Panel (a) plots the normalized price difference $2(\bar{p}^{\text{RI}} - \bar{p}^{\text{CI}})/(\bar{p}^{\text{RI}} + \bar{p}^{\text{CI}})$, where $\bar{p}$ is the product-level mean transaction price under each policy. As in the pooled sample, a larger share of products still show higher prices under CI than under RI. Panel (b) reports $\bar{r}^{\text{RI}} - \bar{r}^{\text{CI}}$, where $\bar{r}$ is the product-level mean return rate under each policy. Most mass lies to the right of zero, indicating higher return rates under RI.



(a) Normalized price difference. (b) Return-rate difference.

Figure 2: Within-product price and return-rate differences (RI − CI).

Taken together, the time series documents repeated within-store policy switches. Pooled and within-

product comparisons show that a larger share of products have higher prices under CI than under RI, and that return rates are higher under RI once products are held fixed.

3 Model

We consider a retail shopping model with RSI in which a monopolistic seller offers apparel category j in period t. The seller faces downward-sloping residual demand, with the outside option absorbing purchases from rivals. Market size N_{jt} varies across categories and seasons (Online Appendix A). Categories are separable, with no cross-category demand spillovers. For notational simplicity, we suppress the category subscript j and period subscript t throughout Sections 3 and 4, reintroducing them in Section 5 for identification across categories and markets.

The game proceeds in three stages. First, the seller sets the retail price p, taking the RSI policy and insurance premiums as given. Second, each consumer decides whether to purchase the product and, under CI, whether to purchase optional insurance. Third, after delivery, consumers observe their idiosyncratic match value shocks and choose whether to keep or return the product.[11]

The insurer provides consumer-specific coverage k_i according to the location-specific 1 kg delivery fee. Under RI, the seller pays premium I^r per transaction, and coverage is automatic for all buyers. Under CI, buyers choose whether to purchase coverage at checkout for premium $I^c(k_i)$.

The seller produces each unit at marginal cost $\theta^c = X\alpha^c + \zeta^c$. The store-level shock ζ^c is common across products and reflects shared inputs such as marketing, customer acquisition and service, labor, utilities, and rent. If a unit is returned, the seller incurs a handling fee and recovers resale value, yielding net return handling cost h. This cost is negative when recovery exceeds the handling cost.

Consumers differ in ex-ante valuation ν_i and location-specific return-shipping coverage k_i. Ex-ante valuation is

$$\nu_i = \mu + \xi + \omega\epsilon_i, \tag{1}$$

where $\mu = X\beta^\mu$ is mean valuation from observed category attributes X, ξ is an unobserved demand shock orthogonal to X, ω is the standard deviation of idiosyncratic ex-ante tastes within the category and is time-invariant, and $\epsilon_i \sim \mathcal{N}(0,1)$, so $\nu_i \sim \mathcal{N}(\mu + \xi, \omega^2)$ with distribution F_ν. Return-shipping coverage k_i is drawn from a location-specific distribution $\Psi(\cdot)$, independent of ν_i.

[11] Hong and Pavlou (2014) find that, for experience goods, returns are driven more by product-fit uncertainty than by quality uncertainty.

After delivery, the idiosyncratic match value shocks of consumer i, $\varepsilon_i^{\text{keep}}$ and $\varepsilon_i^{\text{return}}$, are realized. They follow a Type I extreme value distribution with mean zero and scale $1/\sigma$, with ex-post precision $\sigma = X^{\sigma}\beta^{\sigma}$.[12] Conditional on purchase, consumer i chooses between keeping and returning. On a return, the buyer receives a refund of p, incurs hassle κ and shipping k_i, and is reimbursed only for k_i when coverage applies. Realized utility on each branch is the sum of mean indirect utility and the idiosyncratic shock, $u_i = \overline{u}_i + \varepsilon_i$. Mean utilities under the respective policies are

$$\overline{u}_i^{\text{keep}} = v_i - p - I^c(k_i)BC_i, \tag{2}$$

$$\overline{u}_i^{\text{return}} = -\kappa - k_i(1 - RSI_i) - I^c(k_i)BC_i, \tag{3}$$

where BC_i is an indicator that consumer i purchases CI coverage, and RSI_i is an indicator of insurance, either under RI or under a CI purchase. Under RI, coverage is automatic, so return shipping is reimbursed up to k_i, yielding $\overline{u}_i^{\text{return}} = -\kappa$. Under CI, an insured buyer receives reimbursement and obtains $\overline{u}_i^{\text{return}} = -\kappa - I^c(k_i)$, whereas an uninsured buyer bears the shipping cost, obtaining $\overline{u}_i^{\text{return}} = -(\kappa + k_i)$.

4 Equilibrium Analysis

We focus on equilibrium within a single category and period. Superscript l denotes the RSI scheme (r for RI, c for CI). We analyze consumer behavior under each RSI scheme and then compare the mechanisms that matter for pricing and welfare.

4.1 Consumer Behavior under RSI

Under RI, automatic coverage yields mean utilities $\overline{u}_i^{\text{keep}} = v_i - p$ and $\overline{u}_i^{\text{return}} = -\kappa$. Integrating over Type I extreme value shocks yields expected consumer surplus

$$CS_i^r = \frac{1}{\sigma} \ln\left(e^{\sigma(v_i - p)} + e^{-\sigma\kappa}\right). \tag{4}$$

Consumers purchase if $CS_i^r \geq 0$, or $v_i \geq \hat{v}^r \equiv p + \frac{1}{\sigma}\ln(1 - e^{-\sigma\kappa})$. Conditional on purchase, the return probability is

$$q_i^r \equiv \frac{e^{-\sigma\kappa}}{e^{\sigma(v_i - p)} + e^{-\sigma\kappa}}. \tag{5}$$

High uncertainty (low σ) raises the option value, increasing both consumer surplus and return probability.

Under CI, consumers choose between purchasing with optional insurance at premium $I^c(k_i)$, purchasing without insurance, and not purchasing. Expected consumer surplus under the two purchase options is

[12] The variance is $\pi^2/(6\sigma^2)$.

$$CS_i^{bc} = \frac{1}{\sigma}\ln\left(e^{\sigma(\nu_i - p)} + e^{-\sigma\kappa}\right) - I^c(k_i), \qquad CS_i^{n} = \frac{1}{\sigma}\ln\left(e^{\sigma(\nu_i - p)} + e^{-\sigma(\kappa + k_i)}\right). \tag{6}$$

A purchaser prefers insurance if $CS_i^{bc} \geq CS_i^n$, which holds when valuation is below the insurance-indifference cutoff

$$\nu_i \leq \tilde{\nu}^{bc} \equiv p - \kappa + \frac{1}{\sigma}\ln\left(\frac{e^{-\sigma I^c(k_i)} - e^{-\sigma k_i}}{1 - e^{-\sigma I^c(k_i)}}\right). \tag{7}$$

Because return probability falls with valuation, insurance is valued most by lower-valuation buyers.

Setting $CS_i^{bc} \geq 0$ and $CS_i^n \geq 0$ yields purchase thresholds $\hat{\nu}^{bc} \equiv p + \frac{1}{\sigma}\ln\left(e^{\sigma I^c(k_i)} - e^{-\sigma\kappa}\right)$ and $\hat{\nu}^n \equiv p + \frac{1}{\sigma}\ln\left(1 - e^{-\sigma(\kappa + k_i)}\right)$. These cutoffs characterize optimal choice: consumers purchase with insurance if $\hat{\nu}^{bc} \leq \nu_i < \tilde{\nu}^{bc}$, purchase without insurance if $\nu_i \geq \max\{\tilde{\nu}^{bc}, \hat{\nu}^n\}$, and do not purchase otherwise. An insured segment exists when $\tilde{\nu}^{bc} \geq \hat{\nu}^{bc}$,[13] which holds when coverage k_i is large relative to the premium $I^c(k_i)$ and hassle cost is not too large. If hassle cost is too large, return probability is too low for the option value of insurance to outweigh the premium, so all purchasers remain uninsured.

Insured consumers return with probability q_i^r, as under RI, because the premium is sunk at purchase. Uninsured consumers bear the full return cost k_i and return with lower probability

$$q_i^n \equiv \frac{e^{-\sigma(\kappa + k_i)}}{e^{\sigma(\nu_i - p)} + e^{-\sigma(\kappa + k_i)}}. \tag{8}$$

4.2 Retailer's Decisions

The seller prices against gross demand and returns. Let D_i^l denote the purchase indicator implied by the cutoffs above, and let q_i^l denote the corresponding return probability under scheme $l \in \{r, c\}$ conditional on purchase: q_i^r is given by (5), and q_i^c equals q_i^r for insured purchasers and q_i^n in (8) for uninsured purchasers. Aggregating individual purchase choices and return probabilities over F_ν and Ψ yields gross demand $D^l \equiv N\,\mathbb{E}[D_i^l]$ and return volume $R^l \equiv N\,\mathbb{E}[D_i^l q_i^l]$. Net demand, $D^l - R^l$, is the number of units ultimately kept. Let *RI* denote an indicator equal to one under RI and zero under CI. The seller's aggregate profit is

$$\pi^l(p) = (p - \theta^c - I^r RI)(D^l - R^l) - (h + \theta^c + I^r RI)R^l. \tag{9}$$

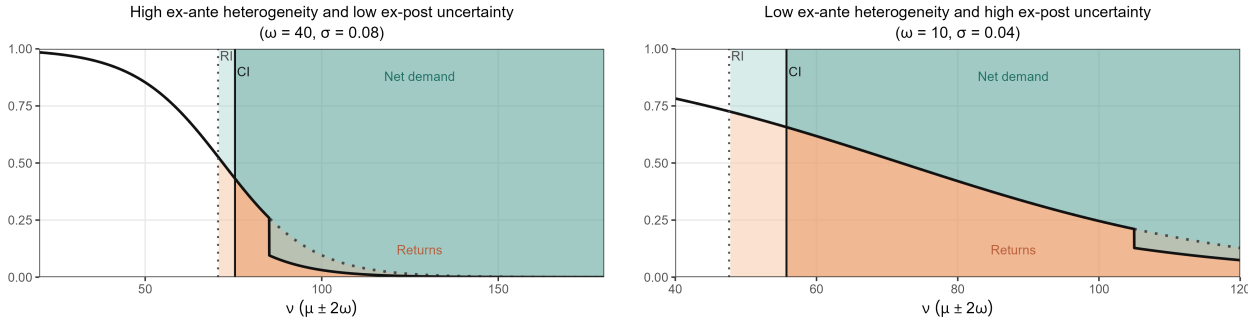
The first term is the margin earned on kept units, while the second is the loss on returned units. At the same price, RI generates higher gross demand than CI because consumers face no explicit premium, but it also attracts more lower-valuation buyers and reimburses return shipping for all buyers. Both raise return volume. The equilibrium price therefore maximizes (9) against net demand $D^l - R^l$ and return volume.

[13] This is equivalent to $e^{-\sigma\kappa}(1 - e^{-\sigma k_i}) \geq e^{\sigma I^c(k_i)} - 1$.

Three channels shape pricing: cost pass-through, surplus extraction, and return mitigation. These forces need not move prices in the same direction.

Cost Pass-through. Under RI, the seller pays the insurance premium I^r, which enters as a marginal cost, so a higher premium raises prices through pass-through. Under CI, the premium does not enter profit directly and affects profit only through demand and returns. Incidence among purchasers also differs. Under CI, consumers pay premiums directly, so insurance costs fall on lower-valuation buyers who return more often. Under RI, premiums are passed through to the product price and borne by all purchasers, so high-valuation keepers effectively subsidize insurance for lower-valuation frequent returners.

Surplus Extraction. Stronger net demand (more keepers) makes a price increase more profitable. We call this force *surplus extraction*. At a common price, RI expands gross demand (no premium) but also raises the return rate (automatic coverage), so net demand need not rise. Figure 3 shows the comparison. The CI purchase cutoff lies to the right of the RI cutoff because gross demand is lower at the same price. Under CI, $q(v)$ also drops sharply at the insurance cutoff, where private return costs jump.



(a) Large ω, large σ ($\omega = 40$, $\sigma = 0.08$).

(b) Small ω, small σ ($\omega = 10$, $\sigma = 0.04$).

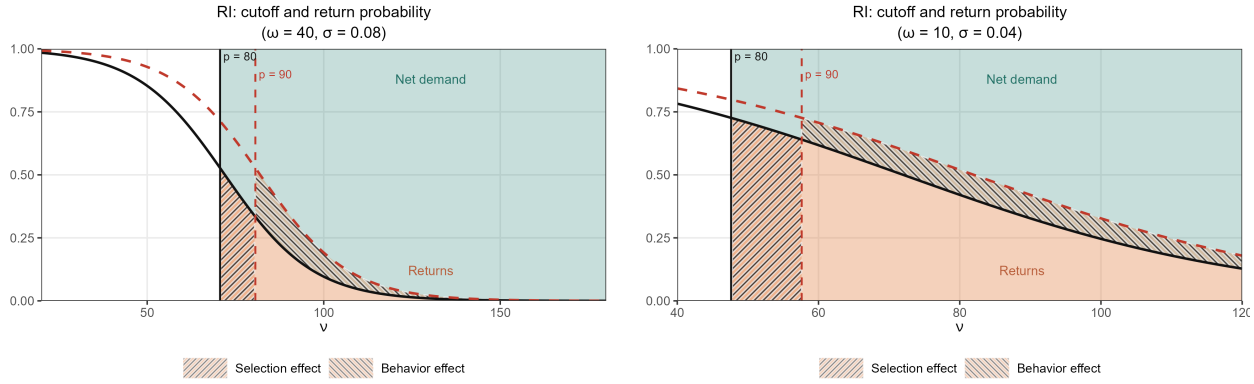
Figure 3: Returns and net demand under RI and CI at a common price. Dotted lines are RI and solid lines are CI. Purchasers lie to the right of each cutoff. Downward-sloping curves are return probabilities $q(v)$. Orange is the return region and teal is net demand. The horizontal axis covers $\mu \pm 2\omega$ with $\mu = 100$.

Whether net demand is higher under RI or CI depends on ex-ante heterogeneity ω and uncertainty. When ex-ante heterogeneity is large and uncertainty is small, as in Figure 3 (left), gross demand is relatively inelastic and return probability is sensitive to reimbursement. Return probability is strongly shifted by reimbursement because mean keep-or-return payoffs dominate idiosyncratic shocks. CI then generates higher net demand because uninsured consumers have substantially lower return rates and the reduction in gross demand from requiring consumers to pay the premium is small. When instead ex-ante heterogeneity is small and uncertainty is large, as in Figure 3 (right), return probabilities respond less to coverage and the premium shifts the purchase cutoff more compared to the concentrated value distribution, so RI's broader

participation can dominate and lead to a higher net demand than CI.

An important pricing implication is that when return probabilities are highly sensitive to coverage, a smaller but better-composed purchaser pool under CI can raise net demand and give the seller an incentive to charge a higher price. The surplus extraction force should not be equated mechanically with a higher equilibrium price, because the observed price also reflects pass-through and return mitigation.

Return Mitigation. A higher price affects return volume through two opposing forces. It screens out lower-valuation buyers, who return more often, which tends to lower return volume (a selection effect). Among remaining purchasers, however, a higher price raises individual return probabilities, because consumers receive a larger refund upon return (a behavior effect). Because these forces pull in opposite directions, the net effect is again ambiguous. Figure 4 illustrates both channels under RI for the same (ω, σ) cases as Figure 3 (CI is similar).



(a) Large ω, large σ ($\omega = 40$, $\sigma = 0.08$).

(b) Small ω, small σ ($\omega = 10$, $\sigma = 0.04$).

Figure 4: Selection and behavior effects of a price increase under RI. Solid curves are the baseline price. Dashed red curves are a higher price. Striped regions mark selection (between cutoffs) and the behavior effect (the upward shift in $q(v)$ among remaining purchasers). Parameters match Figure 3.

When ex-ante heterogeneity is large and uncertainty is small (Figure 4, left), $q(v)$ responds strongly to price. The seller cuts the price to reduce return volume, because a lower price reduces each buyer's return probability without much increase in sales. When instead ex-ante heterogeneity is small and uncertainty is large (right), purchase volume responds strongly to price and $q(v)$ responds less, so the seller prefers a higher price to curb returns. Raising price mainly discourages purchases. When returns are costly, the incentive to mitigate returns is stronger under RI. Whether that incentive raises or lowers the RI price depends on these primitives.

The three channels imply that RI need not have higher equilibrium prices than CI even though the seller pays the premium. When the premium is low and pass-through is weak, surplus extraction and return

mitigation can reverse the ranking. RI prices tend to be higher when ex-ante heterogeneity is small and uncertainty is large, because RI can expand net demand and return mitigation then works mainly through screening out frequent returners. The opposite ranking arises when ex-ante heterogeneity is large and uncertainty is small. The price comparison is therefore an equilibrium outcome rather than a direct measure of insurance incidence.

4.3 Welfare Framework

We evaluate return policies using a welfare framework that separates ex-post return inefficiency from ex-ante purchase inefficiency.

Ex-Post Efficiency. Conditional on purchase, production cost θ^c is sunk. A return is ex-post efficient if and only if the social value of keeping falls short of the social value of returning:

$$v_i + \varepsilon_i^{\text{keep}} \leq -(\kappa + k_i + h) + \varepsilon_i^{\text{return}}, \tag{10}$$

yielding the socially optimal return probability

$$q_i^*(v_i, k_i; \sigma, \kappa, h) \equiv \frac{e^{-\sigma(\kappa+k_i+h)}}{e^{\sigma v_i} + e^{-\sigma(\kappa+k_i+h)}}. \tag{11}$$

Because product prices and insurance payments are transfers, q_i^* evaluates returns against valuation v_i and internalizes net handling cost h. Under RSI, buyers instead compare $v_i - p$ against private return costs, ignoring the seller's return losses and generating excessive returns ($q_i^l > q_i^*$).

