The Welfare Effects of Behavior-based Promotion Discrimination in E-commerce*

Francesco Gabriele†

Version: November 24, 2025

Abstract

E-commerce retailers have the power to personalize promotions based on users' online past purchase behavior. This paper develops a structural model of consumer demand and a policy model of personalized promotions to quantify the welfare effects of behavior-based promotion discrimination (BBPD). Using data from a randomized controlled field experiment with coupons on a cosmetics e-commerce site, the structural estimation reveals an elastic and heterogeneous demand to price discount treatments. With nonparametric estimation via machine learning, the counterfactual analysis tests different targeting algorithms and shows that promotion personalization increases e-commerce profit by 24% and consumer surplus by 4% relative to the uniform pricing benchmark. Exploiting past purchase information is profitable for the e-commerce: BBPD improves personalized coupons algorithm by generating an additional 11% gain in producer surplus without harming loyal customers. This paper contributes to the current public policy debate about personalized pricing strategies in digital markets as the welfare analysis has implications for privacy and fairness concerns.

Keywords: Personalized Promotions, Targeted Marketing, Field Experiment, Structural Demand, Policy Model, Causal Machine Learning, Welfare Analysis, Data Privacy

*Acknowledgments This paper is part of my Ph.D. dissertation. I am truly indebted to Paulina Oliva, Robert Metcalfe, Davide Proserpio and Giorgio Coricelli for their invaluable advising during my doctoral studies. I am grateful to Timothy Armstrong, Tommaso Bondi, David Byrne, Thomas Chaney, Stefano Colombo, Dante Donati, Anthony Dukes, Domenico Fabrizi, Chiara Farronato, Michele Fioretti, Matthew Gentzkow, Dong Woo Hahm, Brett Hollenbeck, Song Lin, Marco Lo Faso, Noriaki Matsushima, Alex Miller, Olivia Natan, Hashem Pesaran, Geert Ridder, Esteban Rossi-Hansberg, Stephan Seiler, Mengze Shi, Benjamin Shiller, Daniel Sokol, Guofu Tan, Miguel Villas-Boas and Angela Zhou. This paper benefited from the comments of the participants to seminars at USC, USC Marshall, USF, CalState Long Beach, Università Bocconi, Università Cattolica, Compass LexEcon London and at the 2024 ACLEC, 2024 AFE, 51st EARIE, 13th IBEO, 22nd ZEW Conference on the Economics of Information, 7th Doctoral Workshop on The Economics of Digitization, 22nd IIOC, 2024 Imperial College PhD Conference, 2024 WoPA, 16th TSE Digital Economics Conference, 2023 SOCAE, 2023 Milan PhD Economics Workshop and 2023 TADC, where previous versions of this paper were presented. I acknowledge funding from USC Department of Economics and USC Marshall Initiative on Digital Competition, and technical support from USC Stevens Center for Innovation. At last, I thank collaborators at *Alpha*. All errors and omissions are my own.

†Affiliation and e-mail Hong Kong University of Science and Technology, Busiess School, Department of Marketing. E-mail address: fgabriele@ust.hk.

1 Introduction

With the growing digitization of market transactions, the development of customers' recognition technologies has enabled digital companies to collect thorough data about their users. As digital information has immense economic potential, nowadays users' tracking is predominant in many industries ranging from online retailers to tech apps.[1] The ability of online platforms to identify customers allows them to offer different prices to different individuals or segments based on their observable characteristics.[2] Behavioral information underpins targeted marketing strategies such as personalized pricing and promotions. In particular, the marketing strategies that offer price discounts based on customers' past purchase history are regarded as behavior-based promotion discrimination (BBPD).[3]

A new era in the regulation of personal data privacy and deceptive conduct of digital companies began in 2012 with the FTC's passage of the E-STOP Act.[4] Lately, the mounting concerns and public debate over the commercial use of customers' data, led to the introduction of the California Consumer Privacy Act in US, and the EU adoption of the General Data Protection Regulation in 2018. Although these regulations aim at preserving consumers' privacy by enhancing the data protection obligations for digital businesses, discriminatory practices are not utterly prevented. These policies indeed require the clients to explicitly give their consent before receiving personalized contents, which limits access to users' information and hinders the effective implementation of marketing strategies. However, while setting different prices based on customers' personal data violates these obligations, offering personalized promotions or targeted discounts does not constitute an infringement (Baik & Larson, 2023).

Further disagreement concerns the implications of these practices on the customers' welfare. The recent academic debate has questioned to what extent personalized promotional strategies are equivalent to price discrimination and, consequently, what are the welfare implications of such marketing practices. If on the one hand price discrimination has been shown to benefit the monopolistic producer, on the other hand it is still a matter of research to quantify the consequences on consumer surplus (e.g., Bergemann et al., 2015; Kallus & Zhou, 2021; De Cornière et al., 2023). Different fields of research have studied personalized pricing. A large body of theoretical works has examined how behavior-based price

[1] As reported by Colombo (2016), examples of companies operating in different industries which price discriminate across consumers on the basis of their online behavior are Amazon, eBay, Wal-Mart, Netflix, and Uber. Also Pazgal and Soberman (2008) and Esteves (2014) discuss various examples of industries in which companies collect, process and utilize users' data for profiling and customization purposes.

[2] Price discrimination refers to the firms' power to effectively price different market segments based on directly observable consumers' characteristics. A recent episode, for instance, saw McDonald's being accused of price profiling. Taunton (2023) reports that the fast food chain was criticized for price-gouging after some customers were offered menus at higher prices though the McDonald's app. The comparison of two screenshots reveals that the same meal deal was offered at $14 and at $17. In addition, the combo of McFlurry desserts had two different prices: $7 or $8.50 (Figure D.1). It was also reported that McDonald's spokesperson replied that app offers differ between users, based on a variety of factors. McDonald's spokesperson expressly claimed: "*Due to the personalization of our app, not all customers will see the same deals, and as an example a deal may be offered to encourage use of the app on the customer's next visit*".

[3] An example of this practice is illustrated by Streitfeld (2000) reporting how in 2000 Amazon was accused of selling the same DVDs at different prices based on customers' purchase histories. By deleting the online cookies identifying them as regular clients, customers could observe the DVDs' price fall from $26.24 to $22.74.

[4] The Ensuring Shoppers Transparency in Online Pricing (E-STOP) Act of 2012 directs the Federal Trade Commission to promulgate rules requiring an Internet merchant with annual gross revenue of more than $1 million to disclose to each consumer, prior to the final purchase of any good or service, the use of a price-altering computer program, defined as one that: (i) accesses a consumer's personal information; and (ii) uses such information to alter the merchant's selling price of a good or service.

discrimination affects profit and welfare in competitive markets (e.g., Esteves, 2010; Chen & Pearcy, 2010; Choe et al., 2018; Esteves et al., 2022).[5] Whereas, a more limited applied literature has attempted to empirically test such theoretical predictions using controlled experiments (e.g., Rossi et al., 1996; Mahmood, 2014; Brokesova et al., 2014).[6] As discussed by Mahmood (2014), although very sophisticated models have studied price discrimination under different market hypotheses, yet a clear understanding of the mechanisms driving the pricing strategies in real world settings, is still deficient. This paper investigates personalized promotion strategies by studying the welfare implications of behavior-based promotion discrimination in e-commerce.[7]

The main contribution of this work is to develop an empirical framework to assess the welfare effects of online price discrimination with promotional coupons. Such framework illustrates how e-commerce platforms leverage users' information to personalize promotional content, and can explain why some targeting algorithms fail to improve revenues. However, quantifying the welfare consequences of behavior-based promotion discrimination requires information about the company's personalization strategy and the consumers' demand, which are both hardly disclosed in the digital retail business. Even with access to observational data, the endogeneity of firms' management guidelines, marketing decisions and privacy policy, poses a serious obstacle for the estimation of demand and supply functions. Using large-scale randomized controlled trials (RCTs) to evaluate BBPD effectiveness can overcome these challenges (e.g., Gordon et al., 2019; Gosnell et al., 2020). This research developed thanks to the collaboration with so-called partner *Alpha*, a tech company that sells customers' segmentation technology to digital businesses.[8] The partner carried out a testing phase on a beauty products e-commerce site via RCT. Specifically, it randomized two different classes of coupon treatments: unconditional discounts and minimum purchase level discounts. Experimental treatments result salient to e-commerce buyers. Randomizing coupons in the field experiment helps to address the endogeneity problem, but poses two main challenges. The first is that the heterogeneity of demand is not directly observable from the data. The second, more significant, is that the company assigns the discounts randomly, without implementing any promotion discrimination algorithm that maps user characteristics into promotional discounts. These challenges motive the development of a framework for demand and supply. This paper estimates a structural demand model

[5] Esteves (2009) surveys the economics of behavior-based price discrimination, while Fudenberg and Villas-Boas (2012) offers a general discussion of theoretical models and Belleflamme and Peitz (2015) provides textbook presentation of BBPD. Theoretical works date back to Villas-Boas (1999) and Fudenberg and Tirole (2000), and address the economic implications of price discrimination on firms' profit and welfare effects. Theoretical results on personalized pricing based on consumers' history suggest that the dominant strategy of competitive firms is to offer discounts with the twofold purpose of poaching competitors' customers and extracting profits from loyal clients.

[6] A canonical example is Rossi et al. (1996) which presents an novel empirical method using panel dataset of purchases to investigate the value of consumer-level information. The paper estimates a Bayesian hierarchical choice model to derive optimal targeted coupons. By empirically analyzing the net gain in revenues with different information sets, the potential of purchase history for targeted discounts is analyzed. Instead, Mahmood (2014) and Brokesova et al. (2014) provide an experimental approach with a focus on loyalty rewards' profitability. These papers represent the first studies of behavioral price discrimination adopting controlled laboratory experiments.

[7] Caillaud and De Nijs (2014) defines behavior-based price discrimination as "*the practice of offering different prices to different customers according to their past purchase history*". In the related literature BBPD has also been referred to as customer relationship management based pricing (e.g., Shin & Sudhir, 2010), pricing with customer recognition (e.g., Fudenberg & Villas-Boas, 2006), or one-to-one pricing (e.g., Rossi et al., 1996).

[8] *Alpha* is a fictitious name of the company in this research partnership. *Alpha* is a tech company based in Silicon Valley with offices in Europe, Asia, South America and Oceania. It was reported that *Alpha* generates an estimated annual revenue of $9M.

and adopts a pricing policy model to investigate the e-commerce's optimal promotion strategy. Both models take advantage of the exogenous variation in pricing treatments induced by RCT to estimate the demand primitives and the causal effects of treatments on demand's margins. By simulating different personalization algorithms, the counterfactual analysis shows that uniform pricing results sub-optimal. In comparison, behavioral-based personalized promotions increase welfare on both sides of the market as elastic users benefit from price discounts, while targeting coupons stimulate demand and, consequently, profits.

The estimated demand model uses a discrete-continuous choice framework (Lee & Allenby, 2014). Users enter the e-commerce platform, randomly receive a discount treatment, and choose how much to spend, or not to purchase at all. The demand model features utility-maximizing consumers, who allocate their exogenous endowment between the e-commerce aggregate purchase and the outside option (Kim et al., 2023). Demand heterogeneity is accounted by the addition of an idiosyncratic shock in a quasi-linear utility functional form. The identification strategy leverages concavity assumptions to define a sufficient condition of optimality, which yields a lower and upper bound on the support of the error distribution for the empirical demand. Since the support thresholds are nonlinear in the demand parameters, variation in the distribution of e-commerce sales and randomized treatments identifies preferences. The estimation of demand primitives reveals that discounts are salient to customers (e.g., Esteves & Reggiani, 2014; Hitsch et al., 2021). Indeed, the model shows that consumers not receiving the coupon, are more responsive to price changes of e-commerce goods. Furthermore, demand elasticity increases in the magnitude of discounts, and depends on the interaction between discount percentages and minimum purchase levels.

On the supply side, a pricing policy model formalizes how the monopolistic policy-maker develops its optimal pricing strategy (Miller & Hosanagar, 2020). The problem of the e-commerce firm is to select, conditional on individual-specific information, the combination of discount and minimum purchase threshold to be assigned to each customers in order to maximize expected profit. Pricing treatments differ in their causal impact on both the probability of conversion (extensive margin) and the expected spending (intensive margin). A treatment is defined optimal if its causal effect on expected revenue dominates the treatment's expected cost. Among different optimal treatments, the policy-maker assigns to each user the treatment which maximizes the product between the extensive and intensive margin, weighted by the treatment. For a finite set of treatments, this assignment criterion produces a optimal treatment policy, which simulates e-commerce's personalized promotion algorithm. Analytically, heterogeneous average treatment effects on conversion rate and online spending are estimated via supervised machine learning. Counterfactual analysis introduces alternative pricing policies, whose assignment rules differ from the optimal treatment policy's. Policy evaluation compares counterfactual policies through a consistent estimator of profit (e.g., Hitsch et al., 2023; Smith et al., 2023). Since it accounts for the causal effect on both demand margins, the optimal treatment algorithm is the most effective policy, with a 39% increase in the expected profit, compared to the uniform pricing benchmark. Information about users' past purchase history is relevant to enhance the effectiveness of discount targeting algorithms. Behavioral analysis decomposes demand between new users, individuals who have never visited the e-commerce before, and returning users, who have already browsed the platform (Chu et al., 2019). Returning users are more profitable than new users: the optimal pricing policy increases the expected profit of returning

users by 43%, whereas new users exhibit a 29% increase. Nevertheless, new users' demand appears to be elastic both at the extensive and intensive margin, while returning users are very inelastic at the intensive margin. Such heterogeneity between the two demand segments can be exploited by the policy-maker. By targeting only the extensive margin, behavior-based personalized policy generates an additional 11% profit gain compared to the aggregate optimal counterfactual for the returning users' market share (e.g., Shiller, 2020). Thus, if e-commerce had monopolistic access to information about users' past purchases data, it would have an incentive to discriminate its pricing policy based on the customers' history.