Ex-Ante Efficiency. Ex-ante efficiency concerns the purchase decision. Expected seller profit from consumer i is $\pi_i^l \equiv (1 - q_i^l)(p - \theta^c - I^r RI) - q_i^l(h + \theta^c + I^r RI)$, and insurer profit $\tilde{\pi}_i^l$ equals premiums collected minus expected reimbursements paid. Summing these with consumer surplus gives the social surplus generated by consumer i, $S_i^l \equiv CS_i^l + \pi_i^l + \tilde{\pi}_i^l$. A transaction is ex-ante efficient if and only if $S_i^l \geq 0$. We evaluate welfare per potential consumer and hence abstract from market size N. Non-purchasers receive zero. Expected consumer surplus and social welfare under scheme l are $CS^l \equiv \mathbb{E}[CS_i^l D_i^l]$ and $W^l \equiv \mathbb{E}[S_i^l D_i^l]$, where the expectation is over v_i and k_i. Under CI, insured and uninsured purchases enter the expectation as separate branches.

Welfare Decomposition and Extensive Margin. To separate purchase and return inefficiency, we compare welfare W^l to first best using an intermediate benchmark that holds observed purchase decisions D_i^l fixed but enforces optimal return behavior q_i^*. Because S_i^l evaluates surplus at private return behavior q_i^l, whether a purchase is efficient depends on return behavior, which can differ across policies. The intermediate

benchmark measures ex-ante inefficiency against a common return rule. Substituting q_i^* into individual social surplus yields the closed form

$$S_i^*(v_i, k_i) \equiv -\theta^c + \frac{1}{\sigma} \ln\left(e^{\sigma v_i} + e^{-\sigma(\kappa + k_i + h)}\right). \tag{12}$$

Let $\hat{v}^*(k_i)$ denote the first-best purchase cutoff solving $S_i^*(v_i, k_i) = 0$, and let $D_i^* \equiv \mathbf{1}\big(v_i \geq \hat{v}^*(k_i)\big)$. Defining $W^{OR} \equiv \mathbb{E}[S_i^*(v_i, k_i) D_i^l]$ as welfare with optimal returns at observed purchases and $W^{FB} \equiv \mathbb{E}[S_i^*(v_i, k_i) D_i^*]$ as first-best welfare, the welfare gap decomposes into

$$W^{FB} - W^l = \underbrace{\left(W^{OR} - W^l\right)}_{\Delta^{ex\text{-}post}} + \underbrace{\left(W^{FB} - W^{OR}\right)}_{\Delta^{ex\text{-}ante}}. \tag{13}$$

The first term, $\Delta^{ex\text{-}post}$, measures surplus lost to inefficient returns among purchasers under scheme l, holding D_i^l fixed. The second term, $\Delta^{ex\text{-}ante} = \mathbb{E}[S_i^*(v_i, k_i)(D_i^* - D_i^l)] = W^{FB} - W^{OR} \geq 0$, measures the remaining loss from distorted market participation relative to first best. While this participation gap can theoretically reflect either under- or overprovision, the option value of return places the first-best cutoff below production cost by that option value ($\hat{v}^* < \theta^c$), while the seller's market power pushes the purchase cutoff well above it ($\hat{v}^l > \hat{v}^*$). In our counterfactuals, $\Delta^{ex\text{-}ante}$ therefore always reflects underprovision relative to first best.

One limitation of $\Delta^{ex\text{-}ante}$ is that it evaluates entry holding returns at the hypothetical first-best rule q_i^*. It therefore does not show whether admitting marginal buyers under actual return behavior creates or destroys surplus. To evaluate the policy-relevant welfare effect of participation, we introduce another measure, the *extensive margin*. It evaluates individual surplus S_i^l at consumers' actual return behavior q_i^l. We quantify this margin using the purchase-efficient gap, the purchase cutoff $\hat{v}^l$ minus the efficient cutoff ($S_i^l(v_i) = 0$). If the gap is positive, the seller underprovides by excluding buyers with $S_i^l \geq 0$. If the gap is negative, the seller overprovides by admitting buyers with $S_i^l < 0$. Section 7 applies this framework at estimated parameters to evaluate the optimal-refund benchmark, RI versus CI, and counterfactual RPD menus.

5 Identification

The primitives to be identified are the coverage distribution $\Psi(\cdot)$, category-specific ex-ante heterogeneity ω_j, mean ex-ante value μ_t, ex-post precision σ_t, and marginal production cost θ_t^c. Net return handling cost h_t is calibrated in Section 6. As in Anderson et al. (2009), high uncertainty and low hassle cost both increase purchases and returns. We normalize hassle cost to 5 RMB so that uncertainty can be separated from hassle

cost.[14]

Under RI, fully reimbursed shipping makes the purchase threshold independent of k_i, so purchase does not select on coverage and $\Psi(\cdot)$ is non-parametrically identified from the empirical coverage distribution.

Let $\bar{w}_t^r$ be mean willingness to pay. Under RI, the purchase probability is $\mathbb{E}[D_{it}^r(k_i, p_t)] = 1 - \Phi\left(\frac{p_t - \bar{w}_t^r}{\omega_j}\right)$. Given a price instrument, variation in p_t across category-months traces out this purchase schedule. How fast the purchase probability falls in price identifies ex-ante heterogeneity ω_j, and where the schedule lies identifies mean willingness to pay $\bar{w}_t^r$. $\bar{w}_t^r$ decomposes into intrinsic value and return option value:

$$\bar{w}_t^r \equiv \underbrace{\mu_t + \xi_t}_{\text{Intrinsic value}} + \underbrace{\left[-\frac{1}{\sigma_t}\ln(1 - e^{-\sigma_t \kappa})\right]}_{\text{Return option value}}.$$

While purchase variation identifies $\bar{w}_t^r$, it cannot separate intrinsic value from return option value. Expected returns provide the additional identification that separates them. Holding purchase probability (and thus $\bar{w}_t^r$) fixed, expected returns strictly increase in uncertainty, so joint variation in purchases and returns across markets identifies σ_t, which backs out the option value and isolates intrinsic value $\mu_t + \xi_t$ (Online Appendix B). CI purchase, return, and insurance moments further restrict the same demand primitives through endogenous insurance selection under CI.

Finally, given calibrated net handling cost h_t, inverting the retailer's pricing first-order condition at observed prices maps estimated demand responses into marginal cost θ_t^c. Cross-market variation in cost shifters identifies cost parameters α^c and store-level cost shocks ζ_t^c via $\theta_t^c = X_t \alpha^c + \zeta_t^c$.

6 Estimation

We estimate the model by GMM, pooling moment conditions across categories j. The estimation sample contains 6 apparel categories and 313 category × month markets after smoothing, resampling, and the 5% purchase-probability screen described in Online Appendix A. Smoothing and resampling capture stable long-run demand. The screen keeps purchase probabilities away from zero to avoid the zero-share problem.[15] Market size is defined at the category × season level as twice the maximum monthly sales within that category-season, and monthly purchase probabilities equal smoothed sales divided by that size. Because the outside option is unobserved, absolute market size cannot be separately identified from the mean utility level. Setting market size to twice peak sales keeps purchase probabilities interior and comparable across

[14] Normalizing $\kappa = 0$ would imply that, under RI, the purchase probability is always one.

[15] A low purchase probability makes the purchase-probability inversion numerically unstable.

seasons. Category-specific ex-ante heterogeneity ω_j is fixed over time.

Each category-month market has a single list price and a single net handling cost, including months in which RI and CI both appear. Purchase, return, and insurance-purchase counts under each scheme are scaled by the time share of that scheme, so each scheme is treated as if it had operated for the full month at that common price. That price balances the seller's pricing incentives under both schemes in that category-month.

6.1 Return-handling calibration

Net return handling cost is h_t = gross handling − scrap value$_t$. We set gross handling to 15 RMB per return, consistent with industry reports indicating that each apparel return costs sellers approximately 15 RMB in sunk marketing, packaging, and reverse logistics.[16] For scrap value, we calibrate a three-way disposition of returned inventory using industry surveys. About one-half of returned apparel units are resold at full retail price (e.g., 48%). Condition grading assigns roughly 20–30% of units to discount or outlet channels and 20–30% to secondary liquidation; we combine these into a 40% markdown tail. Roughly 10% of returns are written off.[17] Weighting across these channels yields an average net scrap value of $0.5\,p_t + 0.4\times(0.5\text{–}0.7)\,p_t \approx 0.75\,p_t$, implying baseline net handling cost $h_t = 15 - 0.75\,p_t$ (RMB). Online Appendix E reports robustness checks under alternative scrap shares of $0.65\,p_t$ and $0.85\,p_t$.

6.2 Demand Side

We estimate demand parameters following the nested fixed-point approach of Berry et al. (1995). The nonlinear parameters $\Theta = (\beta^\omega, \beta^\sigma)$ determine ex-ante heterogeneity, $\omega_j = C_j'\beta^\omega$ for a category-indicator vector C_j, and ex-post precision, $\sigma_t = X_t^\sigma \beta^\sigma$ for a market-covariate vector X_t^σ. Mean utility is $\mu_t = X_t\beta^\mu$, where X_t contains category-month demand shifters: assortment size, in-season share, newly published share, category indicators, and garment style attributes (length, cut, and material texture). Because β^σ is identified from cross-market variation in σ_t, we keep X_t^σ parsimonious, with only category indicators and newly published share. Assortment size and in-season share shift ex-ante value and production cost (μ_t and θ_t^c)

[16] China Venture News and Zhou Jiabao (2025), reprinting *Times Finance*, cite Yinman founder Fang Jianhua on per-return seller costs for Chinese e-commerce apparel. Yinman operates reverse-logistics services for online retailers; the 15 RMB figure is described as the cost of a return that generates no resale revenue, net of any recovery. We treat this as a per-unit processing charge (reverse shipping, inspection, repackaging, and a share of customer-acquisition and fulfillment costs allocated to the return event) and hold it fixed across categories and months.

[17] 48% full-price resale: Eightx (2026, Average eCommerce Return Rate by Category), a U.S. e-commerce apparel benchmark. SELVANE (2024, The hidden cost of fast fashion returns) sorts returns into full-price, discount or outlet (about 20–30%), and secondary liquidation (about 20–30%), with 25–30% never returning at list price; we map the last two channels into the 40% markdown tail. McKinsey & Company (2020, Returning to order) estimates that about 10% of apparel returns are discarded; remaining units that fail full-price resale enter the markdown tail and recover 50–70% of list price. The same shares apply in every category, so h_t varies only through p_t.

rather than uncertainty, and are therefore excluded.

For each candidate Θ, we recover the mean-utility index $\delta_t(\Theta) = \mu_t + \xi_t$ by inverting model purchase probabilities. In single-policy markets, the inversion matches observed coverage-weighted purchase probabilities. When RI and CI coexist within a market, a single δ_t minimizes the sum of squared discrepancies in coverage-weighted purchase probabilities across both policies. We estimate $\beta^\mu(\Theta)$ by IV regression of $\delta_t(\Theta)$ on X_t and define structural demand shocks $\xi_t(\Theta) = \delta_t(\Theta) - X_t\beta^\mu(\Theta)$.

The instrument vector Z_t includes X_t and a Hausman-style cost instrument: the contemporaneous average price of other apparel categories sold by the retailer. Common input costs (textiles, labor, warehousing, and platform logistics) induce joint price variation across categories. Normalizing purchase probabilities by category-season market size removes the common seasonal demand shock across categories, and averaging prices within a category lessens the demand-shock links between individual garments.[18] With these adjustments, category-specific demand shocks ξ_{jt} remain largely orthogonal to demand shocks in other categories, so other-category prices pick up little own-category demand. The first-stage partial F-statistic of this instrument for own price is 100.15 ($p = 0.000$; $N = 313$), well above conventional thresholds for weak instruments.

Given candidate Θ, we form two sets of GMM moments: orthogonality conditions $E[Z_t\xi_t(\Theta)] = 0$ and squared prediction errors $m_t(\Theta)^2$ matching sample return and insurance rates to their model-implied counterparts, where $m_t(\Theta)$ is the coverage-weighted gap.[19] We choose $\hat{\Theta}$ to minimize $\bar{g}_n(\Theta)'W\bar{g}_n(\Theta)$, iterating the weighting matrix W to convergence, and set $\hat{\beta}^\mu = \beta^\mu(\hat{\Theta})$.

6.3 Supply Side

Given $\hat{\Theta}$, we compute model-implied demand and its price derivative under each policy in each market. Substituting these and the observed price into the retailer's pricing first-order condition gives one linear equation in θ_t^c, which we solve to recover the marginal production cost. When only one policy is observed in a market, inversion directly recovers θ_t^c. When both policies coexist, we invert each first-order condition and set θ_t^c to the weighted average of the two implied costs that minimizes the sum of squared pricing residuals.[20]

[18] A pair of pants may go especially well with a certain shirt, which links the demand shocks for those two garments. Averaging prices within each category lessens those links. Market size equals twice peak monthly sales in the category-season and is held fixed within that season, so a larger season raises the denominator and the common seasonal scale does not enter purchase probabilities or ξ_{jt}.

[19] Squaring each market's return- and insurance-rate errors prevents large positive and negative errors from canceling across markets. At the true parameter vector Θ_0 under correct specification, each $m_t(\Theta_0)$ equals zero.

[20] The weights are $(dD_t^r/dp_t)^2$ and $(dD_t^c/dp_t)^2$, because each first-order condition is linear in θ_t^c with slope dD_t^l/dp_t.

This yields 309 implied cost observations.[21] Regressing θ_t^c on X_t and time fixed effects without an intercept recovers cost parameters $\hat{\alpha}^c$ and store-level cost shocks $\hat{\zeta}_t^c$ from $\theta_t^c = X_t\alpha^c + \zeta_t^c$.

6.4 Parameter Estimates

Table 2 reports selected estimates associated with ex-ante heterogeneity β^ω, ex-post uncertainty β^σ, mean value β^μ, and production cost α^c. Standard errors and 95% percentile confidence intervals are computed from $B = 500$ transaction-level bootstrap replicates. Each draw resamples transactions with replacement and re-estimates the model, recentering GMM moment conditions at their full-sample values evaluated at $\hat{\Theta}$. Bootstrap distributions of the main demand and cost parameters appear in Online Appendix F.1.

Ex-ante heterogeneity ω_j is widest for dresses (59.125) and narrowest for tops (18.401). Ex-post uncertainty follows the same order: it is highest for dresses ($\sigma_t = 0.003$) and lowest for tops ($\sigma_t = 0.089$). A one-percentage-point increase in the newly published share raises mean utility μ_t by about 1.6 RMB. On the supply side, in-season listings have higher marginal production costs (13.704).

Table 2: Selected parameter estimates.

Parameter	Estimate	SE	95% CI
β^ω (dress)	59.125	27.485	[58.756 , 72.342]
β^ω (two-piece)	46.162	5.638	[33.979 , 55.141]
β^ω (outerwear)	45.050	13.794	[22.434 , 63.008]
β^ω (pants)	34.611	12.883	[28.793 , 46.593]
β^ω (sweater)	23.447	6.386	[16.540 , 29.890]
β^ω (tops)	18.401	7.964	[2.602 , 34.680]
β^σ (newly published share)	0.073	0.006	[0.065 , 0.081]
β^σ (dress)	0.003	0.000	[0.002 , 0.003]
β^σ (two-piece)	0.051	0.006	[0.044 , 0.071]
β^σ (outerwear)	0.040	0.009	[0.031 , 0.065]
β^σ (pants)	0.048	0.005	[0.039 , 0.053]
β^σ (sweater)	0.050	0.006	[0.042 , 0.060]
β^σ (tops)	0.089	0.037	[0.050 , 0.203]
β^μ (Assortment size)	-0.056	0.052	[-0.185 , 0.001]
β^μ (Newly published share)	155.304	21.784	[152.345 , 201.149]
β^μ (In-season share)	2.835	5.874	[-6.991 , 6.855]
α^c (Assortment size)	-0.115	0.020	[-0.151 , -0.073]
α^c (Newly published share)	-2.035	9.387	[-27.202 , 1.457]
α^c (In-season share)	13.704	1.644	[12.117 , 18.577]

[21] All 313 category-month markets enter demand estimation; supply inversion drops four cells where $|\partial D^l/\partial p| < 10^{-5}$, where dividing by a near-zero derivative is numerically unstable.

6.5 Model Fit

Table 3 compares model-implied and sample values of prices, purchase probabilities, and return and insurance rates at the category × month level. Panel A reports unconditional rates per potential consumer, counting both purchasers and non-purchasers in the denominator. These are the moments targeted in estimation,[22] while Panel B reports the corresponding rates per purchase. Each entry is an unweighted mean or standard deviation across the 313 category × month markets.

Table 3: Model vs. sample at $\hat{\Theta}$ (category × month cell means and SDs).

		Mean		SD	
		Model	Sample	Model	Sample
Price	RI (RMB)	106.76	99.83	53.55	39.56
	CI (RMB)	98.18	95.22	49.32	33.12
Purchase	RI (%)	18.90	26.20	19.59	15.51
	CI (%)	15.11	23.04	17.68	12.83
Panel A: Per potential consumer (targeted moments)					
Return	RI return (%)	6.63	7.69	9.47	6.27
	CI return (%)	2.94	5.55	8.48	4.65
	CI insurance (%)	1.55	0.55	8.83	0.79
Panel B: Per purchase (conditional)					
Return	RI return (%)	37.21	29.33	16.54	15.29
	CI return (%)	17.72	22.90	15.69	12.97
	CI insurance (%)	8.29	2.61	20.46	7.77

The model reproduces the key empirical patterns. Mean RI list prices, purchase probabilities, and return rates exceed their CI counterparts in both model and sample.[23] Purchase probabilities are underpredicted by about 7 percentage points under RI and about 8 under CI. This systematic underprediction can be due to the true valuation distribution having a heavier tail than the assumed normal.

At $\hat{\Theta}$, unconditional return rates average 6.6% under RI and 2.9% under CI in the model, compared with 7.7% and 5.6% in the sample. The implied RI-CI gap is 3.7 percentage points in the model and 2.1 percentage points in the sample, compared with 19.5 and 6.4 percentage points for the conditional gaps in Panel B. There, the model overpredicts RI returns and underpredicts CI returns. Return fit is therefore less poor than the conditional comparison alone suggests, although CI unconditional returns remain underpredicted by more than RI returns. Unconditional CI insurance rates are likewise closer than conditional take-up (1.6% vs. 0.6% in Panel A, against 8.3% vs. 2.6% in Panel B). Level fit remains imperfect. We still use the estimates to

[22] Estimation targets unconditional rates per potential consumer, which prevents markets with low purchase rates from destabilizing the moment conditions by dividing by near-zero denominators.

[23] In Table 1, pooled raw prices weight each transaction equally. RI transactions place more weight on lower-priced categories such as tops (31% of listings) and less on higher-priced categories such as outerwear than CI transactions do, so pooled CI prices exceed pooled RI prices even though RI prices are higher within every category except tops. Here, each category–month market contributes equally to the mean, downweighting the high-volume categories that drive the pooled ranking.

compare relative welfare across policies within each category, because welfare is determined mostly by the unconditional rates per potential consumer. The error in the unconditional purchase rate is large, but about the same under both policies, so it does not favor either one. The model overstates the unconditional return gap by about 1.6 percentage points, which slightly favors CI on the ex-post margin.

7 Counterfactual Simulations

We apply the welfare framework at the estimated parameters to evaluate return policies across categories. We first evaluate an optimal-refund benchmark to isolate the baseline market-structure distortion from monopoly pricing. We then compare the two observed policies, RI and CI, and evaluate an alternative return-penalty discount (RPD) mechanism. Section 7.4 examines the distribution of surplus across agents, and Section 7.5 conducts robustness checks.

For the welfare counterfactuals, we construct one *representative* snapshot per category. Covariates, premium-related inputs, demand shocks, and production-cost shocks are averaged across periods using consumer-count weights, and total market size is the sum of those counts. We then solve for equilibrium prices under each policy and evaluate purchase probabilities, return rates, and welfare by aggregating over valuations and coverage. Averaging primitives first, rather than averaging period-by-period welfare, targets long-run performance and keeps noisy thin months from dominating a highly nonlinear welfare functional.

7.1 Optimal-Refund Benchmark

In our model, marginal cost pricing ($p = \theta^c$) and a return refund equal to net salvage recovery $-h$ = scrap value − gross handling implement the first best. On return, the buyer receives net salvage recovery rather than recovering full product price p, fully internalizing the seller's net loss on return.[24] Return policy can improve ex-ante and ex-post efficiency, but once returns are efficient the remaining distortion is monopoly pricing above the production cost, which stems from market structure.

To isolate this residual market-structure distortion (a loss that return policy design itself cannot correct), we evaluate an *optimal-refund benchmark* (Online Appendix D.2). Setting the consumer refund equal to net salvage value eliminates ex-post return inefficiency, as return probabilities coincide with the first-best rule q_i^*. The remaining gap to the first best is entirely ex-ante inefficiency from monopoly pricing. Neither RI nor CI attains this benchmark, because both refund p on return.

[24] Expected seller profit from consumer i is then zero on every path, so $S_i^* = CS_i^*$ and efficient purchase and return decisions maximize consumer surplus as well as social surplus.

This benchmark is a welfare diagnostic rather than a direct policy prescription. Net handling cost h combines logistics, inspection, and resale value that vary across sellers and categories and are imperfectly observed, so mandating refund $-h$ listing by listing is impractical. Furthermore, distance-selling laws generally mandate full product-price reimbursement on no-reason returns, as under China's seven-day no-reason return rules and the EU Consumer Rights Directive.[25] Product-refund forfeiture beyond these statutory floors raises consumer-protection concerns, although restocking fees and partial refunds are tolerated when disclosed at checkout and tied to processing costs.

Table 4 reports the optimal-refund benchmark across the six apparel categories. Efficient returns set $\Delta^{ex\text{-}post} = 0$ by construction, so the remaining welfare shortfall is entirely ex-ante: about 29.4% of first-best welfare for dresses and 23.8%–25.6% in the other five categories, with welfare under this benchmark attaining 70.7% to 76.2% of first best.

Table 4: Optimal-refund benchmark: welfare decomposition and extensive-margin cutoffs

Category	Welfare (% FB)	$\Delta^{ex\text{-}ante}$ (% FB)	Purch. cutoff	Eff. cutoff	Marg. SS
dress	70.65	29.35	55.48	15.38	20.84
outerwear	74.59	25.41	113.55	78.36	32.09
pants	74.94	25.06	91.98	63.05	26.17
sweater	76.22	23.78	109.18	86.81	21.28
top	75.06	24.94	48.37	35.07	13.04
twopiece	74.36	25.64	102.31	69.36	30.47

Columns 4–6 report the purchase cutoff, the first-best efficient cutoff, and marginal social surplus at the purchase cutoff. Marginal social surplus is strictly positive in every category, ranging from 13.04 RMB for tops to 32.09 RMB for outerwear. This confirms that the ex-ante shortfall is monopoly underprovision on the extensive margin. The seller's market power pushes the purchase cutoff above the efficient cutoff, excluding marginal buyers whose purchase would generate positive social surplus. The purchase-efficient gap is largest for dresses (about 40.1) and smallest for tops (about 13.3).

7.2 RI vs. CI

7.2.1 Equilibrium primitives

Table 5 reports equilibrium primitives for RI and CI by category. The equilibrium price is higher under RI in every category. We explain it by revisiting the surplus-extraction and return-mitigation forces from Section 4.2, using sweaters as an illustration. The other categories are in Appendix C.

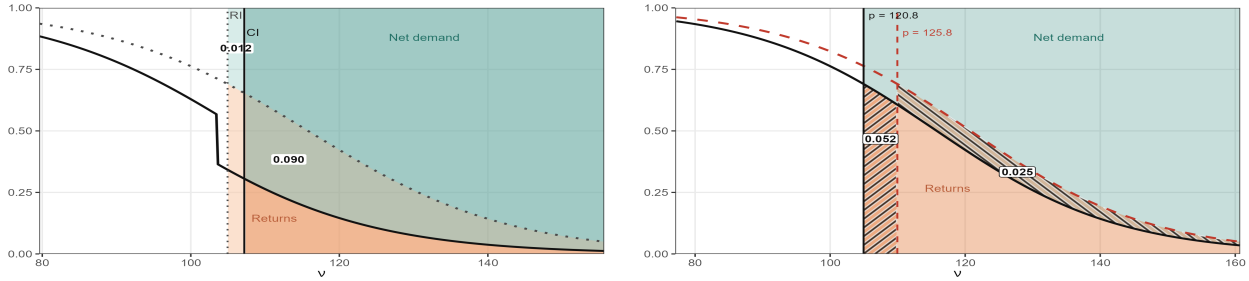
[25] China's *Law on Protection of Consumer Rights and Interests* (Art. 25) and *Interim Measures for Seven-Day No-Reason Returns on Online Purchases*. EU Directive 2011/83/EU, Art. 13. Exceptions and prior buyer confirmation narrow these rights.

Table 5: Category-level equilibrium primitives

Category	Production cost	Return handling	Ex-post SD	RI price	CI price	Purchase probability RI	Purchase probability CI	Return probability RI	Return probability CI
dress	42.22	-48.85	29.84	71.64	67.97	0.11	0.07	0.56	0.35
outerwear	81.76	-76.36	17.65	120.29	115.43	0.28	0.23	0.32	0.12
pants	66.26	-59.65	17.92	100.18	94.43	0.33	0.28	0.37	0.14
sweater	88.17	-71.95	17.32	120.79	112.13	0.41	0.37	0.42	0.15
top	35.39	-25.52	10.05	54.06	49.05	0.20	0.21	0.28	0.06
twopiece	72.08	-66.14	16.55	108.83	103.91	0.22	0.19	0.31	0.11

Notes: Return probabilities are conditional on purchase.

In Figure 5a, the return probability is lower under CI even among insured buyers, because CI's price is lower. Weaker reimbursement under CI excludes net demand of 0.012 despite a lower price, and its lower return probability adds net demand of 0.090, both as shares of market size. Combined, net demand is higher under CI, which raises the surplus-extraction incentive. The same pattern holds in the other categories except dresses, where net demand is slightly higher under RI. Figure 5b shows that selection is stronger than the behavior effect. When the price rises by 5 RMB, the mass of low-valuation frequent returners screened out is 0.052, while the higher return probability among remaining buyers raises return volume by 0.025. The seller mitigates returns by raising the price to screen out frequent returners. This effect is stronger under RI, and the same ranking holds in every category. In summary, the higher price under RI is attributed to stronger pass-through and return mitigation, which dominate the surplus-extraction effect.



(a) Surplus Extraction (Sweater) (b) Return Mitigation (Sweater)

Notes: Masses are normalized to a market of size 1. Coverage is set to the median of Ψ.

Simulated purchase probability is higher under RI than under CI in every category except tops, where it is about 1 percentage point higher under CI. This ranking for tops is inconsistent with the endogenous-moment fit table (Table 3) and the raw data.[26] We discuss its implication for the welfare comparison in the next subsection. Across all categories, the return probability conditional on purchase is lower under CI than under RI.

[26] In the sample, the purchase probability for tops is higher under RI in 7 of the 9 months observed under both policies, and it averages about 23% in the 16 RI-only months against about 15% in the 11 CI-only months.

Category-level variation begins with margins, which are related to the social surplus generated when a product is kept. The dress category exhibits the highest margin, $(p - \theta^c)/p$, at approximately 38%–41% across pricing policies. Sweaters are the lowest, at about 21%–27%. The other categories (tops, outerwear, pants, and two-piece sets) are about 28%–35%. Economically, a higher margin reflects a wider wedge between ex-ante consumer valuations and production costs. That wedge is the social surplus of a kept purchase, so the gains from expanding participation are larger where the margin is high.

When a consumer returns the product, the welfare outcome depends on the net return-handling cost (h_t), which captures salvage value net of shipping and processing frictions. Although a high return probability typically generates ex-post inefficiency, this effect is moderated by the efficiency of the return channel. In the dress category, net handling cost is $h_t = -48.85$ RMB (a net salvage recovery of 48.85 RMB), preserving 68%–72% of the retail price, so the social cost of a return is relatively low. For tops, net handling cost is −25.52 RMB (recovering 25.52 RMB), which covers 47%–52% of the price, so the social cost of a return is much higher. For sweaters, outerwear, pants, and two-piece sets, net salvage recovery covers about 60%–66% of the retail price.

Uncertainty governs how strongly reimbursement shifts probability between keeping and returning.[27] When uncertainty is low, mean keep-or-return payoffs dominate idiosyncratic shocks, and the return shipping fee shifts that probability more strongly. Dresses have high uncertainty (about 30 RMB). Outerwear, pants, sweaters, and two-piece sets have medium uncertainty (about 17 RMB). Tops have low uncertainty (about 10 RMB).

Conditional on a purchase, welfare depends on the margin, net salvage recovery, and the magnitude of uncertainty. The margin is related to the social surplus of a kept purchase, and net salvage recovery is what remains on a return. Uncertainty determines how strongly reimbursement shifts probability between keeping and returning. Price determines whether that purchase is realized. The welfare ranking of RI and CI depends jointly on the price and these three primitives.

7.2.2 Welfare decomposition

Figure 6 shows how these primitives aggregate into welfare. Each stacked bar expresses welfare and the shortfalls from ex-post and ex-ante inefficiency as percentages of first-best welfare. Ex-post inefficiency is larger under RI than under CI in every category, consistent with the lower return probabilities under optional

[27] Equations (5) and (8) imply that reimbursement multiplies return odds $q_i/(1 - q_i)$ by $e^{\sigma k_i}$; higher uncertainty lowers σ and therefore attenuates this shift.

insurance. Full-forfeiture RPD is included for later comparison.

On the ex-post margin, both RSI schemes refund p on return and are more lenient than the optimal-refund benchmark. On return, the seller incurs net cost $p + h$ (refunding p and absorbing net handling cost h), and the insurer bears return shipping k_i when reimbursed. Because the buyer does not face these costs when deciding to return, the uninternalized loss is $p + h + k_i$ under RI and for insured CI purchasers, but only $p + h$ for uninsured purchasers who bear k_i themselves. CI achieves better ex-post return discipline than RI.

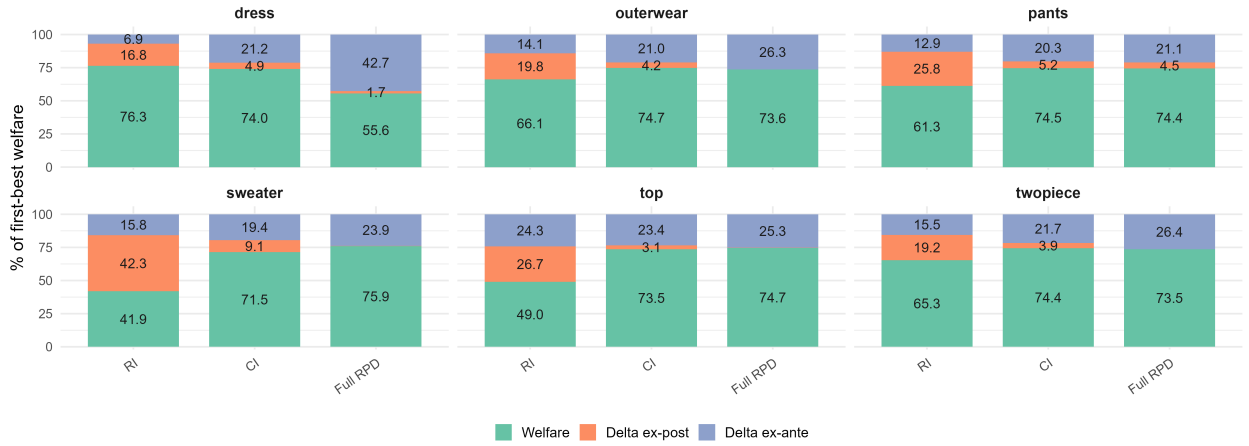


Figure 6: Welfare decomposition by category, normalized as a share of first-best welfare.

For dresses, RI attains a larger overall share of first-best welfare than CI (76% vs. 74%). A high margin and high net scrap value raise the surplus from a kept sale and limit the social cost of a return. High uncertainty adds option value that makes a marginal purchase socially valuable when the return rate is not too high. High uncertainty also yields fewer inefficient returns: the return decision depends mainly on the realized match value rather than on whether the shipping fee is reimbursed, so coverage barely changes the return rate. Automatic coverage and weaker surplus extraction then attract additional buyers who remain socially valuable on net.[28] The gain in ex-ante efficiency dominates the incremental ex-post inefficiency from returns.

Outside the dress category, the ranking reverses. For outerwear, pants, sweaters, and two-piece sets, margins are low to medium, salvage recovery is medium, and uncertainty is medium. Reimbursement therefore shifts more probability from keeping toward returning than it does when uncertainty is high, and more weight falls on the return, where salvage recovery is medium, so some purchases no longer generate positive social surplus. CI performs better by curbing that ex-post over-return and, where RI overprovides, by screening out inefficient purchases. The model overstates the RI–CI return gap relative to the sample

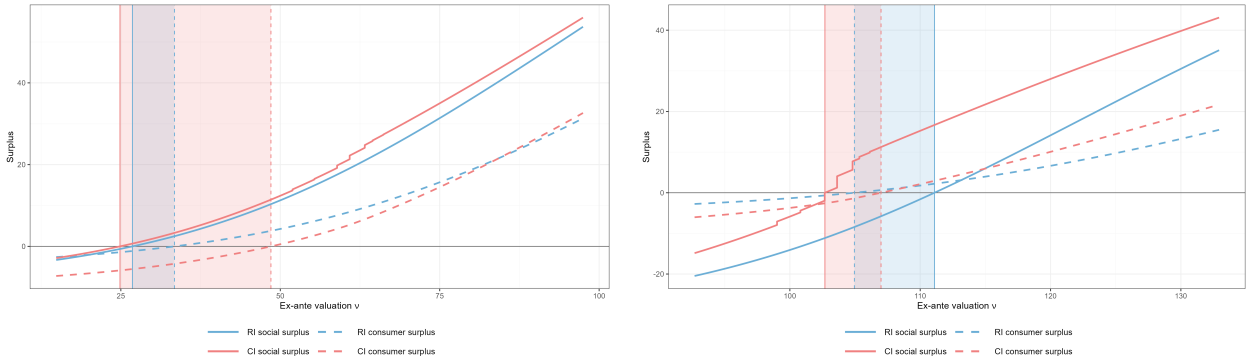
[28] Weaker surplus extraction need not show up as a lower RI price, because other pricing forces also operate.

(Table 3), so CI's estimated advantage outside dresses is, if anything, overstated, while RI's advantage in dresses is conservative.

For sweaters, CI outperforms RI by about 33 percentage points of first-best welfare ex post and underperforms by about 4 percentage points ex ante. The ex-post gap is large because uncertainty is medium, so reimbursement shifts more probability from keeping toward returning than when uncertainty is high. The margin is low and salvage recovery is medium. For tops, the margin is intermediate, while salvage recovery and uncertainty are both low. Relative to sweaters, the higher margin and the lower salvage recovery offset each other. RI underperforms CI on both margins, and about 51% of first-best welfare is dissipated under RI through poor ex-ante screening and ex-post over-return. For tops, purchase probability is higher under CI than under RI in the counterfactual, but lower in the sample, so CI's ex-ante advantage from screening out inefficient purchases is likely understated.

7.2.3 Extensive versus ex-post margin trade-off

The trade-off is clearer on the extensive margin. Figure 7 illustrates the comparison for dresses and sweaters at estimated equilibria. Solid curves are individual social surplus and dashed curves are consumer surplus, each evaluated at the equilibrium price. Solid vertical lines mark efficient cutoffs where social surplus is zero, and dashed vertical lines mark purchase cutoffs where consumer surplus is zero. The shaded band between them is the purchase-efficient gap.[29] Curves and cutoffs average across empirical coverage masses.