In the welfare analysis, the demand of e-commerce customers is combined with the pricing policy of the digital retailer to quantify both the producer and consumer surplus. Simulated demands under counterfactual targeting policies are associated to different distributions of profit and indirect utility. Relative to the uniform policy, the optimal treatment policy increases e-commerce profit by 24%. The algorithm accounts for the demand's heterogeneity by assigning to different users the treatment with the greatest causal impact on both the extensive and intensive margin. Consumers also benefit from targeting price discounts as, at the aggregate level, consumer surplus rises by 4%. Under the uniform pricing benchmark no customer receives any discount, while under the optimal treatment policy the users who are responsive to coupons, they receive the treatment and access the e-commerce goods' basket at a discounted price, contributing to the overall increase in the consumer surplus (Gehrig et al., 2012; Carroni, 2018; Esteves & Shuai, 2023). If the retailer could observe the users' purchasing history, it would be profit-increasing for the monopolist to integrate this information in its personalized promotion algorithm. Unlike the aggregate promotion discrimination, which results from the optimal treatment policy, behavior-based promotion discrimination involves adopting different treatment policies for different segments based on past purchase history. First-time individuals would still be assigned to the optimal treatment policy, which remains the profit-maximizing policy and generates a 18% increase in producer surplus. Conversely, returning users would be allocated to the conversion-targeted policy, which further drives an increase in e-commerce profitability by 11% (+28%), relative to the optimal treatment policy (+25%).

Since the consumer surplus of returning users increases more (+4.9%) than new users' (+3.7%), loyal customers benefit more from the optimal personalized promotion policy. Furthermore, behavior-based promotion discrimination generates additional profits at no cost to this segment. At the aggregate level, the welfare analysis shows that switching from optimal treatment policy to conversion-targeted policy neither decreases nor increases the returning consumers' welfare. By complementing personalized discounts with users' history information, behavioral discount discrimination increases total welfare, particularly for users with past purchase experience (e.g., Colombo, 2018). In this sense, BBPD effectively serves as a market segmentation mechanism: the e-commerce takes advantage of heterogeneity in demand elasticity revealed by past purchase information, to assign behavior-specific pricing algorithms to different consumers' segments. From a theoretical perspective, behavior-based promotion discrimination resembles third-degree price discrimination (e.g., Cowan, 2012; Bergemann et al., 2022). However, different consumers' groups are not charged different prices, but rather are assigned to different targeting algorithms. Within the same segment and for the same, corresponding pricing policy, two users with different covariates might receive different coupon treatments. This mechanism is thus closer to imperfect

first-degree price discrimination, which, at least theoretically, converges to perfect price discrimination as the set of covariates increases and the treatment grid expands (e.g., Shiller et al., 2013).[9]

This paper contributes to the development of a recent empirical literature focusing on targeted promotions in e-commerce retail. It expands the research on personalized marketing by investigating how past purchase records affect the pricing strategy of digital sellers with monopolistic access to users' information.[10] Recent empirical works on personalization have developed new quantitative research methods through advances in machine learning (ML). Supervised ML has contributed to a growing literature on targeting implementation and pricing policy evaluation. This study contributes to an emerging marketing literature which combines ML techniques with randomized field experiments. Simester et al. (2020) clarifies the advantage of using "randomized-by-action" instead of "randomized-by-policy" experiments. Since randomized-by-action design allows the evaluation of new policies without additional experiments, an off-policy approach is preferable for evaluating alternative policies that are not observable in the data. Random assignment of policy treatments combined with off-policy profit evaluation have also been employed in latest publications (e.g., Yoganarasimhan et al., 2020; Liu, 2023; Hitsch et al., 2023; Dubé & Misra, 2023).

Notably, Dubé and Misra (2023) is the first paper to investigate the welfare implications of personalized prices through scalable machine learning tools. Using data from a pricing experiment, they estimate the demand model and build a targeted pricing policy to quantify expected profits with personalized prices, relative to a uniform pricing policy. This paper shares with Dubé and Misra (2023) the purpose of welfare analysis, although there are some considerable differences. This paper focuses on personalized promotions in the context of e-commerce retail business, instead Dubé and Misra (2023) investigates the personalization of prices for a online job search service platform. This first, main difference results in the type of interventions in the field experiments. In Dubé and Misra (2023) different experimental prices are randomly assigned in their first RCT, while this research takes the prices of thousands e-commerce products as given, and leverages the exogenous variation of the promotional components of coupons, namely discounts and minimum purchase thresholds. Second, since only new, prospective customers are included in their samples, Dubé and Misra (2023) studies the welfare implications of personalized pricing, but does not explore how consumer purchase information affects targeting strategies. This paper contributes by specifically studying behavior-based promotion discrimination. Third, while controlled experiments in Dubé and Misra (2023) randomize the monthly fee for an employer-employee matching online platform, this paper leverages variation in percentage discounts of e-commerce check-outs, providing a more general framework for digital retailers. Furthermore, the focus on business-to-client transactions corroborates the external validity of this work, compared to Dubé and Misra (2023) which explore business-to-business pricing decisions. Lastly, both papers measure consumer welfare by simu-

[9] As argued in Dubé and Misra (2023), *"Statistical uncertainty typically limits the segmentation to an imperfect form of targetability. The approximation is also typically closer under unit demand since personalization typically cannot target a different price to each inframarginal unit purchased by a consumer"*.

[10] A summary by Arora et al. (2008) presents the main research challenges and the literature gaps related to both firms and customers behavior in personalization contexts. More recently, through a comprehensive review of personalized marketing literature, Chandra et al. (2022) defines personalization as *"a strategy to gain a competitive advantage, encompassing learning, matching, and delivering products and services to customers"*. The authors argue that personalization per se aims at improving customer satisfaction, thereby inducing loyalty. In this sense, personalized marketing can improve customer welfare.

lating pricing policies using machine learning techniques. However, Dubé and Misra (2023) estimates a logit demand model using a weighted likelihood bootstrap algorithm and then calibrates the personalized pricing policy. Instead, this paper first develops a structural model of consumer demand which accounts for baskets of goods, as typical of e-commerce transactions, rather than single purchases. Pricing policies are simulated through ML and policy evaluation adopts a consistent estimator of expected profit. At the current stage of research, this paper is the first to explore behavior-based personalized promotions combining field experimental data, theoretical models and ML estimation methods. The results have important implications for data privacy policy and the regulation of retail digital market.

This paper is structured as follows. Section 2 presents the partner company and data. Section 3 outlines the field experiment and discusses the empirical analysis. Sections 4 and 5 develop the demand model and the pricing policy model, respectively. Section 6 implements the welfare analysis and examines the welfare consequences of behavior-based promotion discrimination. Section 7 concludes the paper.

2 Partner Company and Data

Alpha is an American tech company based in Silicon Valley. It supports digital retailers by implementing its market segmentation technology to improve the performance of its customers. Its patented software analyzes real-time consumers' behavior to deliver personalized contents with the goal of optimizing targeted business outcomes (Figure D.2). Business of *Alpha*'s customers varies considerably, ranging from electronics companies and car manufacturers to fashion websites. This paper focuses on the case study between *Alpha* and a cosmetics brand. *Alpha* ran several randomized controlled trials on the brand's e-commerce site, and the access to the dataset of RCTs was provided under confidentiality partnership.[11]

Field experiment. In the experiment each session is randomly assigned to the control or the treatment group. The RCT treatment t is a promotional pop-up that appears on the screen of the user's device. There are two components for each treatment t: a percentage discount d_t on the e-commerce gross order and a minimum purchase threshold MPT_t. The customer benefits the discount on the e-commerce order only if the amount spent is at least equal to the minimum threshold. Combinations of d_t and MPT_t generate a multi-armed experimental setting, summarized by a control group and two treatment groups, namely control C if $d_C = 0$ and $MPT_C = 0$ (Figure D.3a), semi-treatment ST if $d_{ST} > 0$ and $MPT_{ST} = 0$ (Figure D.3b) and full-treatment FT if $d_{FT} > 0$ and $MPT_{FT} > 0$ (Figure D.3c). Control approximates uniform pricing, while semi-treatment and full-treatment resemble linear and nonlinear pricing schemes. The discounts d_t and the purchase levels MPT_t are orthogonal and discrete. The treatment grid consists of a 5×5 matrix such that the percentage discount $d_t \in \{0, 5, 10, 15, 20\}$ and the minimum purchase threshold $MPT_t \in \{0, 15, 25, 35, 45\}$. Propensity scores of control, semi-treatment and full-treatment are 15%, 50% and 35%, respectively.

Collaborators at *Alpha* disclosed relevant information regarding the RCT experimental features. First, the size of d_t depends on the industry in which the *Alpha*'s client operates. For beauty products, whose

[11] Reciprocal collaboration was signed by means of mutual nondisclosure agreement (MNDA), which considers confidential information all the information regarding *Alpha*'s identity, products, services, customers, business, finances, and IT infrastructure. In the MNDA, purposes to use *Alpha* confidential information are stated as potential mutual business opportunities.

price varies between €5 and €40, the discount percentage was capped at 20% to contain the cost of the experiment. Second, the magnitude of MPT_t is associated with the distribution of the cart values when the customer reaches the cart page for the first time, and also relates with the final gross order q_i. Once the treatment is received, it is possible, but unlikely, that the minimum purchase is lower than the cart value. This information is important to assure that the demand constraint imposed by MPT_t is actually binding as the user learns about it. Moreover, the probabilities of assignment to each treatment group are specified ex-ante the experiment, which guarantees experimental control for weighting estimators in the context of causal inference. To control for potential learning opportunities and strategic behavior of returning users, RCT sampling imposes that participants are randomly assigned to an experimental group only once. Accounting for the session's IP address, if a user has already received a treatment and visits again the e-commerce, these further sessions are automatically disregarded in the collection of the field experiment's sample.

Conversion funnel. The shopping experience that each user takes by browsing the e-commerce site and ultimately by converting the session into a purchase, can be divided into several steps. With a conversion funnel of three stages, the shopping session is the unit of analysis.[12] Each session begins as user i logs-in the e-commerce. No registration is required to access the website platform. The identification of different users relies on the IP address of any device connected to the e-commerce computer network. As the session begins, treatment t is randomly assigned to user i. By the time user i browses to the cart page for the first time, both discount d_t and purchase threshold MPT_t are received in the form of pop-up of promotional coupon. The discount is automatically added to the order summary page. Once attained the check-out page, user i pays $q_i(1 - d_t)$ if the gross order q_i is equal to or greater than MPT_t, otherwise user i gets no discount. Screenshots of the conversion funnel from a representative session are presented in the Online Appendix (Figure D.4).

Data. The representative sample counts approximately 5 million sessions that occurred over 13 months, between January 1st, 2022 and January 31st, 2023. Since the RCT was conducted online, shopping sessions took place in many countries with an 80% predominance of European countries such as France, Italy, Spain, Portugal, United Kingdom, Germany and the Netherlands. For each observation, variables about characteristics of user, treatment, session, history, technology, time and location are available. Data analysis focuses on the two margins of consumers' demand, namely the probability of purchase (extensive margin) and the gross spending (intensive margin). On average, the attainment rate of cart page is 30% and conversion rate is 6%. Conditional on converting the session, the average gross order is €45 with an average of 5 products and 29 minutes per session. Returning users are individuals with at least one previous visit before the start of the field experiment. They count for 26% of the RCT sample, while new users are 74%. Summary statistics between treatment groups and control groups are balanced, as shown in Table B.1. Section B presents the dataset in detail.

[12] Marketing literature (e.g., Seiler & Yao, 2017) usually divides conversion funnel into four levels: awareness, interest, desire, and action. Given the descriptive purpose of this framework and the dataset characteristics, the three steps leading to purchase are properly described by awareness, interest and action. In this setting, user i attains the first stage once the cart page is visited for the first time, Then, the second stage is set as the user i selects a gross order q_i above the minimum purchase threshold MPT_t. At last, by converting the session, user i reaches the third level.

3 Field Experiment

In this context of digital platforms, behavior-based promotion discrimination refers to the ability of the firm to offer different promotional contents to consumers with different past purchase records (Fudenberg & Villas-Boas, 2006). Because of its information advantage about customers, the e-commerce is capable of setting different promotions to increase its profit. The implicit rationale lies in the assumption that users with different purchase histories actually exhibit diverse demand elasticities. This chapter investigates such hypothesis through descriptive evidence (Section 3.1) and casual inference (Section 3.2).

3.1 Descriptive Evidence

In the representative sample there are roughly three new users (74%) for every returning user (26%). However, at the end of the conversion funnel, this ratio becomes one to one: among the customers who converted the session, 54% are new users and 46% are returning users. New users spend on average €42 per session, compared to €49 average spending of loyal customers. In relative terms, the returning users appear to be more likely to purchase, and to spend more than most of the new users (Figure 1).