(a) Dress. (b) Sweater.

Figure 7: Extensive-margin cutoffs under RI and CI. Under CI, social surplus jumps at insurance cutoffs because coverage lowers private return cost. With multiple coverage levels, the averaged curve can appear kinked.

Relative to CI, whether RI performs better depends on whether extensive-margin gains from broader

[29]The purchase-efficient gap measures extensive-margin distortion. In a textbook single-product monopoly without returns, the degree of underprovision increases with the margin and the price. Across return schemes, however, posted price and margin poorly measure extensive-margin distortions. RI and CI attach different return benefits and insurance costs for buyers and shift costs differently for sellers, so they are effectively different products rather than the same good. The gap is therefore a better metric for comparing extensive-margin underprovision across policies.

participation outweigh ex-post losses from over-return. For dresses, both policies underprovide, but CI's gap is much larger (about 24 versus 7 under RI). In Figure 7, the vertical difference between the solid curves of CI and RI captures the combined effect of equilibrium pricing and the reduction in ex-post return inefficiency under CI when buyers remain uninsured. For dresses, CI's solid curve lies only modestly above RI's, so the ex-post margin gain from having uninsured buyers bear return shipping is small. This is because reimbursement shifts less probability between keeping and returning when uncertainty is high (a standard deviation of 29.84 RMB in Table 5). Against that small ex-post gain, RI's smaller extensive-margin gap dominates. The wide band of socially valuable buyers between the two purchase cutoffs raises welfare relative to CI. The purchase cutoffs differ substantially because high return rates make shipping coverage especially valuable for dresses. RI provides this coverage automatically, whereas CI requires consumers to pay a separate premium.

For sweaters, the extensive-margin distortion differs between RI and CI. CI still underprovides, whereas RI overprovides. The margin trade-off reverses relative to dresses. Across consumers' valuations, CI's solid curve lies well above RI's. This is because reimbursement shifts more probability between keeping and returning when uncertainty is medium, so having uninsured buyers bear return shipping greatly improves ex-post efficiency. On the other hand, RI's broader participation adds little on the extensive margin. The RI-CI difference in purchase cutoffs is narrow, and some of the buyers RI admits lie in the overprovision band. Return discipline under CI therefore dominates. Under RI, overprovision can arise because the seller does not fully internalize insurer claim losses, admitting marginal, return-prone consumers whose transactions generate negative social surplus.

Table 6 reports the purchase cutoff, the efficient cutoff, and marginal social surplus under RI, CI, and RPD. CI underprovides in every category. RI overprovides for pants, sweaters, and tops, and slightly underprovides for outerwear and two-piece sets. Relative to the gaps under the optimal-refund benchmark in Table 4 (about 13 to 40), RSI softens underprovision on the extensive margin, and RI softens more than CI and can overshoot. The RI purchase-efficient gap is about 7 for dresses and ranges from about -6 to about 0.6 for the other categories, while the CI purchase-efficient gap always stays positive.

The CI gap exceeds RI's because the purchase cutoff is higher and the efficient cutoff is lower. The purchase cutoff is generally higher under CI because CI does not reimburse return shipping and the seller's surplus extraction incentive is stronger. The efficient cutoff is lower because marginal entrants who remain uninsured face stronger ex-post discipline, which raises their social surplus relative to RI. For tops the two

Table 6: Purchase cutoffs, efficient cutoffs, and marginal social surplus by category

	Purchase cutoff			Efficient cutoff			Marginal social surplus		
Category	RI	CI	RPD	RI	CI	RPD	RI	CI	RPD
dress	33.41	48.53	66.19	26.49	24.54	34.59	2.49	11.44	26.07
outerwear	103.94	110.03	114.25	103.36	99.84	81.68	0.62	21.02	32.49
pants	83.41	88.83	89.38	85.04	81.04	74.62	-1.83	16.72	17.09
sweater	104.95	106.96	109.22	111.08	102.98	89.87	-8.42	11.23	19.70
top	48.17	47.87	48.48	49.31	40.04	35.34	-2.00	10.39	13.09
twopiece	94.19	99.29	102.89	93.67	90.09	71.99	0.59	20.49	30.82

Notes: Entries average across coverage types. Efficient cutoffs use each scheme's equilibrium return behavior.

purchase cutoffs are nearly the same (about 48), and the larger CI gap comes from a lower efficient cutoff.

Figure 6 and Figure 7 show the trade-off between ex-ante and ex-post efficiency. When the margin, salvage recovery, and uncertainty are high, such as in dresses, both RI and CI underprovide, and RI performs better by expanding efficient purchases. For outerwear, pants, sweaters, and two-piece sets, where margins are low to medium, salvage recovery is medium, and uncertainty is medium, CI performs better by curbing inefficient returns, with slightly lower ex-ante efficiency. For tops, where salvage recovery and uncertainty are both low despite an intermediate margin, CI performs better on both margins by screening out inefficient purchases and curbing excessive returns.[30]

7.3 Return-Penalty Discount (RPD)

Under return-penalty discount (RPD), each consumer chooses either a full price with full refund rights, or a discounted price with no or partial product refund on return, as in common "final sale" or restocking-fee menus.[31] RPD preserves a full-refund option while adding a second branch with stricter refund terms. On the discount branch, share $\phi \in (0, 1]$ of the discounted price is forfeited on return. Equilibrium conditions and the explicit expressions for the welfare objects appear in Online Appendix D.1.

Relative to RSI, RPD is an alternative for correcting excessive returns, especially when margins are thin and handling cost is high (salvage recovery is low). RSI expands low-end participation by subsidizing returns, thereby creating adverse selection ex-ante and over-return ex-post. In contrast, RPD makes the discount attractive to high-valuation keepers rather than to buyers who expect to return, so it does not expand low-end participation through return subsidies. On that branch, forfeiture also lowers the return probability ex-post.

[30] These rankings use each scheme's private return behavior q_i^l. The efficient cutoff under the optimal-refund benchmark, evaluated at q_i^*, remains below every q_i^l-efficient cutoff, so the first-best residual $\Delta^{ex\text{-}ante}$ is still underprovision even where RI overprovides on the extensive margin for some categories (Table 4).

[31] Amazon allows a restocking fee of up to 20 percent after the return window.

RPD also places screening with the party that holds product-level information. Under CI, the insurer sets standardized premiums while the seller chooses only product price. Even when the return rate is low, a premium set from return probability can be too low if handling costs are high, so optional coverage underprices the social cost of returns. Cheap coverage can then generate additional purchases and returns that impose a large negative externality on the seller, which cannot be corrected through product price alone. In contrast, under RPD the platform sets the refund terms, and the seller chooses both the list price and the discount. These two instruments let the seller incorporate private information about return costs into the menu, which can improve welfare relative to CI. Nevertheless, the additional instruments can also strengthen the seller's ability to extract surplus and enlarge the purchase-efficient gap.

7.3.1 Full-forfeiture RPD

We first report full-forfeiture RPD ($\phi = 1$), the leading case that delivers the strictest return discipline and is straightforward to implement. Ex ante, stricter refund terms can exclude socially valuable marginal buyers. Ex post, no product refund narrows the return subsidy to an externality gap of $-h$ alone. When net scrap value exceeds handling ($h < 0$), this gap is positive, so full forfeiture can overshoot into under-return because consumers do not internalize the positive externality that net recovery confers on the seller. Table 7 reports the list price, discount d, effective price, and segment-level purchases and returns.

Depending on the wedge between the purchase cutoff and the cutoff that is indifferent between the discount and the full price, buyers may end up only on the discount branch. In most categories full-price demand is negligible, purchases concentrate on the discount branch, and the return probability there is about zero. Pants are the exception: the full-price purchase probability is 0.15, against 0.12 on the discount branch, and the full-price return probability is 0.21. In contrast to RSI, ex-post inefficiency on the discount branch takes the form of under-return rather than over-return. Ex-post inefficiency is nevertheless smaller under full-forfeiture RPD than under either RSI policy (Figure 6).

Table 7: RPD counterfactual: prices and segment-level purchase probabilities, returns, and return probabilities by category

				Discount option			Full price		
Category	Price	Discount	Effective price	Purchase prob.	Return	Return prob.	Purchase prob.	Return	Return prob.
dress	104.85	37.98	66.88	0.04	0.00	0.01	0.00	0.00	
outerwear	196.74	82.49	114.25	0.20	0.00	0.00	0.00	0.00	
pants	94.98	1.56	94.29	0.12	0.00	0.00	0.15	0.03	0.21
sweater	115.72	6.49	109.22	0.34	0.00	0.00	0.00	0.00	0.00
top	83.91	35.43	48.48	0.20	0.00	0.00	0.00	0.00	
twopiece	174.98	72.09	102.89	0.17	0.00	0.00	0.00	0.00	

Notes: Effective price is $p - d$. Full-price return probabilities are left blank for segments with negligible full-price volume.

Participation, on the other hand, is weaker. By driving return probabilities on the discount branch close

to zero, refund forfeiture leaves the seller with a pool of confident keepers who rarely return, sharpening the surplus-extraction incentive. For dresses, the RPD purchase cutoff is about 66, against about 49 under CI and 33 under RI, even though RPD's effective price (about 67) lies slightly below the CI price (about 68). Forfeiting the product payment on return removes the implicit full-refund subsidy, so the purchase cutoff can rise above the CI cutoff even when the demand-weighted price paid is lower.

Relative to the RI and CI purchase-efficient gaps above, full-forfeiture RPD underprovides by more. For dresses the gap rises to about 32 and marginal social surplus to about 26 (Table 6). Outside dresses, RPD's gap exceeds CI's in every category.

Relative to RI, this stronger screening raises welfare in every category other than dresses. Relative to CI, it raises welfare for sweaters, where the margin is thin. For tops, outerwear, pants, and two-piece sets, welfare is almost indistinguishable from CI, with tops about one percentage point higher, outerwear and two-piece sets about one percentage point lower, and pants tied. In the dress category, where the extensive margin is especially valuable, RPD underperforms both RI and CI. Marginal entrants often have relatively low valuations and high return propensity, yet many such purchases remain socially worth encouraging because kept sales generate high surplus. RSI preserves a channel for those buyers. The discount branch under full-forfeiture RPD largely shuts it down, so the participation loss outweighs the ex-post gain and welfare falls to about 56% of first best.

7.3.2 Partial forfeiture and restocking fees

Full forfeiture is extreme. Softening the return penalty on the discount branch to a partial forfeiture $\phi \in (0, 1)$ restores some participation while retaining most of the return discipline. A natural retail implementation is a restocking fee. On return the buyer pays $\phi(p - d)$ and recovers share $1 - \phi$ of the discounted payment. We focus on restocking fees in the 10–15% range, which are common in retail practice on eligible returns. Figure 8 shows the welfare decomposition across categories as ϕ varies over $\{0.05, 0.10, 0.15, 0.20\}$.

A lower ϕ reduces under-return and raises participation relative to full forfeiture while retaining most of the return discipline. Outside dresses, welfare at a 10%–15% restocking fee is about 75–78% of first best. Partial-forfeiture RPD exceeds both RSI policies and full-forfeiture RPD at these fees. It matches or exceeds the optimal-refund benchmark, except for tops, where partial-forfeiture RPD at a 10%–15% restocking fee underperforms the optimal-refund benchmark by about 1 percentage point (about 75% versus about 76%). The advantage over RI is large, about 9–36 percentage points of first best. The advantage over CI is small,

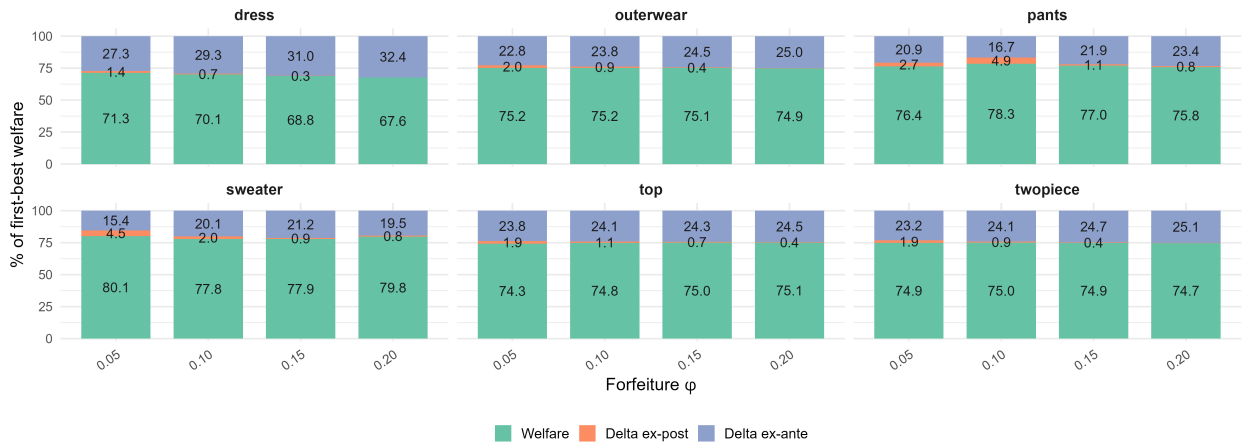


Figure 8: Partial-forfeiture RPD: welfare decomposition by category and $\phi \in \{0.05, 0.10, 0.15, 0.20\}$, normalized within each cell.

about 0.5–4 percentage points, except for sweaters, where the margin is lowest and the advantage is about 6 percentage points. For dresses, welfare at a 10%–15% restocking fee is about 69%–70% of first best. It remains below RI (about 76%), CI (about 74%), and the optimal-refund benchmark (about 71%), and lies well above full-forfeiture RPD (about 56%).

7.4 Welfare across agents

The decompositions above split efficiency losses relative to first best. Figure 9 instead shows how expected social surplus divides among consumers, the seller, and the insurer, as shares of first-best welfare. Partial-forfeiture RPD uses a 10%–15% restocking fee. Expected consumer surplus and seller profit lie above zero. Insurer profit is zero under RPD, which has no insurance, and under CI outside dresses, where no one buys coverage. It is negative under RI or CI when buyers are insured, because expected claims exceed the observed premiums.[32] The three components sum to welfare.

Welfare, expected consumer surplus, and seller profit largely follow the same ranking in every category. The exception is pants, where welfare and expected consumer surplus rank partial-forfeiture RPD, then CI, then full-forfeiture RPD, while seller profit under full-forfeiture RPD and CI is slightly above partial-forfeiture RPD (42.2% and 42.1% versus 41.9%). Outside dresses, partial-forfeiture RPD ranks first. Full-forfeiture RPD and CI are about the same (except for sweaters, where full-forfeiture RPD exceeds CI), and both rank above RI. For dresses, the ranking is RI, then CI, then partial-forfeiture RPD, then full-forfeiture RPD.

[32] Negative insurer profit for a single apparel seller is consistent with reported underwriting losses on Taobao return-shipping insurance (Huatai in 2012; Wang Liangdi, 2017, Peking University).

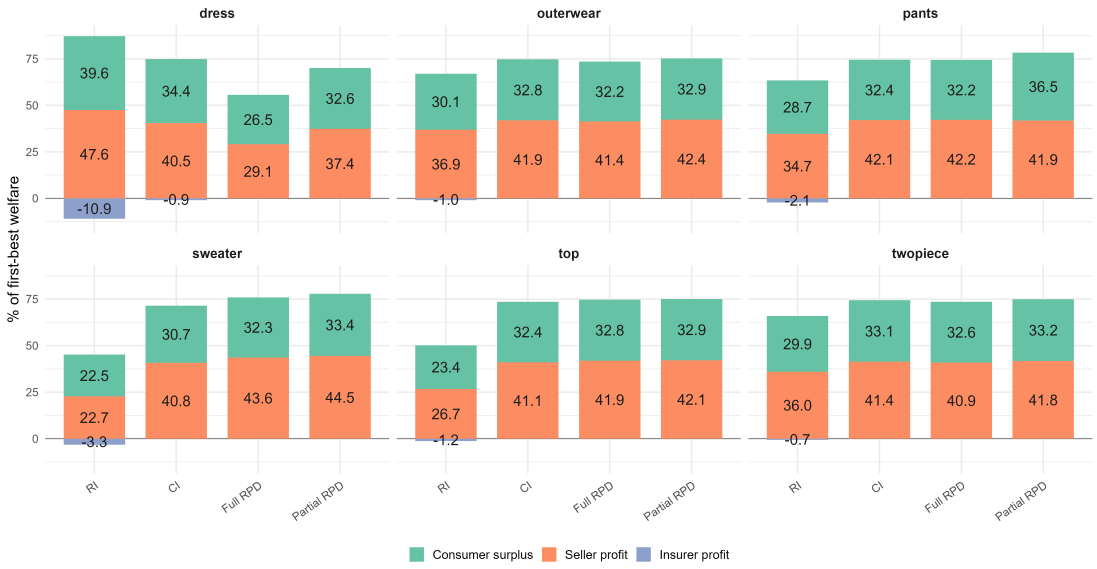


Figure 9: Welfare across agents by category and return policy.

7.5 Robustness check

7.5.1 Scrap-value sensitivity

The baseline decomposition uses the scrap calibration in Section 6.1 (about 75% of price on average). Because $h_t = 15 -$ scrap value, alternative recovery rates shift the return externality and can reorder counterfactuals even when demand primitives are fixed. Online Appendix E gives a table of welfare for each policy at scrap shares of 0.65 p_t and 0.85 p_t, and the welfare decomposition.

The RI–CI tradeoff and the advantage of partial forfeiture largely remain intact, with two exceptions. Outside dresses, CI remains above RI at both scrap rates. For dresses, RI still ranks above CI at 85% scrap, but lies about 6 percentage points below CI when the recovery rate is 65% (about 68% versus about 74%). A 10%–15% restocking fee still beats both RSI policies outside dresses at 65% scrap and, at 85% scrap, for sweaters and tops. For outerwear, pants, and two-piece sets, partial RPD beats RI by about 2, 6, and 3 percentage points at 85% scrap and falls short of CI by about 0.5–1.5 percentage points.