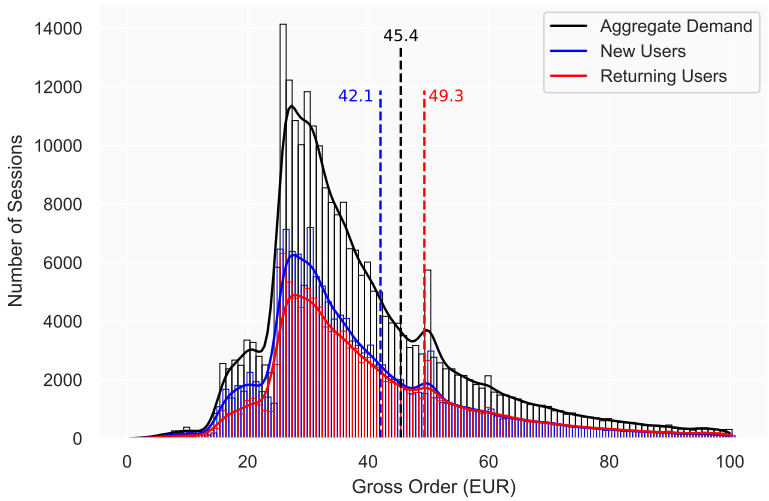


Figure 1: Gross order distribution

Note: The histogram displays the number of converted sessions by gross order. Demand counts 271,918 converted sessions. The plot includes only sessions which were converted with a gross order below €100. Black histogram depicts the aggregate demand, which has 261,157 sessions with €45.4 mean. Of the aggregate demand, 46% are returning users (red plot, €49.3 mean), while 54% are new users (blue plot, €42.1 mean). The plotted densities are computed using a Gaussian kernel smoother.

The following regression tests the association between purchase histories and the conversion probability:

$$Y_i = \gamma_0 + \mathbf{x}_i\gamma + \gamma_U U_i + \gamma_H H_i + \epsilon_i, \tag{3.1}$$

where i indexes the individual-specific session, U_i is the new user dummy, H_i is the purchase history dummy and Y_i is the conversion dummy.[13] Table C.1 displays the estimated coefficients for control, semi-treatment and full-treatment groups. Each column includes additional controls in x_i, which account for

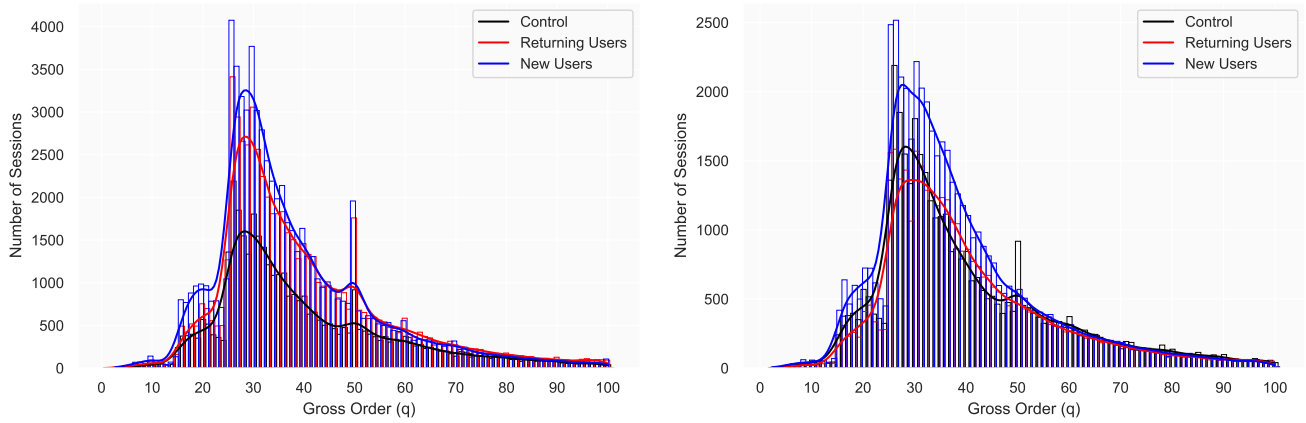
[13] $U = 1$ if the user has no history, namely it is the first session on the e-commerce website. While, for the returning user $U = 0$, $H = 0$ if the user has never purchased before, otherwise $H = 1$ if the user has converted the session at least one.

the characteristics of user, session, history, technology, time and location. For returning users with no past purchases, the conversion likelihood is around 10%. Relative to this group, the estimates indicate that costumers with purchase history exhibit between 60% and 92% more odds of buying. Conversely, new users are associated with a reduction in the purchase probability between 10% and 13%. Statistically significant correlation highlights demand heterogeneity among these groups at the extensive margin.

Is heterogeneity of purchase histories also correlated with demand intensive margin? This linear specification examines whether purchase decisions are associated, on average, with higher spending q_i:

$$\log(q_i) = \delta_0 + \mathbf{x}_i\delta + \delta_U U_i + \delta_H H_i + \epsilon_i, \tag{3.2}$$

where x_i is as in (3.1). Coefficients in Table C.2 should be interpreted in terms of percentage variation of gross order. For both control and treatment, how much is spent on the e-commerce is correlated to the demand segments defined between new and returning users. In particular, new users are associated with a decrease of spending between €2.2 and €3.1 (3-5% of the aggregate average) compared to returning customers. Figure 2 compares the spending distributions of new and loyal users, for both semi- and full-treatment. On the one hand, new and returning users show a similar, higher spending pattern relative to the control group, when they receive unconstrained discounts (Figure 2a). On the other hand, new customers show a propensity to spend greater than loyal clients in the case of conditional discounts (Figure 2b). On the whole, preliminary descriptive evidences suggest that new and returning customers seem to differ both at the extensive and intensive margin.



(a) Semi-Treatment

(b) Full-Treatment

Figure 2: Gross order distribution per treatment

Note: The panels show the distributions of converted sessions for semi-treatment and full-treatment, for returning and new users. Sessions with gross order greater than €100 are not depicted. Black histogram shows the control group demand (39,552 sessions with €40.00 mean and €17.36 st.dev.). (a) Semi-treatment counts 134,274 sessions: 47% of returning users (red plot, €39.88 mean) and 53% of new users (blue plot, €37.75 mean). (b) Full-treatment has 87,331 sessions: 43% of loyal users (red plot, €40.40 mean) and 57% of new users (blue plot, €38.02 mean). The plotted densities are computed using a Gaussian kernel smoother. Figures D.5 and D.6 provide further graphical intuition.

3.2 Causal Inference

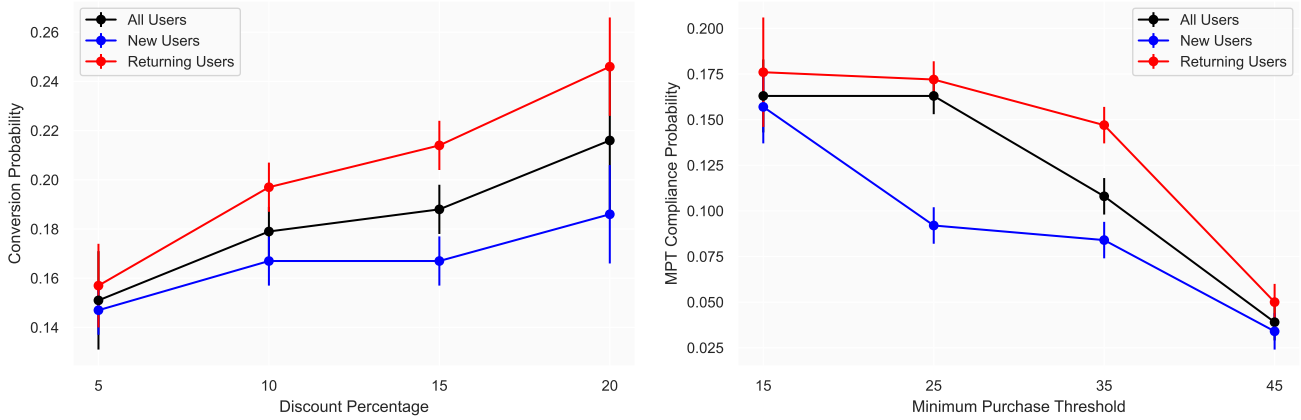
Randomized controlled trials (RCTs) are experiments where the assignment mechanism is independent to individual characteristics, and the experimenter controls for the propensity score, namely the probability of being assigned to the treatment. By the time the users in the treatment group browse to the cart page for the first time during the session, they receive a promotional pop-up that appears on their device screen. The treatment varies according to two terms: the percentage discount d_t on e-commerce spending and the minimum purchase threshold MPT_t. The combination of different levels of d_t and MPT_t generates a multi-armed experimental setting for promotional contents such as discounts and conditional discounts, or coupons. Since RCTs have the potential to establish a cause-effect relationship between the treatment and the targeted outcome, this analysis aims at assessing the impact of promotional treatments on both the extensive and intensive margins of demand.[14]

ATE. The average treatment effect (ATE) measures the average causal effects of multi-armed treatments on both the conversion probability and the compliance probability, and are estimated as follows:

$$Y_i = \alpha_0 + \alpha_{ATE} T_i + \epsilon_i, \tag{3.3}$$

$$C_i = \beta_0 + \beta_{ATE} T_i + \epsilon_i, \tag{3.4}$$

where Y_i is the conversion dummy, C_i is the MPT_t compliance dummy (i.e., $C_i = \mathbb{1}_{\{q_i \geq MPT_t\}}$) and T_i are the treatment dummies for the percentage discounts and the minimum purchase thresholds. Figures 3a and 3b represent the ATE estimates of Equations 3.3 and 3.4 respectively, for all users, and for new and returning customers, at the discrete levels of d_t and MPT_t.



(a) ATE on Conversion Probability

(b) ATE on Compliance Probability

Figure 3: Average treatment effect (ATE) by discount d_t and minimum purchase threshold MPT_t

Note: Panel (a) reports the estimates of ATE with standard deviation for conversion probability. Panel (b) displays the estimates of ATE with standard deviation for compliance probability. The black line refers to aggregate demand, namely new and returning users together, while the blue line refers to new users and the red line to returning users.

On the one hand, Figures 3a compares the levels of discount percentage d_t with the probability of purchase. The ATE estimates exhibit an upward pattern, which implies that higher discounts increase the

[14]For a comprehensive review about field experiment and randomized controlled trials refer to Athey and Imbens (2017). For a presentation about program evaluation, ITT, ATE and LATE see papers like Angrist and Imbens (1995) or Rubin (2005).

conversion rate. Conditional on purchase history, loyal users result more responsive to discounts, with a conversion probability between 16 and 26%. The greater the discount d_t, the greater the difference of ATEs between new and returning clients. On the other hand, Figures 3b relates the minimum spending level of coupons MPT_t with the compliance rate, namely the probability that the user's aggregate spending is greater than MPT_t. Typically coupon discounts involve a compliance cost that needs to be borne in order to benefit the price reduction. At the aggregate level, the greater the demand constraint MPT_t to access the discount, the lower the compliance rate. Moreover, returning users and new users exhibit different patterns. Since new users spend less on average, they are also less likely that their demand q_i will exceed the minimum purchase threshold, which acts as a constraint on the demand's intensive margin.

ITT. The potential of RCTs as causal inference instrument relies on the randomization of the treatments' assignment. The differences in the probability of conversion between groups can be causally attributed, on average, to the assignment of treatments. The intent to treat effect (ITT) measures the average causal effect of assigning the treatment on the conversion probability, and is estimated in the specification:

$$Y_i = \beta_0 + \beta_{ITT} A_t + \epsilon_i, \tag{3.5}$$

where Y_i is the conversion dummy and A_t is the treatments' assignment dummy. In Table 1 the estimates suggest that the probability of converting the session is 5% for the control group, while being assigned to the semi-treatment (full-treatment) induces a 15% (5%) increase in the odds ratio of converting the session. Nevertheless, the coefficient β_{ITT} might not be consistently estimated if the implementation of RCT compromises the randomization of the treatments' assignment. The identification problem is related to the issue of self-selection by participants in the experiment, which represents a potential source of endogeneity that flawed RCTs are unable to control. Self-selection arises when unobserved characteristics of participants or their own behaviors influence the mechanism for assigning promotional treatments, resulting in some users not receiving the treatments they were assigned to. In the experimental setting, the fact that treatment is received by the time the cart page is visualized, might lead to a self-selection bias. It is indeed possible that some individuals assigned to the treatment group may not receive the treatment because they abandon the session before reaching the shopping cart page. The unaccounted heterogeneity associated with the behavior of participants during the shopping session could potentially flaw the exogenous variation induced by the randomization of treatments' assignment, making more difficult to determine the causal link between promotions and demands' margins.

LATE. By using the exogenous randomization of treatments' assignment in the RTC as an instrument for the endogeneity of self-selection, causal inference analysis estimates the local average treatment effects. The local average treatment effect (LATE) is the average causal effect of receiving the promotion for the participants who were assigned to the treatment group. The effectiveness of the treatment is evaluated by implementing a two-stage least squares (2SLS) estimation in which randomness of treatments' assignment serves as a valid instrument for the endogenous variation of users visiting the cart page for the first time. In the second stage, the likelihood of purchase is regressed on first stage fitted estimates according to the

following specifications:

$$\begin{aligned} C_i &= \alpha_0^I + \alpha_{IV}^I A_t + \xi_i, \\ Y_i &= \beta_0 + \beta_{LATE}\widehat{C}_i + \epsilon_i \end{aligned} \tag{3.6}$$

where Y_i is the conversion dummy, C_i is the cart attainment dummy which works as a proxy of compliance rate, $\widehat{C}_i$ is first stage fitted estimate, and A_t is the treatments' assignment dummy. Conditionally on attaining the cart page, the probability of users in the control group to convert the session is 5%. Results in Table 1 indicate that being assigned to the semi-treatment (full-treatment) and receiving the promotional content causes a 97% (53%) increase in the odds of converting the session compared to the control group.[15] On average, the users react to discounts by raising their conversion rate. Nonetheless, LATE estimates reveal the pervasiveness of self-selection among the experiment's participants. Due to self-selection, the semi-treatment (full-treatment) LATE indicates how the assigned participants visualizing the cart and receiving the treatment, are 10 (5) times more likely to convert the session compared to those in the control group. In conclusion, although some causal inference effects are identified, the problem of latent endogeneity hinders to consistently estimate the causal impact of discount promotions on demand.

Table 1: Estimates of ITT and LATE

	Combined Treatments		Semi-Treatment		Full-Treatment	
Likelihood of conversion	ITT	LATE	ITT	LATE	ITT	LATE
Intersection (β_0)	-2.962***	-2.291***	-2.962***	-2.215***	-2.962***	-2.497***
	(0.005)	(0.031)	(0.005)	(0.026)	(0.005)	(0.053)
ITT (β_{ITT})	0.105***[11]		0.142***[15]		0.050***[5]	
	(0.005)		(0.006)		(0.006)	
α_0^I		-1.103***		-1.103***		-1.103***
		(0.003)		(0.003)		(0.003)
α_{IV}^I		0.172***		0.210***		0.118***
		(0.003)		(0.003)		(0.003)
LATE (β_{LATE})		0.609***[84]		0.677***[97]		0.422***[53]
		(0.032)		(0.027)		(0.052)
Pseudo R-squared	0.001	0.001	0.001	0.002	0.002	0.002
Observations	5,087,685	5,087,685	3,331,334	3,331,334	2,597,593	2,597,593

Note: * – $p < 0.1$; ** – $p < 0.05$; *** – $p < 0.01$. The results of causal inference analysis for combined treatments, semi-treatment group and full-treatment group as described in (3.5) and (3.6). Standard errors in parenthesis. Odds ratio absolute value percentage variation in brackets.

The results of both the descriptive analysis and the causal inference support the hypothesis that an observable and considerable source of demand endogeneity is attributable to heterogeneity among groups with different purchasing histories. Overall, new users and returning users respond differently to RCT treatments. On average, users react to discounts by raising their conversion rate, and respond to coupons by reducing their compliance percentage. Individual characteristics linked to the past history of purchases reveal demand heterogeneity both at the extensive and intensive margin. The discounts impacting the demand extensive margin, are more salient to returning users, while the spending constraints induced by MPT_t which affect the intensive margin, result pivotal for new users. The following chapter

[15] In the Online Appendix, Table C.3 reports the estimates of ITT and LATE for new users and returning users.

introduces the demand model in which a representative user solves the constrained utility maximization problem with exogenous treatment. The model is then identified to estimate demand primitives and to derive demand elasticity to promotional treatments.

4 Demand: Elasticity to Exogenous Treatments

Offering promotional treatments would be profitable if, as the discount increases or as the minimum purchase constraint decreases, the conversion rate or the e-commerce spending, or both, increased. This chapter first explores this trade-off between d_t and MPT_t in stimulating demand by providing evidences of correlation between the treatments' components and the demand elasticity at the extensive and intensive margins (Section 4.1). The analysis proceeds by developing a structural demand model (Section 4.2), presenting the identification argument and discussing the identification conditions (Section 4.3), and then estimating the demand's parameters and price elasticity (Section 4.4).

4.1 Reduced-form Analysis

Extensive margin. Does increasing the discount (or decreasing the coupon constraint) lead to an increase in the conversion rate? The following logit regression examines the correlation between the discount percentage d_t and the minimum purchase threshold MPT_t, with the likelihood of conversion:

$$Y_i = \gamma_0 + \mathbf{x}_i\gamma + \gamma_d d_t + \gamma_{MPT} MPT_t + \epsilon_i, \tag{4.1}$$

where i and t index the individual-specific session and the associated treatment respectively, Y_i is the conversion dummy, and x_i is as in (3.1). Since the exogenous treatment consists of either a discount (semi-treatment) or a conditional discount (full-treatment), offering an higher discount percentage does not necessarily imply a greater marginal impact on the conversion rate, as the consumer's demand could be overall inelastic, or it could be responsive to a decrease in the minimum threshold instead. The estimates of Table C.4 show that, on average, a unit increase in the discount is positively associated with the conversion probability for semi-treatment (2%) and full-treatment (3%). Results also indicate that there is no correlation between conversion rate and the magnitude of minimum purchase thresholds.

Intensive margin. The following linear specification investigates the correlation between e-commerce spending, measured by gross order q_i, and both the discount d_t and the purchase threshold MPT_t:

$$\log(q_i) = \delta_0 + \mathbf{x}_i\delta + \delta_d d_t + \delta_{MPT} MPT_t + \epsilon_i, \tag{4.2}$$

where x_i is as in (3.1), and coefficients are interpreted in terms of expenditure percentage variation. According to the estimates in Table C.5, there are no statistically significant results that support the hypothesis that either d_t ot MPT_t are correlated to gross order q_i.

4.2 Demand Model

This section extends the general framework developed by Lee and Allenby (2014) in order to accommodate the dataset of the field experiment run by the partner company on one of its clients' e-commerce website.

The demand model considers the purchase decision of a representative user i with the goal of estimating the primitives of the demand of e-commerce goods that maximize their utility (Dubé, 2019).

The combination (d_t, MPT_t) specifies the randomized treatment of the RCTs. Both the econometrician and user i observe the exogenous treatment (d_t, MPT_t). In their choice problem, user i considers good low l and good high h, together with an outside good x which is the numeraire of the economy. Good l is a bundle of e-commerce products defined by the L-vector products $\sum_{l=1}^{L} p_l q_l$ such that the demand of l q^l is lower than MPT_t ($\sum_{l=1}^{L} p_l q_l < MPT_t$). Good h is a bundle of e-commerce items defined by the H-vector products $\sum_{h=1}^{H} p_h q_h$ such that the demand of h q^h is at least equal to MPT_t ($\sum_{h=1}^{H} p_h q_h \geq MPT_t$). The unit price of good l is normalized to 1, the unit price of good h is $(1 - d_t)$. Preferences are represented by a quasi-linear functional form of direct utility.[16] This assumption guarantees that the utility function is concave for both goods' demand q^l and q^h defined over a non-negative domain.[17] In models of consumer choice, two further assumptions are introduced: linearity and additive separability. The utility function $\mathcal{U}$ is linear in the outside good x, and is additively separable for the interior goods l and h. Linearity ensures that the value of the outside good x does not satiate, while additive separability simplifies the integration of the likelihood function. The quasi-linear additively separable direct utility is

$$\mathcal{U}(q_i^l, q_i^h, x_i^0) = \frac{\alpha_l e^{\epsilon_l}}{\gamma_l} \ln(\gamma_l q_i^l + 1) + \frac{\alpha_h e^{\epsilon_h}}{\gamma_h} \ln(\gamma_h q_i^h + 1) + \alpha_0 x_i^0, \tag{4.3}$$

where q_i^l, q_i^h and x_i^0 represent the purchased quantities of good l, h and outside good of user i.[18] Parameters α_l and α_h correspond to the baseline marginal utility, or the marginal utility for no consumption of goods l and h, respectively. Therefore, a higher marginal utility $\alpha_{l(h)}$ implies less likelihood of a corner solution for interior good l (h). The baseline marginal utility of the outside good is captured by α_0. Additionally, parameters γ_l and γ_h represent a satiation parameter. Flexibility in the degree of satiation allows for greater satiation effect in the consumption of l (h) associated to higher values of $\gamma_{l(h)}$. Error terms ϵ_l and ϵ_h indicate a latent idiosyncratic shock which are known to user i and unobserved by the econometrician. Good-specific heterogeneity is introduced through normally distributed errors, that is $\epsilon_l \sim \mathcal{N}(\mu_l, \sigma_l^2)$ and $\epsilon_h \sim \mathcal{N}(\mu_h, \sigma_h^2)$. Errors ϵ_l and ϵ_h are assumed independent and identically distributed.

User i solves their utility maximization problem subject to a set of constraints. The first corresponds to the non-negative constraints imposed on purchased quantities. The demand of outside good x_i^0 is assumed strictly positive, implying that the indirect utility of user i is non-negative at the corner solution. The demand of good l (h) is assumed non-negative to accommodate for corner solutions in case of no purchase by user i:

$$x_i^0 > 0, q_i^l \geq 0, q_i^h \geq 0. \tag{4.4}$$

[16] Quasi-linear utility functions are linear in one argument, generally the numeraire. Quasi-linear preferences can be represented by the utility function $\mathcal{U}(x_1, x_2, \ldots, x_n) = x_1 + \theta(x_2, \ldots, x_n)$ where x_1 is the numeraire commodity and θ is strictly concave. In (4.3) θ is logarithmic, which is increasing and concave.

[17] In the model marginal utility is thus continuous, positive and decreasing both in good l and in good h, i.e., $\frac{\partial \mathcal{U}}{\partial q_i^{l(h)}} = \frac{\alpha_{l(h)} e^{\epsilon_{l(h)}}}{\gamma_{l(h)} q_i^{l(h)} + 1} > 0$ and $\frac{\partial \mathcal{U}}{\partial q_i^{l(h)} \partial q_i^{l(h)}} = -\frac{\alpha_{l(h)} \gamma_{l(h)} e^{\epsilon_{l(h)}}}{(\gamma_{l(h)} q_i^{l(h)} + 1)^2} < 0$.

[18] The quasi-linear additively separable utility specification in (4.3) corresponds to a linear expenditure system form of the generalized constant elasticity of substitution (CES) utility function. Hanemann (1984) introduces the CES utility functional form $\mathcal{U}(x) = \sum_{k=1}^{K} (\frac{\gamma_k}{\alpha_k}) \psi_k \left((\frac{x_k}{\psi_k})^{\alpha_k} - 1 \right)$ which converges to (4.3) as α_k approaches zero. Bhat (2008) shows that the quasi-linear form closely approximates a subutility function profile based on a combination of γ_k and α_k values for a given ψ_k, with a subutility function based exclusively on γ_k or α_k values.

By definition of interior good low l and good high h, the treatment constraint reflects how the demand of good l *(h)* is defined with respect to the exogenous minimum purchase threshold MPT_t:

$$q_i^l < MPT_t, q_i^h \geq MPT_t. \tag{4.5}$$

Related to the demand constraint imposed by MPT_t, the interior goods l and h are perfect substitutes in case of no corner solution. A strictly positive demand q_i^l implies no demand of good h and vice versa:

$$q_i^l \cdot q_i^h = 0. \tag{4.6}$$

A further constraint comes from the assumption of discreteness of the interior goods' demand. Purchased quantities q_i^l and q_i^h belong to the set of natural numbers $\mathbb{N}_0$ capped at Q, where $Q = M_i - 1$. Such upper bound guarantees non-negative demand for the outside option:

$$q_i^l, q_i^h \in \{0, 1, 2, \ldots, Q\}. \tag{4.7}$$

The last constraint specifies the linear combination of purchased quantities of good l and h and the numeraire of user i weighted by their unit prices. The budget constraint imposes that the exogenous endowment M_i is allocated between the outside good, good l with unit price normalized to 1 and good h weighted by the treatment percentage discount $(1 - d_t)$[19]:

$$q_i^l + (1 - d_t)q_i^h + x_i^0 \leq M_i. \tag{4.8}$$

Ultimately, user i takes two decisions. First, they decide whether to convert the session or not. If they do not, then $q_i^l = 0$, $q_i^h = 0$ and trivially $x_i^0 = M_i$. Otherwise, if they buy, they select either $q_i^l < MPT_t$ or $q_i^h \geq MPT_t$. Then, $q_i^l > 0$ and $x_i^0 = M_i - q_i^l$ or $q_i^h > 0$ and $x_i^0 = M_i - (1 - d_t)q_i^h$ given treatment (d_t, MPT_t).