7.5.2 Inference on welfare ranking and decomposition

Online Appendix F shows the bootstrap distributions of welfare and its decomposition by policy within each category. The interpretations below are based on a visual reading of those distributions and on the pairwise win rates, which count the share of draws in which one policy ranks above another.

The results on how the welfare ranking differs between dresses and the other categories, and on the ex-ante and ex-post efficiency tradeoff, are stable. RI ranks above the other policies for dresses. Outside

dresses, partial-forfeiture RPD, evaluated at the welfare-maximizing restocking fee over $\phi \in \{0.10, 0.15\}$, ranks above the other policies for two-piece sets, outerwear, pants, and sweaters, as in Figure 9. For tops, it ranks above RI and CI, while full RPD and partial RPD are close. Relative to CI, RI has lower ex-ante inefficiency and higher ex-post inefficiency in every category except tops, where it is worse on both margins. Relative to full forfeiture, partial RPD lowers ex-ante inefficiency. Its ex-post inefficiency is lower for dresses, similar for pants, and higher for the other categories.

Overall, the bootstrap evidence corroborates the welfare rankings and decomposition patterns discussed above.

8 Conclusion

This paper studies the welfare consequences of RSI and RPD in online retail. Using transaction-level data from an apparel retailer on Taobao that repeatedly switched between RI and CI, we estimate a structural model of purchase, return, and insurance decisions under monopoly pricing across six garment categories. The welfare ranking among return policies is category-specific, governed by the trade-off between ex-ante efficiency and ex-post efficiency.

For categories in which the margin, net scrap value, and uncertainty are high, a generous return policy (RI) performs better by greatly improving ex-ante efficiency at the cost of ex-post efficiency, as for dresses. For categories in which the margin, net scrap value, and uncertainty are lower, a stricter return policy (CI) performs better by greatly improving ex-post efficiency at the cost of ex-ante efficiency, as for sweaters. When net scrap value and uncertainty are both low, as for tops, a generous return policy strongly shifts the return rate and makes many purchases inefficient to begin with. A stricter return policy then improves both ex-ante efficiency and ex-post efficiency by screening out inefficient purchases and curbing inefficient returns.

The platform should set the return policy according to these three category-specific primitives. When the platform does not observe these primitives, partial-forfeiture RPD with a moderate restocking fee (10%–15%) outperforms both RSI policies outside dresses at the baseline and when net scrap value is low. When net scrap value is high, partial RPD still outperforms both RSI policies for sweaters and tops. For outerwear, pants, and two-piece sets, partial RPD beats RI but falls slightly short of CI. For dresses, partial RPD remains below RI at the baseline and when net scrap value is high.

These findings challenge prevailing e-commerce platform practices. Platforms frequently compel mer-

chants to offer automatic return insurance indiscriminately, even on low-cost merchandise with a thin margin and low net scrap value, or require them to adopt RI in order to join major sales events such as Singles' Day. Mandatory RI is inefficient in thin-margin or low-net-scrap-value categories, dissipating about 58% of first-best welfare for sweaters and about 51% for tops. The social cost of uniform reimbursement is larger under emerging practices such as refund-without-return policies, because net scrap value can fall below our baseline calibration.

8.1 Limitations and future research

We treat RSI as given when recovering demand and inverting seller pricing. Conditional on the observed scheme, endogenizing enrollment would leave demand estimation largely unchanged. Purchase, insurance, and return still identify the same primitives. The supply side would change, because optimality of the chosen scheme becomes an additional inequality constraint on total profit at the store level. Recovered costs would have to rationalize both within-scheme pricing and the seller's preference for RI versus CI at the store level. We deliberately avoid that inequality.[33] Total profit at the store level depends on the relative market size of each category, whereas production cost recovered only from the observed price does not. The model treats each category on its own, and the welfare comparison is within a category, so relative market size should not enter the recovered costs.

Several extensions would change the policy ranking. First, under oligopoly, buyers screened out by a high price or a strict refund can divert to other products or sellers. That diversion mitigates the social cost of exclusion relative to our single-seller benchmark and therefore tilts welfare toward stricter return terms. Second, we study platform-regulated refund menus (RPD). Delegating refund choice to sellers adds a screening instrument that can expand participation for some types while strengthening surplus extraction, with an ambiguous net welfare effect. Third, our counterfactuals are risk-neutral; risk aversion would raise the value of automatic coverage. Fourth, we take insurer premiums as exogenous. Endogenizing them would reintroduce double marginalization between the insurer and the seller, with an ambiguous effect on the ranking across categories.[34] Fifth, on a two-sided platform, return policy also shifts cross-network externalities between buyers and sellers; characterizing the platform-optimal refund then requires modeling

[33] Both schemes are observed in 18 of 73 months. In those months the store changes scheme while category fundamentals are expected to be relatively unchanged, so a ranking of total profit cannot justify the observed scheme.

[34] Double marginalization pushes the retail price above the vertically integrated benchmark. The wedge is weaker than in a fixed-cost insurance model because expected claims scale with the return probability, so the insurer sets a lower premium than when claims are a fixed cost. Remaining pass-through can still widen the private-social gap in return costs. Ex ante, higher prices are costly when marginal purchases have positive social surplus (as for dresses) but can help in thin-margin categories where RSI already overprovides. Ex post, a higher price can worsen over-return.

both sides jointly, which we leave for future research.

References

Abdulla, H., M. Ketzenberg, and J. D. Abbey (2019): "Taking Stock of Consumer Returns: A Review and Classification of the Literature," *Journal of Operations Management*, 65, 560–605.

Altug, M. S. and T. Aydinliyim (2016): "Counteracting Strategic Purchase Deferrals: The Impact of Online Retailers' Return Policy Decisions," *Manufacturing & Service Operations Management*, 18, 376–392.

Altug, M. S., T. Aydinliyim, and A. Jain (2021): "Managing Opportunistic Consumer Returns in Retail Operations," *Management Science*, 67, 5660–5678.

Anderson, E. T., K. Hansen, and D. Simester (2009): "The Option Value of Returns: Theory and Empirical Evidence," *Marketing Science*, 28, 405–423.

Berry, S., J. Levinsohn, and A. Pakes (1995): "Automobile Prices in Market Equilibrium," *Econometrica*, 63, 841–890.

Bower, A. B. and J. G. Maxham (2012): "Return Shipping Policies of Online Retailers: Normative Assumptions and the Long-Term Consequences of Fee and Free Returns," *Journal of Marketing*, 76, 110–124.

Che, Y.-K. (1996): "Customer Return Policies for Experience Goods," *The Journal of Industrial Economics*, 44, 17–24.

Chen, J. (2021): "A Structural Model of Purchases, Returns, and Return-Based Targeting Strategies," *Journal of the Academy of Marketing Science*, 49, 695–718.

Chen, J. and R. Grewal (2013): "Competing in a Supply Chain via Full-Refund and No-Refund Customer Returns Policies," *International Journal of Production Economics*, 146, 246–258.

Chen, T., K. Ma, and C. Zheng (2018): "Does Freight Insurance Really Work? Effect of Freight Insurance on Consumers' Attitude in Online Shopping," *International Journal of Internet Marketing and Advertising*, 12, 209–232.

Chesnokova, T. (2007): "Return Policies, Market Outcomes, and Consumer Welfare," *Canadian Journal of Economics/Revue canadienne d'économique*, 40, 296–316.

Costello, J. P. and C. J. Bechler (2026): "Just Keep It: When and Why Returnless Product Returns Foster Brand Support," *Journal of Marketing Research*, 63, 191–209.

Courty, P. and L. Hao (2000): "Sequential Screening," *The Review of Economic Studies*, 67, 697–717.

Davis, S., E. Gerstner, and M. Hagerty (1995): "Money Back Guarantees in Retailing: Matching Products to Consumer Tastes," *Journal of Retailing*, 71, 7–22.

Davis, S., M. Hagerty, and E. Gerstner (1998): "Return Policies and the Optimal Level of "Hassle"," *Journal of Economics and Business*, 50, 445–460.

Fan, Z.-P. and Z. Chen (2020): "When should the e-tailer offer complimentary return-freight insurance?" *International Journal of Production Economics*, 230, 107890.

Fang, L., C. Wu, Z. Yuan, and K. Zhang (2026): "Returnless Partial Refunds: A Large-Scale Field Experiment in E-Commerce," *Marketing Science*.

Griffis, S. E., S. Rao, T. J. Goldsby, and T. T. Niranjan (2012): "The Customer Consequences of Returns in Online Retailing: An Empirical Analysis," *Journal of Operations Management*, 30, 282–294.

Hess, J. D., W. Chu, and E. Gerstner (1996): "Controlling Product Returns in Direct Marketing,"

Marketing Letters, 7, 307–317.

Hjort, K. and B. Lantz (2016): "The Impact of Returns Policies on Profitability: A Fashion E-Commerce Case," *Journal of Business Research*, 69, 4980–4985.

Hong, Y. and P. A. Pavlou (2014): "Product Fit Uncertainty in Online Markets: Nature, Effects, and Antecedents," *Information Systems Research*, 25, 328–344.

Hsiao, L. and Y.-J. Chen (2012): "Returns Policy and Quality Risk in E-Business," *Production and Operations Management*, 21, 489–503.

Huang, X. and D. Zhang (2020): "Service Product Design and Consumer Refund Policies," *Marketing Science*, 39, 366–381.

Ibragimov, M., S. El Kihal, J. Hauser, and R. Ursu (2026): "Return-Aware Search: Jointly Modeling Clicks, Purchases, and Returns in Online Retail," MIT Sloan Working Paper 6967-24.

Inderst, R. and M. Ottaviani (2013): "Sales Talk, Cancellation Terms and the Role of Consumer Protection," *The Review of Economic Studies*, 80, 1002–1026.

Inderst, R. and G. Tirosh (2015): "Refunds and Returns in a Vertically Differentiated Industry," *International Journal of Industrial Organization*, 38, 44–51.

Janakiraman, N., H. A. Syrdal, and R. Freling (2016): "The Effect of Return Policy Leniency on Consumer Purchase and Return Decisions: A Meta-analytic Review," *Journal of Retailing*, 92, 226–235.

Janssen, M. and C. Williams (2024): "Consumer Search and Product Returns in E-Commerce," *American Economic Journal: Microeconomics*, 16, 387–419.

Jerath, K. and Q. Ren (2024): "Consumer Search and Product Returns," *Marketing Science*.

Li, S., J. Yu, and M. Zhang (2024): "Search Prominence with Costly Product Returns," Working paper. https://arxiv.org/abs/2410.06791.

Li, X., J. Gao, and Y. Bian (2023a): "Return Freight Insurance Strategies for the Online Retailer and Insurance Company," *International Journal of Production Economics*, 256, 108752.

Li, Y., G. Li, and X. A. Pan (2023b): "Optimal Return Shipping Insurance Policy with Consumers' Anticipated Regret," *Production and Operations Management*, 32, 3209–3226.

Li, Y., G. Li, G. K. Tayi, and T. C. E. Cheng (2021): "Return Shipping Insurance: Free versus for-a-fee?" *International Journal of Production Economics*, 235, 108110.

Matthews, S. A. and N. Persico (2007): "Information Acquisition and Refunds for Returns," PIER Working Paper 07-021. Earlier draft: Information Acquisition and the Excess Refund Puzzle (PIER 05-015, 2005). https://ssrn.com/abstract=1004280.

McWilliams, B. (2012): "Money-Back Guarantees: Helping the Low-Quality Retailer," *Management Science*, 58, 1521–1524.

Montaguti, E., S. Neslin, and S. Valentini (2026): "Free Riding on Product Returns to Drive Profits," *Journal of Marketing*, 1–21, advance online publication; printed as Journal of Marketing 0(0). Copyright American Marketing Association 2026.

Moorthy, S. and K. Srinivasan (1995): "Signaling Quality with a Money-Back Guarantee: The Role of Transaction Costs," *Marketing Science*, 14, 442–466.

Najafi, S., I. Duenyas, and A. Singla (2024): "Price Markdowns to Induce Customers to Opt out of Free Returns," *IISE Transactions*, 0, 1–15.

Newberry, P. and X. Zhou (2024): "The Effects of Return Policies in E-commerce," Working paper, University of Georgia and Amazon. https://www.colorado.edu/economics/sites/default/files/

attached-files/newberry.pdf.

Padmanabhan, V. and I. P. L. Png (1997): "Manufacturer's Returns Policies and Retail Competition," *Marketing Science*, 16, 81–94.

Patel, P., C. Baldauf, S. Karlsson, and P. Oghazi (2021): "The Impact of Free Returns on Online Purchase Behavior: Evidence from an Intervention at an Online Retailer," *Journal of Operations Management*, 67, 511–555.

Petersen, J. A. and V. Kumar (2009): "Are Product Returns a Necessary Evil? Antecedents and Consequences," *Journal of Marketing*, 73, 35–51.

——— (2015): "Perceived Risk, Product Returns, and Optimal Resource Allocation: Evidence from a Field Experiment," *Journal of Marketing Research*, 52, 268–285.

Petrikaitė, V. (2018): "A Search Model of Costly Product Returns," *International Journal of Industrial Organization*, 58, 236–251.

Ren, M., J. Liu, S. Feng, and A. Yang (2021): "Pricing and Return Strategy of Online Retailers Based on Return Insurance," *Journal of Retailing and Consumer Services*, 59, 102350.

Shang, G., B. P. Ghosh, and M. R. Galbreth (2017a): "Optimal Retail Return Policies with Wardrobing," *Production and Operations Management*, 26, 1315–1332.

Shang, G., P. Pekgün, M. Ferguson, and M. Galbreth (2017b): "How Much Do Online Consumers Really Value Free Product Returns? Evidence from eBay," *Journal of Operations Management*, 53-56, 45–62.

Shehu, E., D. Papies, and S. A. Neslin (2020): "Free Shipping Promotions and Product Returns," *Journal of Marketing Research*, 57, 640–658.

Shulman, J. D., A. T. Coughlan, and R. C. Savaskan (2009): "Optimal Restocking Fees and Information Provision in an Integrated Demand-Supply Model of Product Returns," *Manufacturing & Service Operations Management*, 11, 577–594.

——— (2010): "Optimal Reverse Channel Structure for Consumer Product Returns," *Marketing Science*, 29, 1071–1085.

——— (2011): "Managing Consumer Returns in a Competitive Environment," *Management Science*, 57, 347–362.

Su, X. (2009): "Consumer Returns Policies and Supply Chain Performance," *Manufacturing & Service Operations Management*, 11, 595–612.

Tran, T., H. Gurnani, and R. Desiraju (2018): "Optimal Design of Return Policies," *Marketing Science*, 37, 649–667.

Vollmer, D. (2025): "Is Resale Needed in Markets with Refunds? Evidence from College Football Ticket Sales," *The RAND Journal of Economics*, 56, 251–268.

Wood, S. L. (2001): "Remote Purchase Environments: The Influence of Return Policy Leniency on Two-Stage Decision Processes," *Journal of Marketing Research*, 38, 157–169.

Zhang, C., M. Yu, and J. Chen (2022): "Signaling Quality with Return Insurance: Theory and Empirical Evidence," *Management Science*, 68, 5847–5867.

Online Appendix

A Data Construction

This appendix constructs the analysis sample from raw transactions in three steps: filters, RSI classification and CI imputation, and aggregation.

A.1 Sample Filtering

We keep 2017–2023 transactions and drop multi-item or multi-quantity orders (complementarity and lower marginal return cost), orders without a logistic ID (typically cancelled before shipment), orders with buyer-paid premium shipping services (a separate surcharge outside the product price), partial refunds for minor defects, and non-product listings. We also drop single-day RSI switches for promotions such as Singles' Day, so the sample reflects the store's usual scheme.

A.2 RSI Scheme Classification and CI Imputation

Missing insurance fields are ambiguous (RI ineligibility versus CI non-purchase). Complete records with premium and coverage are RI. Coverage only is CI. We propagate the day-level majority label, treat long runs of missing records as CI, and drop remaining ambiguous cases and confirmed RI ineligibility (1.8% of RI transactions).

CI coverage is imputed from the nearest same-province observation by date. Premiums are not recorded in 2017–2023, so we impute $I^c(k_i) = ak_i + b$ from 2025 Taobao data. Agents do not internalize current returns into future premiums.[35]

A.3 Aggregation, Filtering and Smoothing

We aggregate the cleaned transactions to category × month cells, counting sales and averaging transaction prices within each cell. We keep categories whose average monthly sales over the sample reach at least fifty transactions. The six qualifying categories are dress, tops, pants, sweater, two-piece, and outerwear.

Monthly sales and prices are smoothed with a five-month centered moving average that combines the current month with the two adjacent months on each side, using the available months only at the ends of the sample. Each category-month is then resampled to its smoothed sales count, drawing from the month's observed transactions without replacement when observed sales exceed the smoothed count and

[35] We treat the 2025 premium-coverage slope as stable. It implies premiums of roughly 60% of coverage, consistent with older public documentation (e.g. Zhihu).

with replacement otherwise. Every resampled transaction in a category-month is assigned the smoothed price, so RI and CI transactions in the same month share one list price and therefore one net handling cost $h_t = 15 - 0.75\, p_t$.

Market size is twice the peak smoothed monthly sales within the category-season, with seasons defined as March–May, June–August, September–November, and December–February. Monthly purchase probabilities equal smoothed sales divided by market size. We drop category-months with purchase probability below 5% to avoid the zero-share problem.

When both schemes appear within a month, we compute each scheme's time share as the share of the month's transaction-active hours during which the scheme was in effect, common across categories. The observed purchase, return, and insurance-purchase counts under each scheme are divided by that time share, converting them into full-month counts at the common smoothed price. The scale factor is capped at the market size divided by twice the scheme's observed transaction count, so no scheme's full-month volume exceeds half the market size and purchase probabilities remain interior. Because a single factor multiplies every coverage × purchase × return (× insurance) cell of the scheme-month, the scaling preserves the return rate per purchase and the insurance take-up rate within the scheme-month.