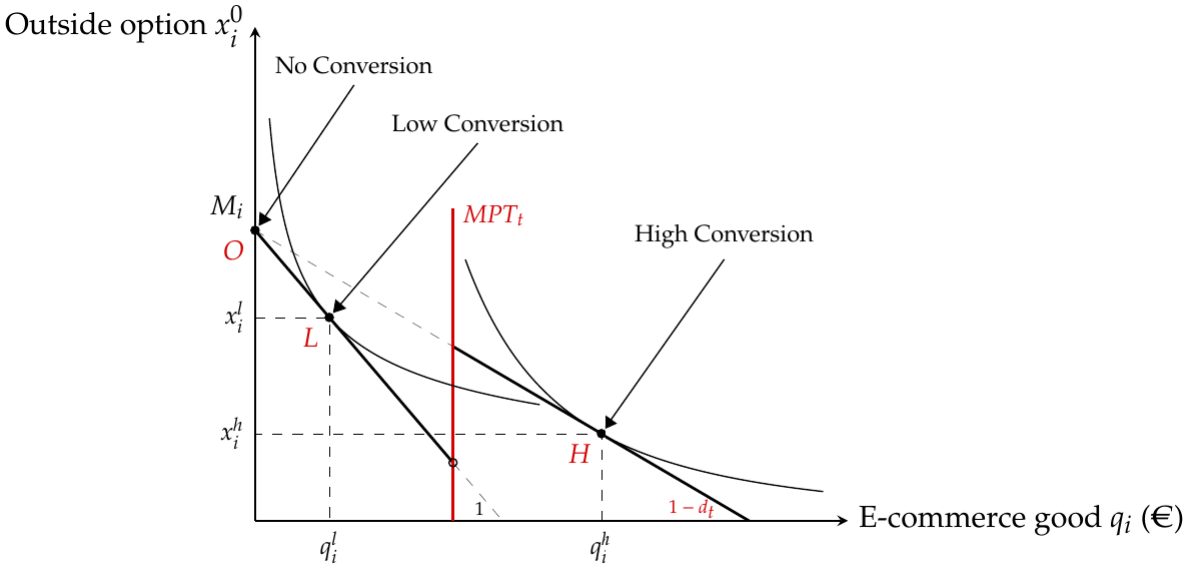


Figure 4: Graphical intuition of constrained utility maximization problem

Note: The figure represents the choice problem in (4.9). Full-treatment group has three demand equilibria: No Conversion, Low Conversion and High Conversion. The slope of budget constraint is 1 if $q_i < MPT_t$ and $1 - d_t$ otherwise.

[19]Because the Marshallian demand of $(x_2, \ldots, x_n)$ for quasi-linear utility functional forms does not depend on M_i, it is not subject to wealth effect, thus no further assumptions are imposed on the exogenous endowment M_i.

Figure 4 provides a visual representation of the demand model's equilibria. Specifically, three demand equilibria arise: No Conversion (also corner solution, i.e., $q_i^l = 0$, $q_i^h = 0$, $x_i^0 = M_i$), Low Conversion (i.e., $q_i^l > 0$, $q_i^h = 0$, $x_i^0 = M_i - q_i^l$) and High Conversion (i.e., $q_i^l = 0$, $q_i^h \geq MPT_t$, $x_i^0 = M_i - (1 - d_t)q_i^l$). In the Online Appendix, Figure D.7 graphically depicts the equilibria for control group, semi-treatment group and full-treatment group. The next section presents the utility maximization problem, derives the sufficient conditions of identification strategy, and estimates the demand's primitives and price elasticity.

4.2.1 Utility Maximization Problem

The utility maximization problem of user i is represented by a constrained optimization linear programming problem. The quasi-linear direct utility in (4.3) defines the objective function, while (4.4), (4.5), (4.6), (4.7) and (4.8) include both equality and inequality constraints. The purchase decision of user i is therefore formulated as

$$\begin{aligned} \max_{(q_i^l, q_i^h, x_i^0)} \; & \mathcal{U}(q_i^l, q_i^h, x_i^0) = \frac{\alpha_l e^{\epsilon_l}}{\gamma_l} \ln(\gamma_l q_i^l + 1) + \frac{\alpha_h e^{\epsilon_h}}{\gamma_h} \ln(\gamma_h q_i^h + 1) + \alpha_0 x_i^0 \\ s.t. \; & q_i^l + (1 - d_t) q_i^h + x_i^0 \leq M_i, \\ & q_i^l, q_i^h \in \{0, 1, 2, \ldots, Q\}, \;\; q_i^l < MPT_t, q_i^h \geq MPT_t, \;\; q_i^l \cdot q_i^h = 0, \\ & x_i^0 > 0. \end{aligned} \tag{4.9}$$

A demand equilibrium is a triplet $\{q_i^l, q_i^h, x_i^0\}^*$ which solves the maximization utility problem (4.9). It depends on an exogenous treatment (d_t, MPT_t), parameters' set $\{\alpha_0, \alpha_l, \gamma_l, \alpha_h, \gamma_h\}$ and the idiosyncratic shock realizations ϵ_l and ϵ_h. At any triplet $\{q_i^l, q_i^h, x_i^0\}^*$, the budget constraint in (4.8) is binding because the marginal utility α_0 of the outside good is strictly positive. By substitution method, the constrained optimization problem described in (4.9) can be equivalently formulated through a unconstrained linear programming problem as follows

$$\begin{aligned} \max_{(q_i^l, q_i^h)} \; & \mathcal{U}^*(q_i^l, q_i^h) = \frac{\alpha_l e^{\epsilon_l}}{\gamma_l} \ln(\gamma_l q_i^l + 1) + \frac{\alpha_h e^{\epsilon_h}}{\gamma_h} \ln(\gamma_h q_i^h + 1) + \alpha_0 \Big(M_i - q_i^l - (1 - d_t) q_i^h \Big) \\ s.t. \; & q_i^l, q_i^h \in \{0, 1, 2, \ldots\}, q_i^l < MPT_t, q_i^h \geq MPT_t, q_i^l q_i^h = 0, \\ & M_i - q_i^l - (1 - d_t) q_i^h \geq 0. \end{aligned} \tag{4.10}$$

Although their equivalence, the crucial difference between (4.9) and (4.10) depends on the degrees of freedom of linear programming. The objective function $\mathcal{U}$ in (4.9) is an unconstrained utility function of goods l and h and the outside good, while $\mathcal{U}^*$ in (4.10) is a constrained utility function of goods l and h satisfying the budget constraint. The objective function $\mathcal{U}^*$ preserves increasing satiation and negative concavity, which correspond to the diminishing marginal utility conditions:

$$\frac{\partial \mathcal{U}^*(q_i^l, q_i^h)}{\partial q_i^{l(h)}} = \frac{\alpha_{l(h)} e^{\epsilon_{l(h)}}}{\gamma_{l(h)} q_i^{l(h)} + 1} - \alpha_0 (1 - \mathbb{1}_{\{q_i^{l(h)} \geq MPT_t\}} d_t) > 0, \tag{4.11a}$$

$$\frac{\partial \mathcal{U}^*(q_i^l, q_i^h)}{\partial q_i^{l(h)} \partial q_i^{l(h)}} = -\frac{\alpha_{l(h)} \gamma_{l(h)} e^{\epsilon_{l(h)}}}{(\gamma_{l(h)} q_i^{l(h)} + 1)^2} < 0. \tag{4.11b}$$

Given the initial assumptions of additively separability of goods l and h and linearity of the outside good, both (4.11a) and (4.11b) depend only the demand of the interior good. This implies that the sufficient conditions to rationalize the observed demand equilibrium $\{q_i^l, q_i^h\}^*$ evaluate each realized purchase of interior goods independently. In addition, the binding mutual exclusivity constraint in (4.6) further simplifies the comparison condition and, consequently, the likelihood function is analytically inferable.

4.3 Identification and Likelihood Function

The identification argument leverages a random utility framework for rationalizing the demand predicted by the constrained utility maximization problem. The introduction of an idiosyncratic shock allows for demand heterogeneity which first order conditions would fail to satisfy. The variation in purchased quantities could potentially be driven by heterogeneity in consumers' preferences, unobservable characteristics of products or exogenous changes in market factors. The error terms included in (4.3) account for any unobserved covariates which would affect the demand of goods l and h. The identification strategy takes advantage of the concavity property of the utility function $\mathcal{U}^*$, which entails the existence of a global maximum at the realized demand $\{q_i^{l^*}, q_i^{h^*}\}$. Since the utility function $\mathcal{U}^*$ is characterized by a unique, global maximum within the proximity of the empirical distribution of purchased quantities, then negative convexity guarantees the condition of optimality by comparing the indirect utility at the realized demand with respect to the counterfactual indirect utility at neighboring discrete points' grid.

Thus, the sufficient condition for the observed demand $\{q_i^{l^*}, q_i^{h^*}\}$ to solve (4.10) derives from the property of utility function $\mathcal{U}^*$ to be concave at each quantity of any demand equilibrium. By considering a neighboring grid $\Delta \in \{-n, n\}$ of $\{q_i^{l^*}, q_i^{h^*}\}$ with $n \in \mathbb{N}_0$, the sufficient condition for optimality is

$$\begin{aligned} \mathcal{U}^*(q_i^{m^*}, q_i^{-m^*}) > \max\Big\{\mathcal{U}^*(q_i^{m^*}+\Delta, q_i^{-m^*}) \mid (q_i^{m^*}+\Delta, q_i^{-m^*}) \in F\Big\}_{\Delta\in\{-n,n\}} \quad \forall m \in \{l,h\} \\ \text{where} \quad F = \Big\{(q_i^m, q_i^{-m}) \mid M_i - q_i^m - (1-d_t)q_i^{-m} \geq 0, \\ q_i^m, q_i^{-m} \in \{0,1,2,\ldots,Q\}, q_i^m \cdot q_i^{-m} = 0\Big\}. \end{aligned} \tag{4.12}$$

Optimality conditions yield a lower bound $\underline{B}_i^m$ and an upper bound $\overline{B}_i^m$ on the support of the distribution of error ϵ_m where $m \in \{l, h\}$, which are consistent with the realized interior goods' demand $\{q_i^{l^*}, q_i^{h^*}\}$ within the neighboring $(-n, n)$ interval.

Proposition 1. *Demand equilibrium $\{q_i^{l^*}, q_i^{h^*}\}$ is defined in the interval $(\underline{B}_i^m, \overline{B}_i^m)$ of the support of error ϵ_m's distribution $\forall m \in \{l, h\}$ where:*

$$\underline{B}_i^m = \ln\left(\frac{\alpha_0(1-\mathbb{1}_{\{q_i^{m^*}\geq MPT_t\}}d_t)\gamma_m\Delta}{\alpha_m}\right) - \ln\left(\ln\left(\frac{\gamma_m q_i^{m^*}+1}{\gamma_m(q_i^{m^*}-\Delta)+1}\right)\right), \tag{4.13a}$$

$$\overline{B}_i^m = \ln\left(\frac{\alpha_0(1-\mathbb{1}_{\{q_i^{m^*}\geq MPT_t\}}d_t)\gamma_m\Delta}{\alpha_m}\right) - \ln\left(\ln\left(\frac{\gamma_m(q_i^{m^*}+\Delta)+1}{\gamma_m q_i^{m^*}+1}\right)\right). \tag{4.13b}$$

Proof. Refer to Section A.1 in the Online Appendix.

Parameter α_m captures the marginal utility for no consumption of good m. Its estimate depends on the share of corner solutions in the empirical demand distribution. The lower bound $\underline{B}_i^m$ goes to $-\infty$ for

$q_i^{m^*} = 0$, whereas for non-zero demand it is inversely proportional to α_m. This results in a small value of α_m for a higher proportion of corner solutions as the smaller α_m gets, the greater upper boundary $\overline{B}_i^m$ becomes in order to account for a larger share of zeros. Likewise, the satiation parameter γ_m implies a weaker preference, or lower satiation, for good m as γ_m decreases. For larger fractions of corner solutions, parameter γ_m reduces to accommodate for a greater upper boundary $\overline{B}_i^m$ in the error's distribution. The relation between the share of corner solutions in the empirical distribution and the marginal changes of preferences' parameters is corroborated by comparative statics.[20] The empirical distribution of purchased quantities contributes in the identification of γ_m. While variation in α_m would shift both $\underline{B}_i^m$ and $\overline{B}_i^m$ without modifying the error interval, different magnitudes in the demand entail different likelihoods given a positive estimate for the satiation parameter γ_m. Within the strictly positive demand distribution, analytical comparative statics show that both thresholds $\underline{B}_i^m$ and $\overline{B}_i^m$ depend on γ_m, but not on α_m.[21] However, satiation γ_m expands the likelihood interval of error ϵ_m in a not trivial manner, since the sign of $\frac{\partial \underline{B}_i^m}{\partial \gamma_m}$ and $\frac{\partial \overline{B}_i^m}{\partial \gamma_m}$ is not uniquely defined.[22]

Given the assumption of independent and identically normally distributed errors ϵ_l and ϵ_h, the likelihood of the empirical demand distribution is computed by integrating the joint density of both errors in the intervals specified in (4.13a) and (4.13b). For the demand parameters' space $\theta \in \Theta$ where $\Theta = \{\alpha_0, \alpha_l, \alpha_h, \gamma_l, \gamma_h, \mu_l, \sigma_l, \mu_h, \sigma_h\}$, the multinomial likelihood function $\mathcal{L}(\Theta)$ is defined as

$$\mathcal{L}(\Theta | q_i^{m^*}, q_i^{-m^*}, d_t, MPT_t) = \prod_{i=1}^{I} \int_{\underline{B}_i^m}^{\overline{B}_i^m} \left[\varphi\left(\frac{\epsilon_m - \mu_m}{\sigma_m} \right) d\epsilon_m \right]^{I_i^m} \cdot \int_{\underline{B}_i^{-m}}^{\overline{B}_i^{-m}} \left[\varphi\left(\frac{\epsilon_{-m} - \mu_{-m}}{\sigma_{-m}} \right) d\epsilon_{-m} \right]^{I_i^{-m}} \quad \forall m \in \{l, h\}, \tag{4.14}$$

where I_i^m is the indicator function $\mathbb{1}_{q_i^{m^*} \geq 0}$ incorporating (4.6). In case $q_i^{m^*} = 0$ the lower boundary $\underline{B}_i^m$ converges to $-\infty$. Similarly, the upper boundary $\overline{B}_i^m$ approaches $+\infty$ if the upper neighboring grid $q_i^{m^*} + \Delta$ violates the budget constraint, that is $q_i^{m^*} = Q$. In Section A.2 treatment-specific log-likelihood functions are derived.

4.4 Estimation of the Demand Model

Demand parameters Θ are estimated through maximum likelihood estimation (MLE). The goal of maximum likelihood estimation is to find the values of the demand primitives that maximize the log-likelihood function over the parameter space Θ. Therefore, treatment-specific MLE is

$$\hat{\theta}_t = \underset{\theta \in \Theta}{\operatorname{argmax}} \ \ell_t\left(\theta | q_i^l, q_i^h, d_t, MPT_t\right) \quad \forall t \in \{C, ST, FT\}, \tag{4.15}$$

where ℓ_t is the log-likelihood function for each combination of (d_t, MPT_t). In addition, the parameters of aggregate demand are estimated through the algorithm as follows

---

[20] $\frac{\partial \overline{B}_i^m}{\partial \alpha_m}|_{q_i^{m^*}=0} = -\frac{1}{\alpha_m} < 0$ and $\frac{\partial \overline{B}_i^m}{\partial \gamma_m}|_{q_i^{m^*}=0} = -\frac{\Delta}{\gamma_m \Delta + 1} \frac{1}{\ln\left(\gamma_m \Delta + 1\right)} < 0.$

[21] $\frac{\partial \underline{B}_i^m}{\partial q_i^{m^*}} = \gamma_m^2 \Delta \left(\frac{\gamma_m (q_i^{m^*} - \Delta) + 1}{\gamma_m q_i^{m^*} + 1} \right) \frac{1}{\ln\left(\frac{\gamma_m q_i^{m^*} + 1}{\gamma_m (q_i^{m^*} - \Delta) + 1} \right)} > 0$ and $\frac{\partial \overline{B}_i^m}{\partial q_i^{m^*}} = \gamma_m^2 \Delta \left(\frac{\gamma_m q_i^{m^*} + 1}{\gamma_m (q_i^{m^*} + \Delta) + 1} \right) \frac{1}{\ln\left(\frac{\gamma_m (q_i^{m^*} + \Delta) + 1}{\gamma_m q_i^{m^*} + 1} \right)} > 0.$

[22] $\frac{\partial \underline{B}_i^m}{\partial \gamma_m} = \frac{1}{\gamma_m} - \Delta \left(\frac{\gamma_m (q_i^{m^*} - \Delta) + 1}{\gamma_m q_i^{m^*} + 1} \right) \frac{1}{\ln\left(\frac{\gamma_m q_i^{m^*} + 1}{\gamma_m (q_i^{m^*} - \Delta) + 1} \right)}$ and $\frac{\partial \overline{B}_i^m}{\partial \gamma_m} = \frac{1}{\gamma_m} - \Delta \left(\frac{\gamma_m q_i^{m^*} + 1}{\gamma_m (q_i^{m^*} + \Delta) + 1} \right) \frac{1}{\ln\left(\frac{\gamma_m (q_i^{m^*} + \Delta) + 1}{\gamma_m q_i^{m^*} + 1} \right)}.$

$$\hat{\theta} = \underset{\theta \in \Theta}{\operatorname{argmax}} \sum_{t=1}^{T} \Pr\left(t = \bar{t}\right) \ell_t\left(\theta | q_i^l, q_i^h, d_t, MPT_t\right), \tag{4.16}$$

where $\bar{t} \in \mathcal{T} = \{C, ST, FT\}$. The identification argument requires to assign a neighboring grid $\Delta \in \{-n, n\}$ with $n \in \mathbb{N}_0$. By default, n is equal to 1.[23]

Table 2 reports the estimates of (4.15) for control, semi-treatment and full-treatment and results of (4.16) for aggregate demand. The aggregate estimates provide evidence that exogenous promotional treatments have an impact the preference primitives of the RCTs' participants. The estimates of marginal utility for no consumption (α_0) are consistent across different treatments' combinations. For each column, the results show how the consumers value one unit of good l or of good h more than one unit of the outside option ($\alpha_{l(h)} > \alpha_0$). At the aggregate level, the marginal utility of good l (α_l) is almost twice than the baseline utility of consuming one unit of the outside good, given the same normalized unit price. While the marginal utility of purchasing one unit of good h (α_h) at discounted unit price $(1 - d_t)$ is four time greater.

Table 2: Estimates of demand parameters

	Control	Semi-Treatment	Full-Treatment	Aggregate Demand
α_0	0.483	0.451	0.623	0.563
	[0.396,0.515]	[0.331,0.458]	[0.559,0.681]	[0.491,0.616]
α_l	1.419		1.300	1.045
	[1.392,1.511]		[1.215,1.333]	[0.973,1.093]
γ_l	0.537		0.726	0.959
	[0.495,0.534]		[1.333,0.736]	[0.937,0.976]
μ_l	0.453		0.302	0.045
	[0.443,0.480]		[0.281,0.321]	[0.018,0.056]
σ_l	1.876		1.738	1.111
	[1.785,1.903]		[1.708,1.840]	[1.060,1.186]
α_h		1.440	1.077	2.196
		[1.363,1.480]	[1.031,1.152]	[2.119,2.247]
γ_h		0.510	0.928	0.636
		[0.486,0.525]	[0.902,0.942]	[0.622,0.659]
μ_h		0.477	0.077	0.392
		[0.462,0.500]	[0.055,0.093]	[0.363,0.403]
σ_h		1.834	1.203	1.683
		[1.770,1.901]	[1.145,1.259]	[1.606,1.732]

Note: The results of the structural estimation of Θ for control group, semi-treatment group, full-treatment group and aggregate demand. Parameters α_0, α_l and α_h represent the marginal utilities of numeraire and of goods l and h, γ_l and γ_h correspond to the satiation parameter for each interior goods, and $\mu_{l(h)}$ and $\sigma_{l(h)}$ are the mean and the standard deviation of the normally distributed error $\epsilon_{l(h)}$ which accounts for good-specific unobserved heterogeneity. 95% bootstrap confidence intervals are in square brackets (501 repetitions). The dataset includes sessions in the control group and only sessions where the treatment was received (1,800,446 observations).

[23]Since the choice of $n = 1$ could be deemed arbitrary, a sensitivity analysis is performed in which n is set equal to 5 and 10 to evaluate the sensitivity of counterfactual estimates. Results show that the estimates with $n = 1$ are robust across different neighboring grid magnitudes.

Furthermore, the good l exhibits a higher satiation rate than the discounted good h ($\gamma_l > \gamma_h$). This result reflects the effect on the demand's preferences of the percentage difference in unit prices of the two interior goods. Because a unit of good l costs more than good h, γ_l implies weaker preference, or greater satiation, for good l with respect to good h. Table 2 also shows the estimates of mean and standard deviation of good-specific heterogeneity. The mean of the idiosyncratic demand shocks' distribution for good l is lower than that for good h ($\mu_l < \mu_h$).

4.4.1 Elasticity of Demand to Treatments

The estimation of demand primitives allows to compute demand elasticity. The own-price and cross-price elasticities are derived in Section A.3. Table 3 reports the price elasticity $\varepsilon_{l(h)}$ and the cross-price elasticity ε_0 where the estimates measure the percentage change in demand associated to a 10% change in the unit price. The exogenous endowment M_i is set to €101 since the upper level of Q_i would be €100, to guarantee that the demand of the outside good is always strictly positive. The aggregate results show that the demand of discounted good h is less elastic than the demand of full-price good l. A 10% reduction of the unit price would increase the demand of good l by 23% and the demand of good h by 12%. These estimates support the hypothesis of diminishing responsiveness in the demand to promotional pricing treatments at the intensive margin. Ceteris paribus, a further decrease in price would affect the demand for the good h, which already has a discounted $(1 - d_t)$ unit price, by 47% less than the demand for the good l. The cross-price elasticity is zero, implying that a 10% discount has no impact on the demand for the numeraire. Changes in the price of interior goods would not affect the demand for the outside good. Table C.6 presents the estimates of different levels of price change. Demand elasticity increases as the magnitude of discounts raises. Overall, a discount percentage between 5 and 20% increases the demand of good l between 21 and 26% and between 12 and 13% for the good h.

Table 3: Estimates of treatment elasticity

	Control	Semi-Treatment	Full-Treatment		Aggregate Demand	
	Low (l)	High (h)	Low (l)	High (h)	Low (l)	High (h)
$\varepsilon_{l(h)}$	-1.312	-1.272	-1.639	-2.434	-2.307	-1.235
	[-1.468,-1.221]	[-1.381,-1.147]	[-1.791,-1.633]	[-2.536,-2.384]	[-2.407,-2.186]	[-1.305,-1.116]
ε_0	0.022	0.023	0.015	0.012	0.011	0.018
	[-0.047,0.204]	[-0.178,0.028]	[-0.046,0.108]	[-0.043,0.092]	[-0.113,0.097]	[-0.029,0.178]

Note: Own-price and cross-price elasticities derived in Section A.3. Discount is set to 10% and endowment M_i to €101. 95% bootstrap confidence intervals are in square brackets (501 repetitions). The dataset includes sessions in the control group and only sessions where the treatment was received with a order gross equal to or lower than €100.

In conclusion, the demand analysis documents that promotional discount treatments are salient to users. On average, a 10% discount leads to a demand increase between 12 and 23%. The customers that pay non-discounted price exhibit a more elastic demand than users that comply to MPT_t and benefit d_t. Furthermore, the sensitivity of demand differs depending on the interaction between discount and minimum purchase constraints. Hence, the heterogeneity of demand's elasticity plays a crucial role in

designing how promotional treatments should be assigned to different groups of customers. To address the motives behind costumers' segmentation and promotion discrimination, the next chapter studies the pricing problem of a monopolistic e-commerce and investigates how personalized treatments' algorithms impact on expected profit.

5 Supply: Endogenous Promotional Policy

This chapter explores how the e-commerce platform can take advantage of the heterogeneity in users' demand elasticity to maximize its profit, and how the customers respond to endogenously assigned promotional treatments. This framework allows to estimate the causal impact of promotional discounts on the expected profit, and to compare the profitability of different pricing algorithms. A policy model empirically simulates personalized discount algorithms (Section 5.1) and exploits variation in the RCT treatments and users' characteristics to estimate expected profit (Section 5.2). The counterfactual analysis compares the theoretically derived, profit-maximizing promotional policy with alternative pricing policies (Section 5.3). At last, the definition of behavior-based promotion discrimination is introduced within this supply framework (Section 5.4).

5.1 Policy Model

This section models the statistical decision problem of optimal pricing policy for an e-commerce company in the spirit of Berger (2013), and deepens the work on personalized discount targeting by Miller and Hosanagar (2020). The model fits the characteristics of the RCT promotional treatments, namely the interaction between the discount percentage d_t and the minimum purchase threshold MPT_t on the e-commerce spending. By making the treatment an endogenous variable of the model, the analysis defines the condition for the optimal assignment of promotional treatments, focuses on the profit-maximizing assignment rule, and estimates the causal impact of promotions on the demand's margins.

The model considers a monopolistic policy-maker firm F and a continuous mass of users, indexed by i. As user i accesses the e-commerce website, firm F costlessly learns the vector of observable covariates X_i and has monopolistic access to user i's information. To simplify the pricing problem and focus on the personalized assignment of promotional treatments, the model assumes that firm F competes on prices à la Bertrand to set the prices of thousands of products available on the e-commerce website, as typical of digital retail platforms. In this sense, e-commerce firm F is price-taker, but given its monopolistic information advantage over user's i past purchase records, firm F is also price treatment-maker. Firm F has to select the treatment T_i in the treatment space $\mathcal{T}$ for each user i. Any treatment T_i is an unordered pair of discount percentage d_t and minimum purchase threshold MPT_t of gross order Q_i. Thus, $T_i = \{d_t, MPT_t\}_i$ defines the treatment combination that user i is assigned to by the e-commerce. Given T_i, user i decides whether to purchase or not, through the binary variable $Y_i \in \{0,1\}$. If $Y_i = 1$, user i converts the session and selects the gross order Q_i which is defined as $Q_i = \sum_{k=1}^{K} p_k Q_{ki}$ and is currency-measurable. Firm F observes y_i and q_i, namely the realization of the stochastic variables Y_i and Q_i. Marginal costs c are assumed zero. The problem of firm F is to select a priori the promotional treatment $T_i^* \in \mathcal{T}$ which maximizes the expected profit.