B Identification of Ex-Post Precision

This appendix derives the expression for expected returns under RI and proves that, holding the purchase schedule fixed, expected returns are strictly decreasing in ex-post precision σ_t.

Under RI, return shipping is fully reimbursed up to k_i, yielding mean return utility $\overline{u}_{it}^{\text{return}} = -\kappa$. A consumer purchases if and only if $\nu_{it} \geq \hat{\nu}_t^r \equiv p_t + \frac{1}{\sigma_t}\ln(1 - e^{-\sigma_t\kappa})$. Defining willingness to pay as $w_{it}^r \equiv \nu_{it} - \frac{1}{\sigma_t}\ln(1-e^{-\sigma_t\kappa})$, the purchase condition is equivalent to $w_{it}^r \geq p_t$. Because $\nu_{it} \sim \mathcal{N}(\mu_t + \xi_t, \omega_j^2)$, w_{it}^r is normally distributed with mean $\bar{w}_t^r = \mu_t + \xi_t - \frac{1}{\sigma_t}\ln(1 - e^{-\sigma_t\kappa})$ and standard deviation ω_j.

Conditional on purchase, the probability of return is

$$q_{it}^r(\nu_{it}) = \frac{e^{-\sigma_t\kappa}}{e^{\sigma_t(\nu_{it}-p_t)} + e^{-\sigma_t\kappa}}.$$

Substituting $\nu_{it} - p_t = (w_{it}^r - p_t) + \frac{1}{\sigma_t}\ln(1 - e^{-\sigma_t\kappa})$ into the denominator and dividing both numerator and denominator by $e^{-\sigma_t\kappa}$ simplifies the individual return probability to

$$q_{it}^r(w_{it}^r) = \frac{1}{1 + (e^{\sigma_t\kappa} - 1)e^{\sigma_t(w_{it}^r - p_t)}}. \tag{14}$$

Integrating over the distribution of w_{it}^r yields the unconditional expected return rate per potential buyer:

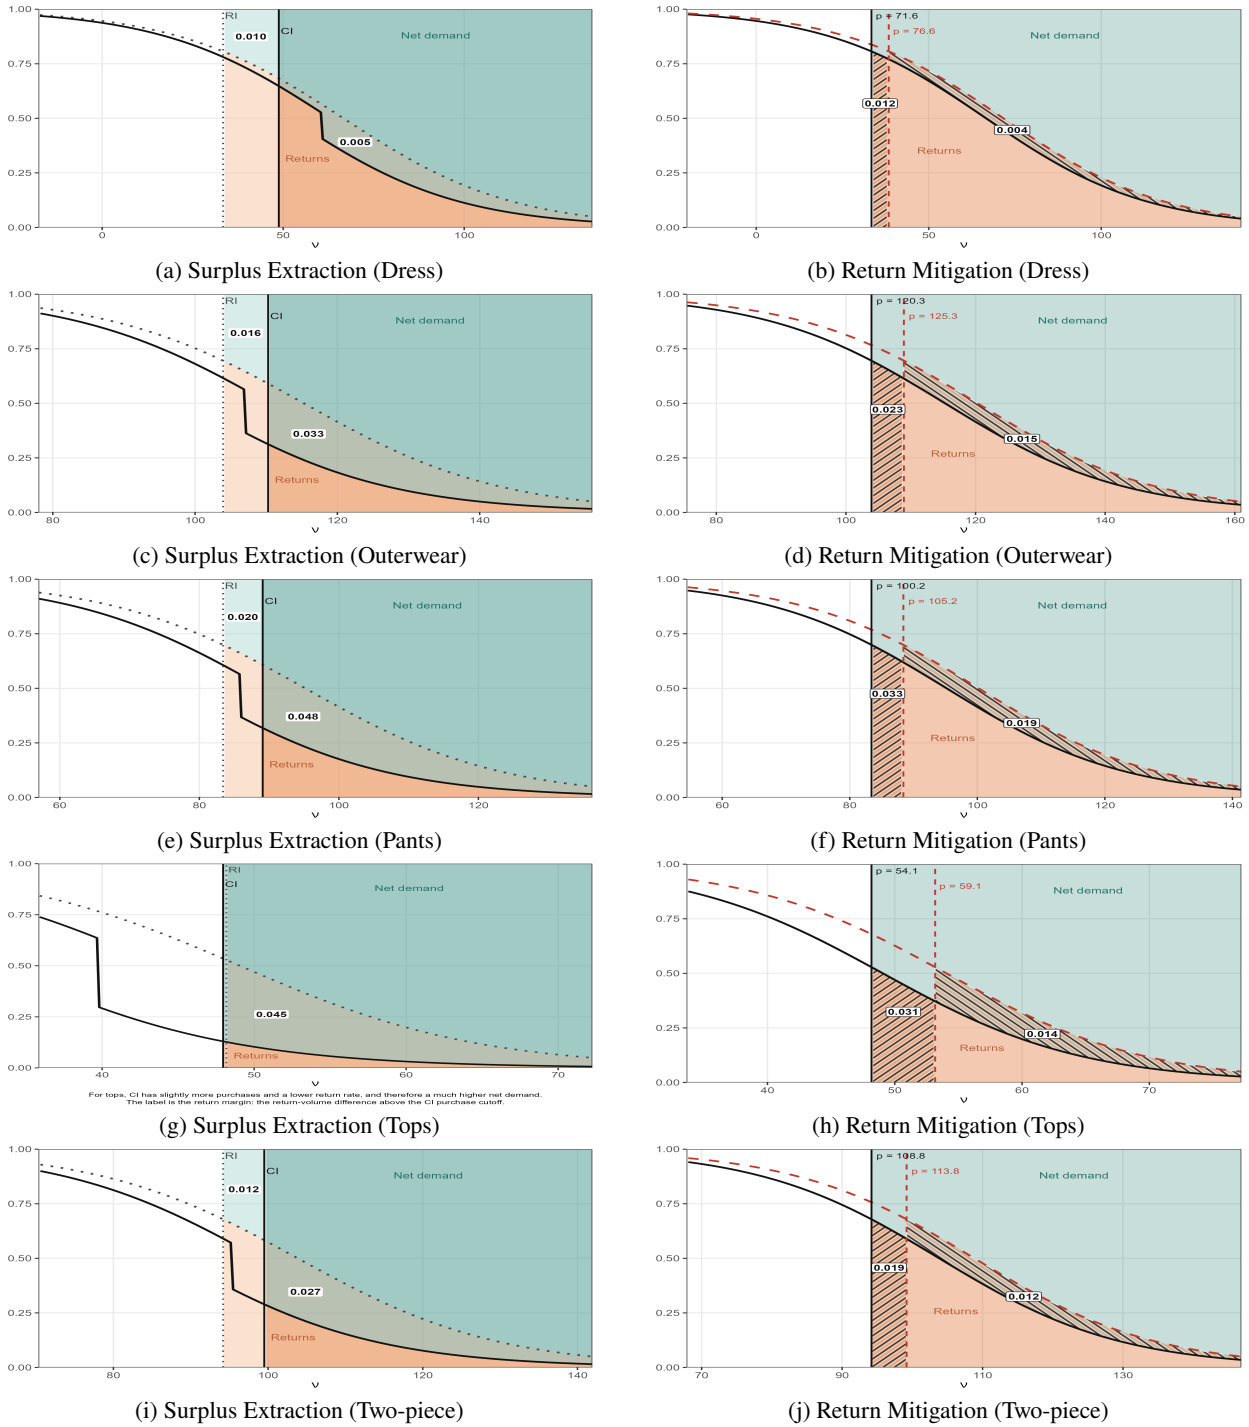
$$\mathbb{E}[R_{it}^r(p_t)] = \mathbb{E}_{w_{it}^r}\left[\frac{1}{1+(e^{\sigma_t\kappa}-1)e^{\sigma_t(w_{it}^r-p_t)}}\,\mathbf{1}\big(w_{it}^r \geq p_t\big)\right]. \tag{15}$$

The return rate conditional on purchase is simply $\mathbb{E}[R_{it}^r(p_t)]/\mathbb{E}[D_{it}^r(p_t)]$.

As shown in Section 5, price variation identifies ω_j and $\bar{w}_t^r$, holding the distribution of w_{it}^r and the purchase probability $\mathbb{E}[D_{it}^r(p_t)] = 1 - \Phi\big(\frac{p_t-\bar{w}_t^r}{\omega_j}\big)$ fixed. For every purchaser ($w_{it}^r \geq p_t$), $w_{it}^r - p_t \geq 0$. Because $\kappa > 0$, the term $e^{\sigma_t\kappa} - 1 > 0$ is strictly increasing in σ_t, while $e^{\sigma_t(w_{it}^r-p_t)} \geq 1$ is weakly increasing in σ_t. The denominator of (14) is therefore strictly increasing in σ_t, so each purchaser's return probability is pointwise strictly decreasing in σ_t. Because the distribution of w_{it}^r and the purchase region $w_{it}^r \geq p_t$ are held fixed, the expected return rate (15) (as well as the return rate conditional on purchase) is strictly decreasing in σ_t. Consequently, σ_t is uniquely identified by the return rate, after which intrinsic mean valuation $\mu_t + \xi_t = \bar{w}_t^r + \frac{1}{\sigma_t}\ln(1 - e^{-\sigma_t\kappa})$ is uniquely recovered.

C Surplus Extraction and Return Mitigation by Category

Figure 10 pairs surplus extraction and return mitigation for the categories other than sweaters. In each row the left panel is surplus extraction at equilibrium prices and the right panel is return mitigation under RI.



(a) Surplus Extraction (Dress)

(b) Return Mitigation (Dress)

(c) Surplus Extraction (Outerwear)

(d) Return Mitigation (Outerwear)

(e) Surplus Extraction (Pants)

(f) Return Mitigation (Pants)

(g) Surplus Extraction (Tops)

(h) Return Mitigation (Tops)

(i) Surplus Extraction (Two-piece)

(j) Return Mitigation (Two-piece)

Figure 10: Left panels are surplus extraction at equilibrium prices; dotted lines are RI and solid lines are CI. Right panels are return mitigation under RI.

D Return-Policy Extensions

D.1 Return-Penalty Discount: Equilibrium

RPD offers a discount branch at $p - d$ with forfeiture share $\phi \in [0, 1]$ of $p - d$ on return, or a full-price branch at p with a full product refund. The buyer bears $\kappa + k_i$ on either return; the discount branch recovers $(1 - \phi)(p - d)$.

Expected surplus on the two branches is

$$CS_i^f = \frac{1}{\sigma} \ln \left(e^{\sigma(\nu_i - p)} + e^{-\sigma(\kappa + k_i)} \right), \tag{16}$$

$$CS_i^d = d - p + \frac{1}{\sigma} \ln \left(e^{\sigma \nu_i} + e^{\sigma[(1-\phi)(p-d) - (\kappa + k_i)]} \right). \tag{17}$$

The consumer prefers the discount branch when $CS_i^d \geq CS_i^f$, which defines

$$\tilde{\nu}^{RPD} = -(\kappa + k_i) + \frac{1}{\sigma} \ln \left(\frac{e^{\sigma(p-d)} - e^{\sigma(1-\phi)(p-d)}}{1 - e^{-\sigma d}} \right). \tag{18}$$

Purchase thresholds are

$$\hat{\nu}^d = \frac{1}{\sigma} \ln \left(e^{\sigma(p-d)} - e^{\sigma[(1-\phi)(p-d) - (\kappa + k_i)]} \right), \tag{19}$$

$$\hat{\nu}^f \equiv p + \frac{1}{\sigma} \ln \left(1 - e^{-\sigma(\kappa + k_i)} \right). \tag{20}$$

The consumer purchases when $\max\{CS_i^d, CS_i^f\} \geq 0$. Both branches are active when $\tilde{\nu}^{RPD} \geq \hat{\nu}^f$, or equivalently (when $1 - e^{-\sigma d} > 0$)

$$\frac{e^{\sigma(p-d)} - e^{\sigma(1-\phi)(p-d)}}{1 - e^{-\sigma d}} \geq e^{\sigma p} \left(e^{\sigma(\kappa + k_i)} - 1 \right). \tag{21}$$

Conditional on purchase,

$$q_i^d = \frac{e^{\sigma[(1-\phi)(p-d) - (\kappa + k_i)]}}{e^{\sigma \nu_i} + e^{\sigma[(1-\phi)(p-d) - (\kappa + k_i)]}}, \tag{22}$$

$$q_i^f = \frac{e^{-\sigma(\kappa + k_i)}}{e^{\sigma(\nu_i - p)} + e^{-\sigma(\kappa + k_i)}}. \tag{23}$$

With branch indicators D_i^d, D_i^f and aggregates D^d, R^d, D^f, R^f, the seller chooses (p, d) to maximize $\pi^{RPD} = \pi^d + \pi^f$, where $\pi^d = (p - d - \theta^c) D^d - [(1 - \phi)(p - d) + h] R^d$ and $\pi^f = (p - \theta^c) D^f - (p + h) R^f$. Individual social surplus on each branch is

$$S_i^d \equiv CS_i^d + (p - d - \theta^c) - q_i^d \left[(1 - \phi)(p - d) + h \right], \tag{24}$$

$$S_i^f \equiv CS_i^f + (1 - q_i^f)(p - \theta^c) - q_i^f (\theta^c + h). \tag{25}$$

Define $CS^{RPD} \equiv \mathbb{E}_{\nu_i,k_i}[D_i^d\, CS_i^d + D_i^f\, CS_i^f]$ and $W^{RPD} \equiv \mathbb{E}_{\nu_i,k_i}[D_i^d\, S_i^d + D_i^f\, S_i^f]$. The main-text full-forfeiture case is $\phi = 1$; partial forfeiture varies ϕ.

D.2 Optimal Refund Policy

With list price p and return refund $-h$,

$$CS_i^{\text{rs}} = -p + \frac{1}{\sigma}\ln\left(e^{\sigma\nu_i} + e^{-\sigma(h+\kappa+k_i)}\right), \tag{26}$$

$$\hat{\nu}^{\text{rs}} = \frac{1}{\sigma}\ln\left(e^{\sigma p} - e^{-\sigma(h+\kappa+k_i)}\right), \tag{27}$$

where $\hat{\nu}^{\text{rs}}$ solves $CS_i^{\text{rs}} = 0$ when $e^{\sigma p} > e^{-\sigma(h+\kappa+k_i)}$. The consumer purchases when $\nu_i \geq \hat{\nu}^{\text{rs}}$ and returns with $q_i^{\text{rs}} = q_i^*$ from (11). Aggregates are

$$D^{\text{rs}} \equiv \mathbb{E}_{\nu_i,k_i}\left[\mathbf{1}\left(\nu_i \geq \hat{\nu}^{\text{rs}}\right)\right], \tag{28}$$

$$R^{\text{rs}} \equiv \mathbb{E}_{\nu_i,k_i}\left[\mathbf{1}\left(\nu_i \geq \hat{\nu}^{\text{rs}}\right) q_i^{\text{rs}}\right], \tag{29}$$

and the seller chooses p to maximize $\pi^{\text{rs}}(p) = (p - \theta^c)D^{\text{rs}}$, with $S_i^{\text{rs}} = CS_i^{\text{rs}} + (p - \theta^c)$.

E Scrap-Value Sensitivity

The main text varies scrap to $0.65\, p_t$ and $0.85\, p_t$. Each panel reports welfare as a percent of first best, then repeats the normalized policy decomposition and the partial-RPD ϕ grid.

E.1 Scrap value 0.65 p_t

Table 8: Welfare as a percent of first best. Scrap 0.65 p_t.

				Partial forfeiture ϕ			
Category	RI	CI	Full RPD	0.05	0.10	0.15	0.20
dress	68.3	74.2	61.6	73.8	73.3	72.5	71.8
outerwear	59.6	73.3	74.6	74.6	75.1	75.3	75.3
pants	53.6	72.7	76.8	74.4	77.4	75.5	79.3
sweater	31.5	68.4	76.4	73.6	77.5	78.0	76.5
top	42.9	72.4	75.0	73.6	74.3	74.7	75.0
twopiece	59.1	73.1	74.3	74.2	74.8	75.0	75.0

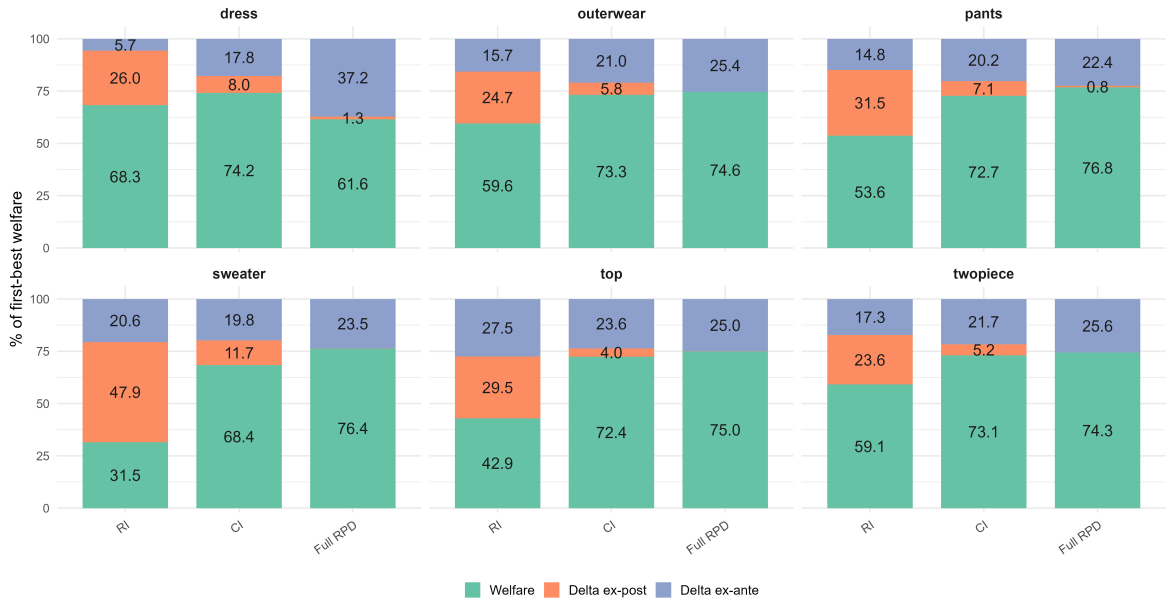


Figure 11: Welfare decomposition as a share of first best. Scrap 0.65 p_t.