The conditional response functions are defined as

$$\mu_{it}(x) = \Pr(Y_i = 1 \mid T_i = t, X_i = x), \tag{5.1}$$

$$q_{it}(x) = \mathbb{E}[Q_i \mid Y_i = 1, T_i = t, X_i = x], \tag{5.2}$$

$$\gamma_{it}(x) = \Pr(Q_i \geq MPT_t \mid Y_i = 1, T_i = t, X_i = x). \tag{5.3}$$

Equation 5.1 describes the probability that, conditional on characteristics x and given treatment t, user i converts the session. Equation 5.2 specifies the expected gross order Q_i as a random variable, conditional on user i conversion, and Equation 5.3 defines the conditional probability that expenditure Q_i is equal to or greater than threshold MPT_t. Given observables x, the expected profit of user i for treatment t is

$$\mathbb{E}\big[\Pi_i(t \mid x)\big] = \mu_{it}(x)\gamma_{it}(x)\Big(q_{it}(x)(1-d_t) - c\Big) + \mu_{it}(x)\Big(1-\gamma_{it}(x)\Big)\Big(q_{it}(x) - c\Big). \tag{5.4}$$

Equation 5.4 captures two mutually exclusive events. Either user i converts the session with a probability $\mu_{it}(x)$ and $q_{it}(x)$ is greater than or equal to MPT_t with a probability $\gamma_{it}(x)$, then they benefit the discount d_t and pay $q_{it}(x)(1-d_t)$. Or user i purchases, selects $q_{it}(x)$ below MPT_t, and pays $q_{it}(x)$. For each user i, firm F selects the optimal treatment $T_i \in \mathcal{T}$ which maximizes the expected profit. The following equation formalizes the assignment rule which captures the promotional treatment policy decision by firm F:

$$\delta_i^*(x) = \underset{t \in \mathcal{T}}{\operatorname{argmax}} \left[\sum_{t=1}^{\dim(\mathcal{T})} \mu_{it}(x)\Big[\gamma_{it}(x)q_{it}(x)(1-d_t) + (1-\gamma_{it}(x))q_{it}(x)\Big] \right]. \tag{5.5}$$

Despite its conciseness, Equation 5.5 explains the logic of counterfactual analysis. Among all the possible pairs T_i in the space $\mathcal{T}$, firm F must select, a priori, that treatment set T_i^* which maximizes the expected profit of user i. Nevertheless, the theoretical challenge addressed by this formulation consists in a counterfactual comparison between the hypothetical optimal scenario and different unobservable scenarios, in terms of expected profit. The next section derives the condition for the profit-maximizing treatment, discusses the fundamental problem of counterfactual evaluation, estimates the optimal assignment rule, and implements policy evaluation for counterfactual analysis.

5.1.1 Condition of Optimal Treatment

The unordered pairs $T_i = \{d_t, MPT_t\}_i$ encompasses different pricing schemes depending on the magnitude of d_t and MPT_t as described in Section 2. If $T_i \in \{0\}^2$, user i receives no promotional treatment and the control group applies, namely $t = C$. It defines a uniform non-promotional policy and represents the benchmark case of any other policy intervention. Instead, the linear pricing policy imposes a linear relationship between discounts and demand, namely $T_i \in \mathbb{R}^+ \times \{0\}$. In this case, semi-treatment $t = ST$ occurs when user i receives discount d_t independently to the realization of gross order Q_i. Conversely, non-linear pricing arises if $T_i \in \mathbb{R}^+ \times \mathbb{R}^+$, that is when the minimum purchase threshold MPT_t imposes a necessary condition on user i's demand intensive margin Q_i in order to benefit the discount d_t. This case defines full-treatment $t = FT$.[24]

The conditional average treatment effects (CATEs) on the conversion rate $\mu_{it}(x)$ and on the gross order

[24] Lastly, $T_i \in \{0\} \times \mathbb{R}^+$ is trivial and not considered in the discussion.

$q_{it}(x)$ are respectively

$$\Delta\mu_{it}(x) = \mu_{it}(x) - \mu_{iC}(x), \tag{5.6a}$$

$$\Delta q_{it}(x) = q_{it}(x) - q_{iC}(x). \tag{5.6b}$$

Equations 5.6a and 5.6b quantify the average causal impact of treatment t on user i demand's extensive and intensive margins, relative to the counterfactual control case $t = C$, where treatment set $\mathcal{T} = \{C, SF, FT\}$ and $t \in \mathcal{T}$. As the control group defines the benchmark of CATEs, the condition for optimal treatment follows an equivalent intuition, namely

$$\mathbb{E}\big[\Pi_i \mid T_i = t, X_i = x\big] > \mathbb{E}\big[\Pi_i \mid T_i = C, X_i = x\big]. \tag{5.7}$$

The inequality 5.7 imposes that firm F selects the treatment $T_i = t$ for customer i with covariates $X_i = x$ as long as the associated expected profit is strictly greater than the expected profit in the counterfactual case of no treatment $t = C$. The CATEs in (5.6a) and (5.6b) allow (5.7) to be reformulated as follows

$$\underbrace{\Big(\Delta\mu_{it}(x)\Big)\Big(\Delta q_{it}(x)\Big) + q_{iC}(x)\Big(\Delta\mu_{it}(x)\Big) + \mu_{iC}(x)\Big(\Delta q_{it}(x)\Big)}_{\text{Causal effect of treatment } t \text{ on expected revenue}} > \underbrace{\Big(\mu_{it}(x)\gamma_{it}(x)q_{it}(x)\Big)d_t}_{\text{Expected cost of treatment } t}. \tag{5.8}$$

The condition for optimal treatment implies that the causal effect of treatment t on the expected revenue has to dominate the expected cost of treatment t. With a positive product of CATEs, the first term of expected revenue raises the left-hand side of (5.8) with concordant signs for the second and third term. Whereas with a negative product of CATEs, the first term of expected profit reduces the left-hand side of (5.8) with opposite signs for the second and third term. The cost of treatment t is unequivocally positive, while the signs and the magnitudes of the terms quantifying the causal impacts of treatment t on conversion probability and e-commerce spending, are the object of empirical estimation.

5.1.2 Optimal Treatment Policy

The optimal treatment policy function incorporates the economic intuition of (5.8) in the following proposition.

Proposition 2. *Optimal treatment policy is defined as*

$$f_{it}(x) = \Delta\mu_{it}(x)\Delta q_{it}(x) + q_{iC}(x)\Delta\mu_{it}(x) + \mu_{iC}(x)\Delta q_{it}(x) - \mu_{it}(x)\gamma_{it}(x)q_{it}(x)d_t. \tag{5.9}$$

Proof. Refer to Section A.4 in the Online Appendix.

The optimal treatment policy applies the necessary condition of optimal promotional treatment. For any treatment t such that $f_{it}(x) > 0$, it is strictly profitable to assign treatment t to user i. Conversely, for all the treatment t such that $f_{it}(x) \leq 0$, firm F maximizes its profit by assigning user i to the control group.

As both CATEs converge to zero, the optimal policy approaches to zero. Thus, for promotional treatments with no causal impact on either the conversion likelihood or the e-commerce spending, the optimal treatment policy would recommend no promotions. Instead, for elastic responses at the extensive margin or at the intensive margin (or both), the variation of $\Delta\mu_{it}(x)$ and $\Delta q_{it}(x)$ is driven by $\mu_{it}(x)$ and $\mu_{iC}(x)$, and by $q_{it}(x)$ and $q_{iC}(x)$, respectively (Figure D.8). This results in a comparative, non-trivial

behavior of $f_{it}(x)$, affected by both the marginal expected gain of treatment on revenue and the expected cost of treatment. Finally, the assignment rule in (5.5) is rearranged in terms of the optimal promotional treatment function $f_{it}(x)$ as follows

$$\delta_i^*(x) = \begin{cases} t & \underset{t \in \mathcal{T}_{-\{C\}}}{\text{argmax}}\, f_{it}(x) > 0 \\ C & \text{if } f_{it}(x) \leq 0 \quad \forall\, T_i \in \mathcal{T}_{-\{C\}}. \end{cases} \tag{5.10}$$

The algorithmic rule defined by Equation 5.10 compares the values of all the strictly positive policies $f_{it}(x)$ among pairs $T_i \in \mathcal{T}_{-\{C\}}$, and selects that T_i^* which maximizes $f_{it}(x)$. If none of the treatment satisfies this condition, then $T_i^* = C$ is the expected profit-maximizing treatment of user i with covariates x. Given a finite set of treatments $\mathcal{T}$, the assignment criterion in (5.10) simulates a personalized promotional content algorithm.

5.2 Estimation of the Policy Model

This section discusses the estimation methodology in three steps. The first step presents the data collected in the online field experiment and the variables used to estimate the policy model. The second step estimates the conditional responses and CATEs for the demand's extensive and intensive margins. The primitives of the model $\mathcal{P} = \{\mu_{it}(x), q_{it}(x), \gamma_{it}(x), \Delta\mu_{it}(x), \Delta q_{it}(x)\}$ for $t \in \mathcal{T}$ are estimated by matching the demand's margins of users with similar covariates and different treatments. In the third step, the optimal assignment rule for a sample subset n of covariates X_n is simulated, given the estimated model primitives $\widehat{\mathcal{P}}$.

Experiment. In the experimental phase firm F runs a RCT experiment. User i is randomly assigned to one of the three treatment groups $\mathcal{T} = \{C, SF, FT\}$. The realization of each pair $T_i = \{d_t, MPT_t\}_i$ depends on exogenously specified probabilities that are known to the econometrician (Section 2). The representative sample is defined by the set $\mathcal{S} = \{Y_i, C_i, Q_i, T_i, X_i\}_{i=1}^{I}$. It reports the purchased decision Y_i, the MPT_t compliance response C_i, the realized gross order Q_i, the assigned treatment T_i, and the vector of covariates X_i of the I users in the experiment. Table 4 summarizes the variables retained in the vector X_i from a broader set of individual characteristics, while Table B.1 provides further information of the experimental dataset $\mathcal{S}$.

Estimation. The estimation procedure applies the method presented in Chernozhukov et al. (2018) to evaluate the primitives of the model via cross-validation. It implements ensemble learning by defining an arbitrary supervised machine learning (ML) algorithm and leveraging repeated splits of the experimental data sample to construct ensemble estimates across different subsamples. For each iteration of the sample split, the data set is randomly divided into calibration (or training) sample and validation (or testing) sample. The supervised learning algorithm is trained on the calibration subsample and employed for out-of-sample predictions on the validation subsample. Furthermore, the obtained predictions are used in a constructed regression on the observed dependent variables of the validation split. Eventually, the estimated coefficients of both predictions and regressions are averaged and recombined to improve the predictive performance of ensembled estimates.

Table 4: Variables in the vector of covariates X_i

Variables	Aggregate Observations	Mean	St.Dev.
Time to session (minute)	5,087,685	12.96	17.64
Number of pages	5,087,685	14.31	15.95
Attained cart (0/1)	5,087,685	0.29	0.45
Time to cart (minute)	1,456,608	9.83	11.90
New user (0/1)	5,087,685	0.74	0.48
Purchase history (0/1)	5,087,685	0.03	0.17
Number of sessions	5,087,685	1.77	3.36
Number of prior visits	1,923,779	8.89	22.12
Number of prior purchases	1,923,779	0.18	1.25
Number of prior abandoned carts	638,710	3.78	8.51

Note: Covariates of RCT dataset used in Algorithm 6 to estimate the primitives $\mathcal{P}$ in classification and regression ensemble ML.

The implementation of estimation uses nonparametric supervised machine learning algorithms. The goal of supervised learning is to learn a function that maps inputs (i.e., independent variables) into a labeled output (i.e., dependent variable). In general, this class of models has three main advantages. First, the lack of assumptions regarding the underlying function increases the predictive power of the algorithm. Second, the flexibility of the model allows a large number of functional forms to fit the underlying data generating process. Third, the predictive performances are improved with respect to constrained, simpler parametric algorithms. The analysis leverages ensemble methods to improve the robustness of statistics by combining the predictions of different estimators built with multiple learning algorithms. Two families of ensemble methods are considered: averaging methods and boosting methods. Averaging methods construct several estimators independently and then average their predictions. Instead, boosting methods built base estimators sequentially and weighted combination of estimators aims at reducing the bias. In particular, bagging method (BAG) and random forest (RF) are considered for averaging methods, while adaptive boosting (AB), gradient boosting (GBDT) and neural network (NN) are evaluated for boosting methods.[25] For each method, the supervised learning is trained on the dataset $\mathcal{S}$ for classification with conversion decision Y_i as the labeled target, and for regression with order spending Q_i as the prediction output. In the Online Appendix, Section E.1 presents a detailed discussion of ensemble ML methods.

The comparison of predictive performance between BAG, RF, AB, GBDT, and NN depends on the distributions of $\widehat{\mu}_{it}(x)$ and $\widehat{q}_{it}(x)$. For classification methods, accuracy score $a(\mu_{it}(x), \widehat{\mu}_{it}(x))$ is computed, while for regression methods the coefficient of determination $R^2(q_{it}(x), \widehat{q}_{it}(x))$ is calculated. The analysis is performed for all the treatment groups and it confirms that `GradientBoostingClassifier` and `GradientBoostingRegressor` are the best-performing prediction models for both classification and regression. Indeed, GBDT Classifier predicts 81% of the variation in validation sample, whereas GBDT Regressor is able to match the gross order distribution in testing sample by 40% (Tables C.7 and C.8).

[25]Surveys such as Sagi and Rokach (2018) and Mienye and Sun (2022) present a comprehensive review of traditional, novel and state-of-the-art ensemble methods together with applications and examples. For more technical discussion about ensemble deep learning Yang et al. (2023) provides a valid reference.