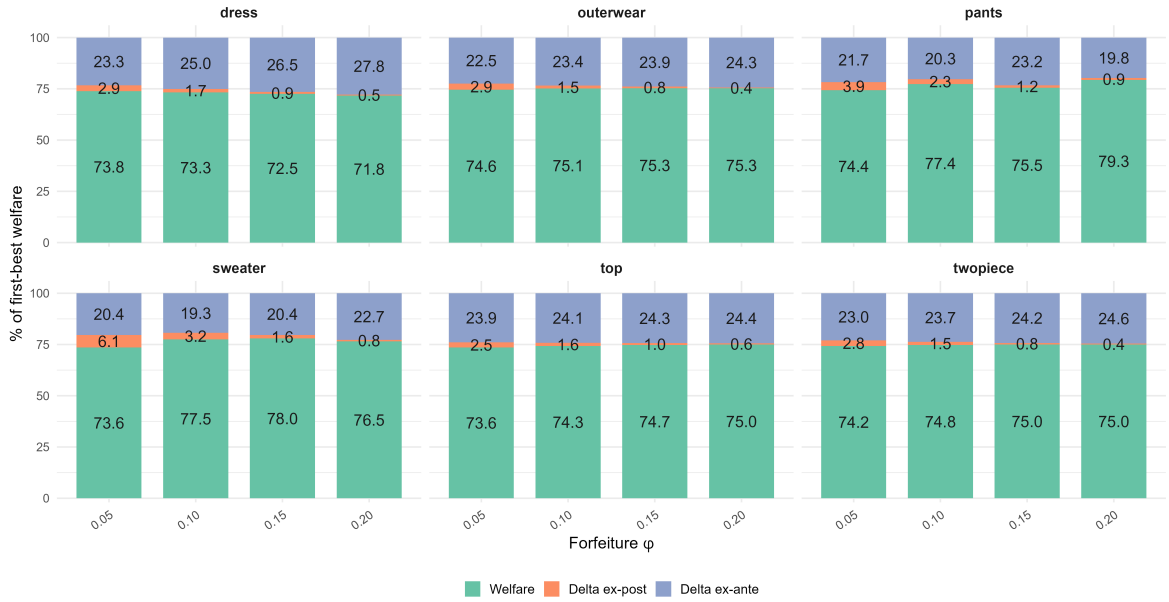


Figure 12: Partial-forfeiture RPD by ϕ. Scrap 0.65 p_t.

E.2 Scrap value 0.85 p_t

Table 9: Welfare as a percent of first best. Scrap 0.85 p_t.

				Partial forfeiture ϕ			
Category	RI	CI	Full RPD	0.05	0.10	0.15	0.20
dress	79.1	68.8	46.7	64.2	62.4	60.7	59.2
outerwear	72.1	75.2	71.5	74.8	74.2	73.7	73.2
pants	68.8	75.5	71.3	75.4	75.0	74.5	74.0
sweater	53.7	74.4	74.4	76.6	77.0	77.0	75.8
top	55.4	74.5	74.1	74.9	75.1	75.1	75.1
twopiece	71.1	74.9	71.8	74.7	74.2	73.8	73.4

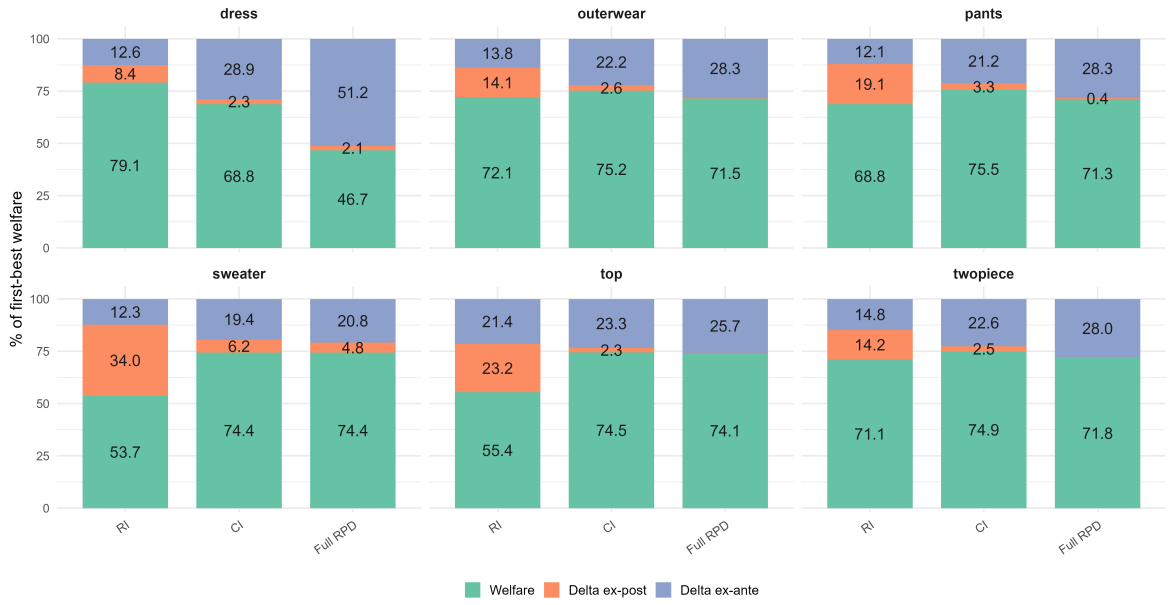


Figure 13: Welfare decomposition as a share of first best. Scrap 0.85 p_t.

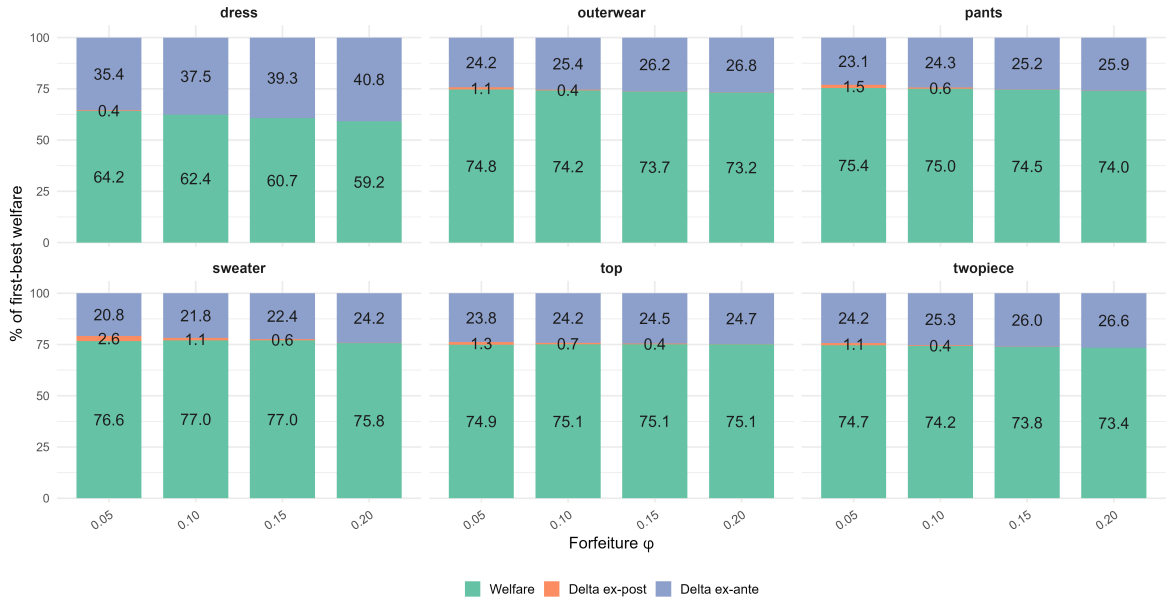
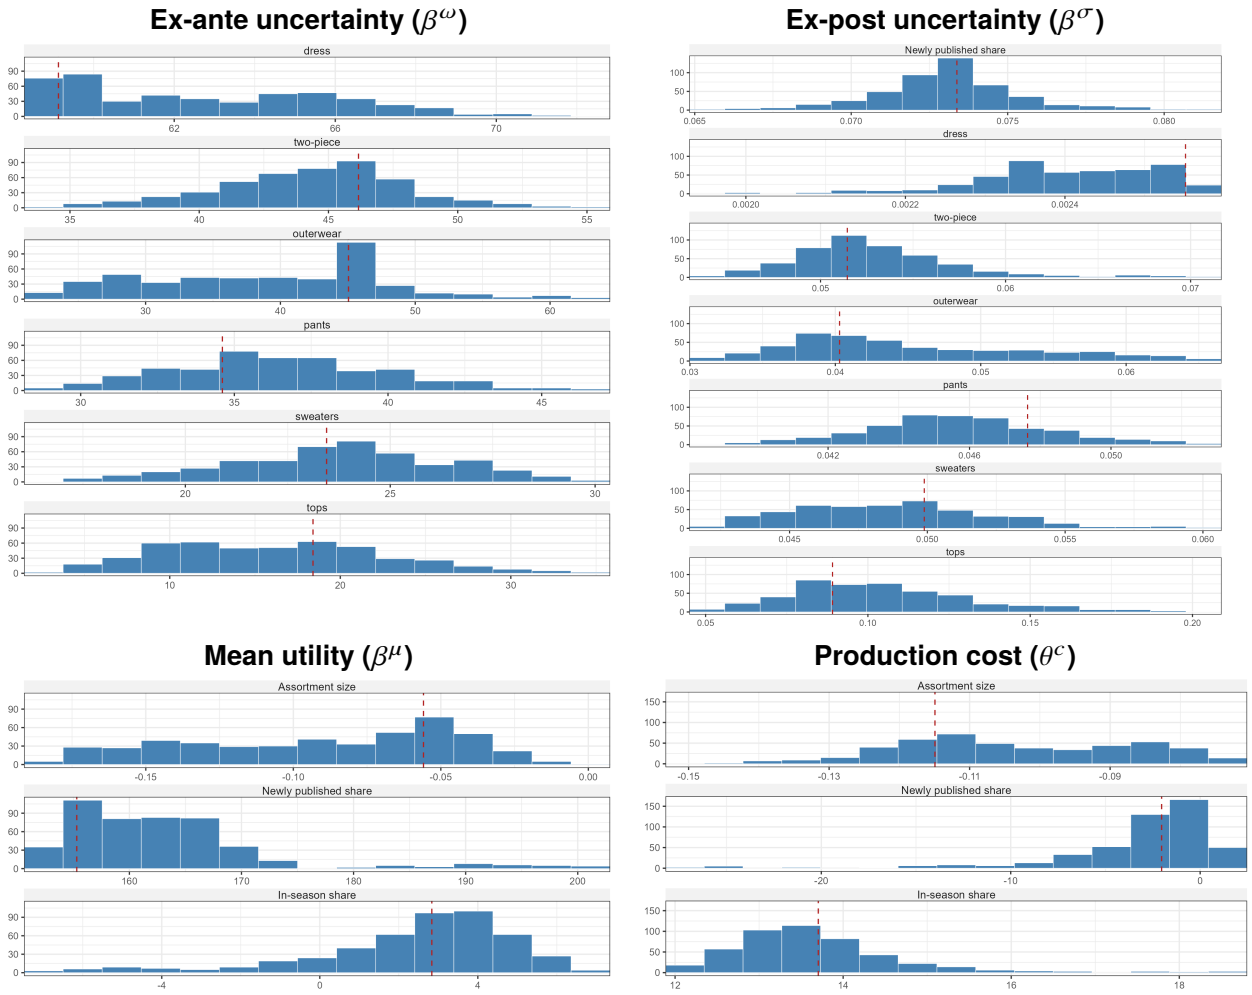


Figure 14: Partial-forfeiture RPD by ϕ. Scrap 0.85 p_t.

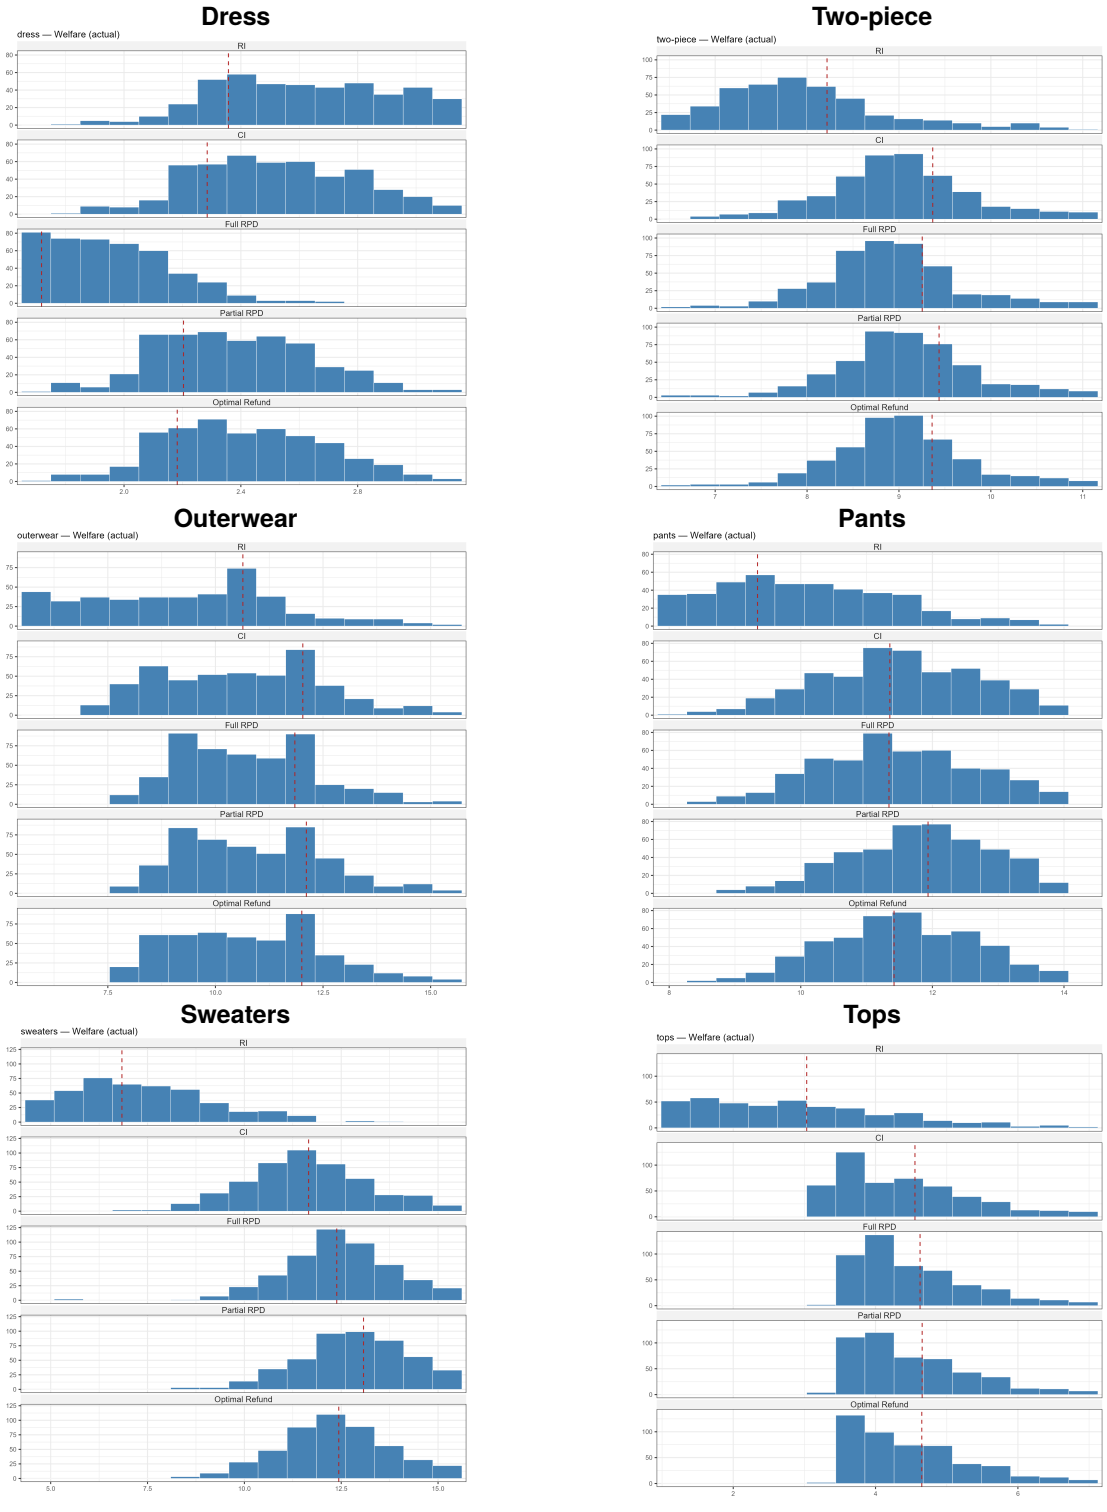
F Bootstrap Inference

Each panel plots bootstrap replicates in 15 bins. Axes span the 2.5th–97.5th percentile of the displayed draws (by coefficient for parameters; pooled across policies within a category for welfare, scaled to first best). The dashed line is the full-sample point estimate. Main-text standard errors use all replicates.

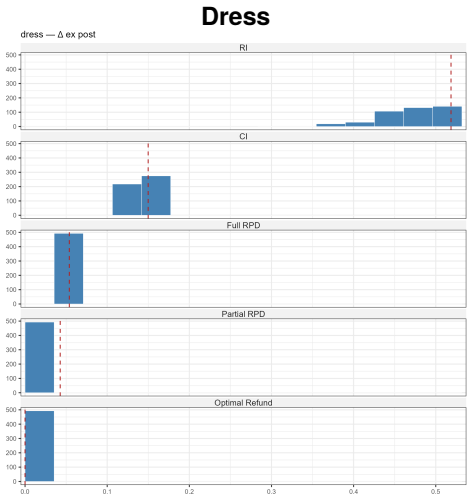
F.1 Point estimates

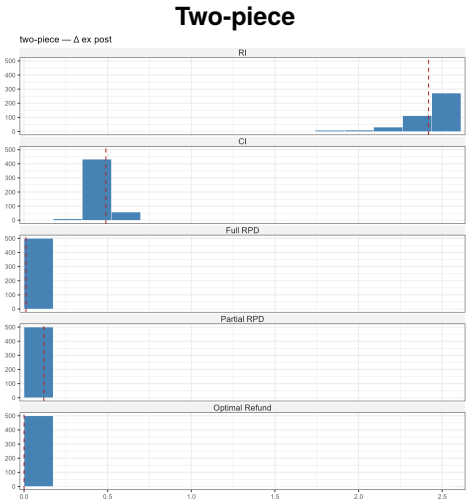


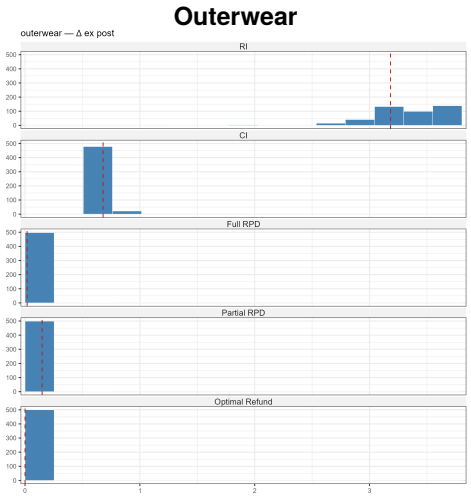
F.2 Actual welfare

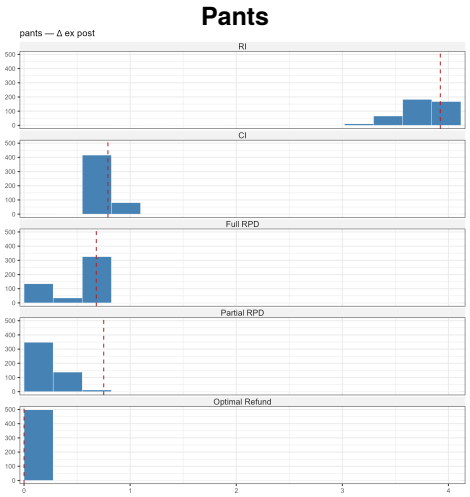


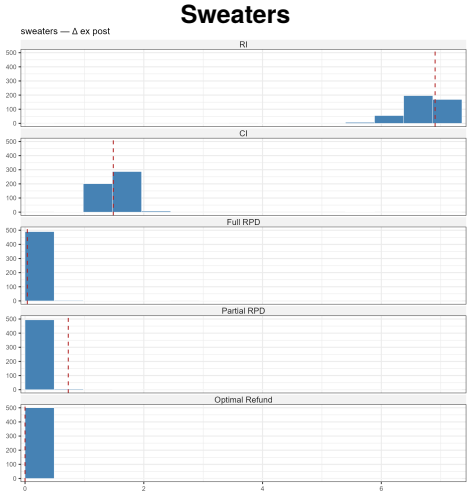
F.3 Δ ex-post

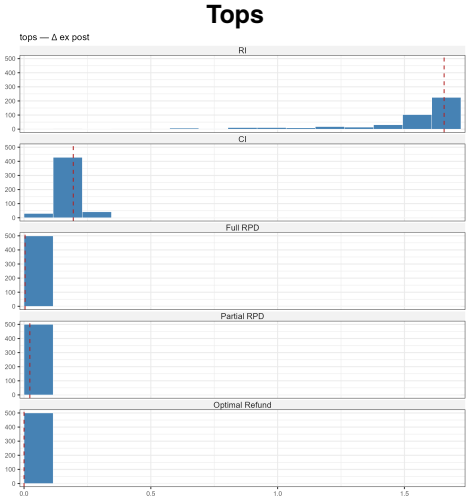

Dress
dress — Δ ex post
RI
CI
Full RPD
Partial RPD
Optimal Refund


Two-piece
two-piece — Δ ex post
RI
CI
Full RPD
Partial RPD
Optimal Refund

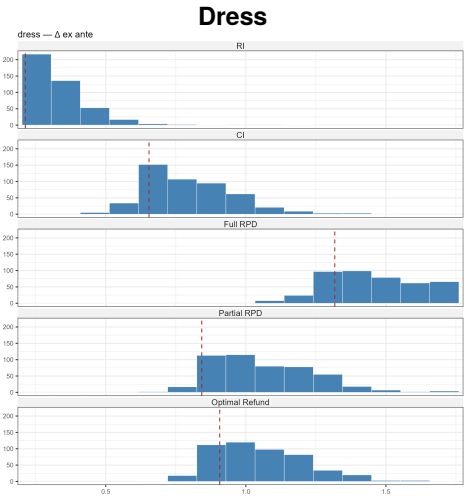

Outerwear
outerwear — Δ ex post
RI
CI
Full RPD
Partial RPD
Optimal Refund

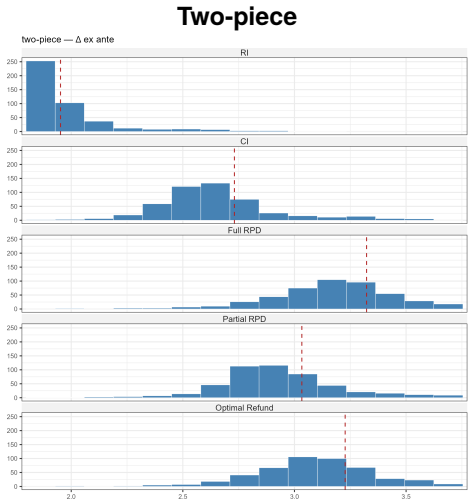

Pants
pants — Δ ex post
RI
CI
Full RPD
Partial RPD
Optimal Refund

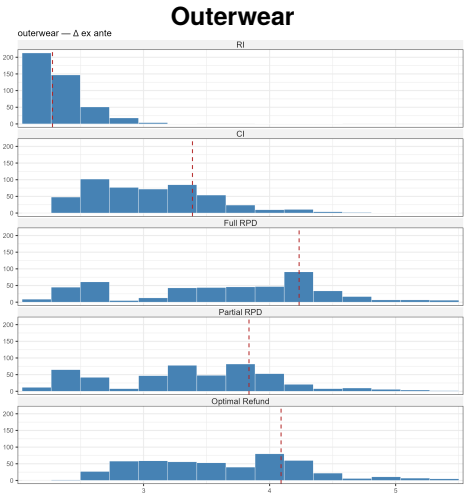

Sweaters
sweaters — Δ ex post
RI
CI
Full RPD
Partial RPD
Optimal Refund

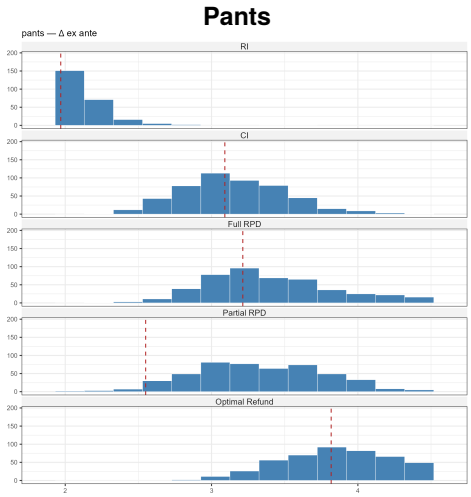

Tops
tops — Δ ex post
RI
CI
Full RPD
Partial RPD
Optimal Refund

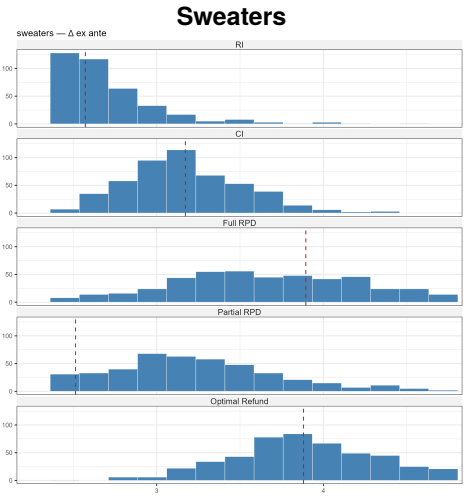
F.4 Δ ex-ante

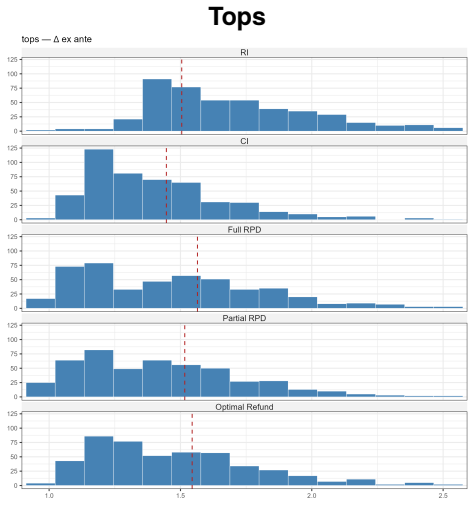

Dress
dress — Δ ex ante
RI
CI
Full RPD
Partial RPD
Optimal Refund


Two-piece
two-piece — Δ ex ante
RI
CI
Full RPD
Partial RPD
Optimal Refund


Outerwear
outerwear — Δ ex ante
RI
CI
Full RPD
Partial RPD
Optimal Refund


Pants
pants — Δ ex ante
RI
CI
Full RPD
Partial RPD
Optimal Refund


Sweaters
sweaters — Δ ex ante
RI
CI
Full RPD
Partial RPD
Optimal Refund


Tops
tops — Δ ex ante
RI
CI
Full RPD
Partial RPD
Optimal Refund

F.5 Pairwise win rates

Table 10: Pairwise bootstrap win rates for actual welfare.

Dress

	RI	CI	Full RPD	Partial RPD	Optimal Refund
RI	—	0.99	1.00	1.00	1.00
CI	0.01	—	1.00	1.00	1.00
Full RPD	0.00	0.00	—	0.00	0.00
Partial RPD	0.00	0.00	1.00	—	0.01
Optimal Refund	0.00	0.00	1.00	0.99	—

Two-piece

	RI	CI	Full RPD	Partial RPD	Optimal Refund
RI	—	0.00	0.01	0.01	0.01
CI	1.00	—	0.83	0.03	0.38
Full RPD	0.99	0.17	—	0.00	0.00
Partial RPD	0.99	0.97	1.00	—	1.00
Optimal Refund	0.99	0.62	1.00	0.00	—

Outerwear

	RI	CI	Full RPD	Partial RPD	Optimal Refund
RI	—	0.02	0.04	0.02	0.03
CI	0.98	—	0.55	0.09	0.38
Full RPD	0.96	0.45	—	0.07	0.23
Partial RPD	0.98	0.91	0.93	—	0.99
Optimal Refund	0.97	0.62	0.77	0.01	—

Pants

	RI	CI	Full RPD	Partial RPD	Optimal Refund
RI	—	0.01	0.01	0.01	0.00
CI	0.99	—	0.40	0.03	0.42
Full RPD	0.99	0.60	—	0.01	0.36
Partial RPD	0.99	0.97	0.99	—	0.99
Optimal Refund	1.00	0.58	0.64	0.01	—

Sweaters

	RI	CI	Full RPD	Partial RPD	Optimal Refund
RI	—	0.01	0.01	0.01	0.01
CI	0.99	—	0.01	0.01	0.01
Full RPD	0.99	0.99	—	0.11	0.54
Partial RPD	0.99	0.99	0.89	—	0.97
Optimal Refund	0.99	0.99	0.46	0.03	—

Tops

	RI	CI	Full RPD	Partial RPD	Optimal Refund
RI	—	0.00	0.00	0.00	0.00
CI	1.00	—	0.14	0.03	0.07
Full RPD	1.00	0.86	—	0.37	0.25
Partial RPD	1.00	0.97	0.63	—	0.56
Optimal Refund	1.00	0.93	0.75	0.44	—

Notes: Each off-diagonal cell is the share of bootstrap draws in which the row policy has higher actual welfare than the column policy (non-missing pairs only). Diagonal cells are blank.

Table 11: Pairwise bootstrap win rates for ex-post inefficiency.

Dress

	RI	CI	Full RPD	Partial RPD	Optimal Refund
RI	—	0.00	0.00	0.00	0.00
CI	1.00	—	0.00	0.00	0.00
Full RPD	1.00	1.00	—	0.00	0.00
Partial RPD	1.00	1.00	1.00	—	0.00
Optimal Refund	1.00	1.00	1.00	1.00	—

Two-piece

	RI	CI	Full RPD	Partial RPD	Optimal Refund
RI	—	0.00	0.00	0.00	0.00
CI	1.00	—	0.00	0.00	0.00
Full RPD	1.00	1.00	—	1.00	0.00
Partial RPD	1.00	1.00	0.00	—	0.00
Optimal Refund	1.00	1.00	1.00	1.00	—

Outerwear

	RI	CI	Full RPD	Partial RPD	Optimal Refund
RI	—	0.00	0.00	0.00	0.00
CI	1.00	—	0.00	0.00	0.00
Full RPD	1.00	1.00	—	0.96	0.00
Partial RPD	1.00	1.00	0.04	—	0.00
Optimal Refund	1.00	1.00	1.00	1.00	—

Pants

	RI	CI	Full RPD	Partial RPD	Optimal Refund
RI	—	0.00	0.00	0.00	0.00
CI	1.00	—	0.01	0.00	0.00
Full RPD	1.00	0.99	—	0.27	0.00
Partial RPD	1.00	1.00	0.73	—	0.00
Optimal Refund	1.00	1.00	1.00	1.00	—

Sweaters

	RI	CI	Full RPD	Partial RPD	Optimal Refund
RI	—	0.00	0.00	0.00	0.00
CI	1.00	—	0.00	0.00	0.00
Full RPD	1.00	1.00	—	0.93	0.00
Partial RPD	1.00	1.00	0.07	—	0.00
Optimal Refund	1.00	1.00	1.00	1.00	—

Tops

	RI	CI	Full RPD	Partial RPD	Optimal Refund
RI	—	0.00	0.00	0.00	0.00
CI	1.00	—	0.00	0.00	0.00
Full RPD	1.00	1.00	—	0.99	0.00
Partial RPD	1.00	1.00	0.01	—	0.00
Optimal Refund	1.00	1.00	1.00	1.00	—

Notes: Each off-diagonal cell is the share of bootstrap draws in which the row policy has lower ex-post inefficiency than the column policy. A higher number means the row policy is more efficient ex post. Diagonal cells are blank.

Table 12: Pairwise bootstrap win rates for ex-ante inefficiency.

Dress

	RI	CI	Full RPD	Partial RPD	Optimal Refund
RI	—	1.00	1.00	1.00	1.00
CI	0.00	—	1.00	1.00	1.00
Full RPD	0.00	0.00	—	0.00	0.00
Partial RPD	0.00	0.00	1.00	—	0.28
Optimal Refund	0.00	0.00	1.00	0.72	—

Two-piece

	RI	CI	Full RPD	Partial RPD	Optimal Refund
RI	—	0.99	1.00	1.00	1.00
CI	0.01	—	1.00	1.00	1.00
Full RPD	0.00	0.00	—	0.00	0.00
Partial RPD	0.00	0.00	1.00	—	1.00
Optimal Refund	0.00	0.00	1.00	0.00	—

Outerwear

	RI	CI	Full RPD	Partial RPD	Optimal Refund
RI	—	0.93	0.91	0.89	0.97
CI	0.07	—	0.83	0.77	1.00
Full RPD	0.09	0.17	—	0.04	0.23
Partial RPD	0.11	0.23	0.96	—	1.00
Optimal Refund	0.03	0.00	0.77	0.00	—

Pants

	RI	CI	Full RPD	Partial RPD	Optimal Refund
RI	—	1.00	0.99	0.99	0.99
CI	0.00	—	0.93	0.79	0.99
Full RPD	0.01	0.07	—	0.38	0.73
Partial RPD	0.01	0.21	0.62	—	0.99
Optimal Refund	0.01	0.01	0.27	0.01	—

Sweaters

	RI	CI	Full RPD	Partial RPD	Optimal Refund
RI	—	0.88	0.91	0.79	0.95
CI	0.12	—	0.92	0.51	1.00
Full RPD	0.09	0.08	—	0.08	0.65
Partial RPD	0.21	0.49	0.92	—	1.00
Optimal Refund	0.05	0.00	0.35	0.00	—

Tops

	RI	CI	Full RPD	Partial RPD	Optimal Refund
RI	—	0.31	0.41	0.38	0.40
CI	0.69	—	0.69	0.67	0.89
Full RPD	0.59	0.31	—	0.19	0.26
Partial RPD	0.62	0.33	0.81	—	0.97
Optimal Refund	0.60	0.11	0.74	0.03	—

Notes: Each off-diagonal cell is the share of bootstrap draws in which the row policy has lower ex-ante inefficiency than the column policy. A higher number means the row policy is more efficient ex ante. Diagonal cells are blank.