Gradient boosting decision tree is used to build ensemble predictions which aggregate the estimators of 100 GBDTs, calibrated on 80% of the training sample of dataset $\mathcal{S}$, stratified by the treatment conditions. The outcomes are estimates of the conditional responses and treatment effects' functions for all the promotional treatments $\widehat{\mathcal{P}} = \{\widehat{\mu}_{it}(x), \widehat{\mu}_{iC}(x), \widehat{q}_{it}(x), \widehat{q}_{iC}(x), \widehat{\Delta\mu}_{it}(x), \widehat{\Delta q}_{it}(x), \widehat{\gamma}_{it}(x)\}$. At last, the estimates of the primitives of the policy model are combined to estimate the discrete realizations of $f_{it}(x)$ to extrapolate the distribution on targeting policy function $\widehat{f}_{iST}(x)$ and $\widehat{f}_{iFT}(x)$ (Figure D.9). The pseudocode of the GBDT estimation procedure is illustrated in the Online Appendix (Algorithm 6).

Simulation. In the policy model, how firm F selects the promotional treatment $T_i^* \in \mathcal{T}$ which maximizes expected profit, responds to the optimal assignment rule. Given estimates $\widehat{\mathcal{P}}$, (5.10) becomes

$$\widehat{\delta}_i^*(x) = \begin{cases} t & \underset{t \in \{ST,FT\}}{\operatorname{argmax}} \widehat{f}_{it}(x) > 0 \\ C & \text{if } \widehat{f}_{it}(x) \leq 0 \quad \forall\, T_i \in \{ST, FT\}, \end{cases} \tag{5.11}$$

where the targeting policy function is replaced by its estimates $\widehat{f}_{iST}(x)$ and $\widehat{f}_{iFT}(x)$, and the treatment set $\mathcal{T} = \{C, SF, FT\}$. The algorithm works by comparing $f_{it}(x)$ among treatment groups by varying the realizations of d_t and MPT_t. Since the realization probabilities of d_t and MPT_t are exogenously determined by the experimental design, then the algorithm normalizes the comparison of treatment policies between groups given an exogenous variation within groups. Equation 5.11 replicates the optimal assignment rule for covariates X_n and allows to simulate the distributions of conversion probability and expected profit under targeting policy function $f_{it}(x)$.[26]

5.3 Counterfactual Analysis

5.3.1 Alternative Policies

In order to compare the impact on the expected profit of the optimal treatment policy in (5.9) through (5.11), alternative treatment policy functions are considered. These functions are defined according to a different necessary condition relative to (5.7). While the optimal treatment policy assumes that any treatment has a potential, causal effect on both the demand's extensive and intensive margins, this section introduces alternative treatment policies that consider different ways in which the demand responds to promotional treatments. First, the uniform, non-treatment policy function is

$$f(x) = -1, \tag{5.12}$$

and represents the non-treatment, non-personalized benchmark of policy evaluation. It corresponds to a uniform pricing scheme where all users are assigned to the control group, irrespective of their characteristics X_i. Regardless the causal effects of promotional treatments on consumers' demand, firm F commits itself not to promotion discriminate and offers no discounts.

Second, the assignment criterion of non-targeted treatment policy averages, for each treatment, the causal effect among all users, and only assigns the promotional treatment with the greatest average impact on the aggregate demand. This treatment function is referred as non-targeted because it disregards the

[26] Simulation subsample size n is set 20% of the whole dataset, such as the validation sample of ML cross-validation.

heterogeneity of demand elasticity to promotional treatments among different users. The non-targeted treatment policy results as follows

$$f_t(x) = \mathbb{E}_x[\mu_{it}(x)]\mathbb{E}_x[q_{it}(x)] - \mathbb{E}_x[\mu_{iC}(x)]\mathbb{E}_x[q_{iC}(x)] - \mathbb{E}_x[\mu_{it}(x)]\mathbb{E}_x[\gamma_{it}(x)]\mathbb{E}_x[q_{it}(x)]d_t, \tag{5.13}$$

where the function is treatment-specific, but not user-specific due to the introduction of the expectation operator in $f_{it}(x)$. Indeed, all the terms varying in x are replaced by expectations, for any treatment t.

In addition, two alternative policies specifically target the extensive and the intensive margin of users' demand. The conversion-targeted treatment policy and the revenue-targeted treatment policy capture the causal impact of promotions only on the conversion probability and on the e-commerce spending, respectively. These two policies correspond to the following functional specifications

$$f_{it}(x) = \Delta\mu_{it}(x), \tag{5.14}$$

$$f_{it}(x) = \Delta q_{it}(x). \tag{5.15}$$

These are two treatment-specific, user-specific policies where the targeting conditions are defined by CATEs $\Delta\mu_{it}(x)$ and $\Delta q_{it}(x)$, respectively. Equations 5.14 and 5.15 assess the effectiveness of treatment T_i, conditional on observables x, based on the extent to which the treatment marginally influences the probability of conversion or the magnitude of expenditure. The differences in treatment elasticity at the extensive and intensive margin are captured by the conversion-targeted and the revenue-targeted policy respectively, while the complementary effect on the two margins is not accounted, similarly to the expected cost in (5.8). The conversion-targeted function targets those customers who are highly reactive in their likelihood of purchasing on e-commerce, whereas the revenue-targeted function focuses on those users who are very responsive in their e-commerce spending propensity.

At last, the conversion-revenue treatment policy function

$$f_{it}(x) = \Big(\mathbb{1}_{\{\Delta\mu_{it}(x)>0\}}\Delta\mu_{it}(x)\Big)\Big(\mathbb{1}_{\{\Delta q_{it}(x)>0\}}\Delta q_{it}(x)\Big)\gamma_{it}(x) \tag{5.16}$$

combines (5.14) and (5.15) weighted by the probability of complying to MPT_t for (5.3).[27] In this case, the assignment condition of Equation 5.16 imposes that both CATEs must be strictly positive, that is the considered treatment has to increase the conversion rate and to shift rightward the expected spending distribution for user i. The conversion-revenue treatment function targets a limited, very responsive set of users that exhibit uplifting marginal effects on both demand's margins.

5.3.2 Policy Evaluation

The set $\mathcal{S} = \{y_i, c_i, q_i, t_i, x_i\}_{i=1}^{I}$ represents the realized RCT sample. It collects the observed conversion decision y_i and compliance response c_i, the realized gross order q_i, the assigned treatment t_i, and the vector of covariates x_i of the I users in the field experiment. The goal of evaluating different policies relates to the problem of counterfactual policy evaluation in the literature on machine learning. If a user randomly received a certain treatment in the RCT, while the targeting policy would have assigned them to a different treatment, the counterfactual response is not observable. Since the realized sample set $\mathcal{S}$

[27] Trivially, $\gamma_{it}(x) = 1$ for $T_i = C$ and $T_i = ST$.

features controlled propensity scores of the treatment groups, the Horvitz–Thompson estimator with known sampling probabilities provides unbiased counterfactual estimation (Horvitz & Thompson, 1952; Huber, 2014). The analysis of policy evaluation takes advantage of inverse probability weighting (IPW) estimator whose weighting reduces the bias of unweighted estimators, yielding a consistent estimate of the profitability of the proposed policy.

Policy evaluation compares the performance, in terms of conversion rate and expected profit, of the alternative policies with different targeting function $f_{it}(x)$ for the same set $\mathcal{S}$. Thus, the inverse probability weighting estimator of probability of conversion is

$$\widehat{\mu}^{IPW}_{f_{it}(x)} = \frac{1}{I}\left[\sum_{i=1}^{I} \mathbb{1}_{\{t=C\}} \frac{y_i \mathbb{1}_{\{\widehat{f}_{it}(x)\leq 0\}}}{\Pr(t=C)} + \left(\frac{1}{\dim(\mathcal{T})-1} \sum_{t=1}^{\dim(\mathcal{T})-1} \mathbb{1}_{\{t=t_i\}} \frac{y_i \mathbb{1}_{\{\widehat{f}_{it}(x)>0\}}}{\Pr(t=t_i)}\right)\right], \tag{5.17}$$

and the IPW estimator of expected profit is

$$\widehat{\pi}^{IPW}_{f_{it}(x)} = \frac{1}{I}\left[\sum_{i=1}^{I} \mathbb{1}_{\{t=C\}} \frac{y_i \mathbb{1}_{\{\widehat{f}_{it}(x)\leq 0\}}}{\Pr(t=C)}(q_i) + \left(\frac{1}{\dim(\mathcal{T})-1} \sum_{t=1}^{\dim(\mathcal{T})-1} \mathbb{1}_{\{t=t_i\}} \frac{y_i \mathbb{1}_{\{\widehat{f}_{it}(x)>0\}}}{\Pr(t=t_i)}(q_i)\Big(1-c_i d_t\Big)\right)\right]. \tag{5.18}$$

Table 5 summarizes the estimates of (5.18) at the aggregate level for each targeting policy. The IPW estimator has the advantage to run counterfactual simulations by controlling for propensity scores and varying the treatment set's dimension. Therefore, Table 5 considers a reduced treatment space in the first and second column, namely semi-treatment and full-treatment respectively. In the third column, the profit estimates of the alternative promotional policies associated to the combined treatment space $\mathcal{T} = \{C, ST, FT\}$ are discussed. To make these estimates more intuitive, the comparison between alternative policies is framed in terms of percentage variation, defined as $\widehat{\Delta\pi}_{f_{it}(x)} = \left(\frac{\widehat{\pi}^{IPW}_{f_{it}(x)} - \widehat{\pi}^{IPW}_{uniform}}{\widehat{\pi}^{IPW}_{uniform}}\right)_{\%}$, where the non-treatment, uniform policy stands for the reference point. Table 6 considers differences in policy evaluation ascribable to demand heterogeneity between new and loyal customers. It reproduces the same counterfactual analysis of Table 5 conditional on the sessions of new users and returning users, respectively. The evaluation discussion focuses on the expected percentage change in profit, with the uniform pricing policy as the benchmark case. In the Online Appendix, Tables C.9 and C.10 replicate the same exercise for (5.17).

All the policies result in an increase in the expected profit, relative to the uniform treatment policy. The non-targeted policy does not account for potential demand heterogeneity in the responses of users, but allows for the assignment criterion to vary within in the treatment set. The estimate indicates that non-targeted promotions increase the expected profit on average by 32%. The reduction in the unit price induced by the non-personalized discount treatment stimulates the demand, and consequently raises the expected profit. Consistent with the estimates of elasticity in Section 4.4, this result confirms that price discounts, at the aggregate level, are salient to the customers. The non-targeted policy increases profitability compared to uniform policy function, but it is unable to capture the causal impact of personalized treatment assignment on profit.

The conversion-targeted policy and the revenue-targeted policy focus on the demand elasticity at the intensive and extensive margin, namely whether a certain promotional treatment has a positive impact on a specific user's conversion rate and e-commerce spending. The estimated effects of the treatment on

Table 5: Estimates of counterfactual simulations of expected profit

	Semi-Treatment		Full-Treatment		Combined Treatments	
Policy	$\widehat{\pi}^{IPW}_{f_{iST}(x)}$	$\widehat{\Delta\pi}_{f_{iST}(x)}$	$\widehat{\pi}^{IPW}_{f_{iFT}(x)}$	$\widehat{\Delta\pi}_{f_{iFT}(x)}$	$\widehat{\pi}^{IPW}_{f_{i\{ST,FT\}}(x)}$	$\widehat{\Delta\pi}_{f_{i\{ST,FT\}}(x)}$
Uniform	0.0030	–	0.0013	–	0.0016	–
Non-targeted	0.0011	–67.93%	0.0013	0.00%	0.0021	+31.77%
		(0.466)		(0.687)		(1.934)
Conversion-targeted	0.0031	+0.321%	0.0014	+8.10%	0.0021	+28.00%
		(1.156)		(0.671)		(1.813)
Revenue-targeted	0.0035	+15.21%	0.0014	+11.38%	0.0023	+38.98%
		(1.212)		(0.677)		(1.876)
Conversion-revenue	0.0035	+15.21%	0.0014	+11.42%	0.0023	+39.01%
		(1.212)		(0.679)		(1.876)
Optimal treatment	0.0036	+18.63%	0.0014	+12.61%	0.0023	+39.63%
		(1.218)		(0.679)		(1.881)

Note: The results of simulation analysis of expected profit of alternative policies and optimal treatment policy. First and second column consider the treatment space $\mathcal{T} = \{C, ST\}$ and $\mathcal{T} = \{C, FT\}$, respectively. The third column simulates the pricing algorithm with the combined treatment space $\mathcal{T} = \{C, ST, FT\}$. Comparison between alternative policies is framed in terms of percentage variation defined as $\widehat{\Delta\pi}_{f_{it}(x)} = \left(\frac{\widehat{\pi}^{IPW}_{f_{it}(x)} - \widehat{\pi}^{IPW}_{uniform}}{\widehat{\pi}^{IPW}_{uniform}}\right)_{\%}$, where the non-treatment, uniform policy stands for the benchmark case. The estimates of expected profit of (5.18) per counterfactual policies and optimal treatment policy uses dataset sessions where the treatment was received. The testing dataset is equal to 20%. Semi-treatment testing dataset features 19,908 observations, full-treatment dataset has 36,748 sessions, and the combined treatments testing dataset counts 56,656 data points. Standard errors in parenthesis.

the probability of purchase and on the gross order's magnitude identify those customers characterized by sufficiently elastic demand, who would be more likely to purchase or would purchase more, given a promotional incentive. The comparison of the two policies highlights how personalized discounts affect the expected profit through the intensive and the extensive margin, respectively. At the aggregate level, the revenue-targeted policy increases profit by 39%, 10 percentage points more than the conversion-targeted case (28%), suggesting that, on average, the consumers are more responsive at the intensive margin rather than at the extensive margin, for different promotional treatments. The algorithm targeting CATE on revenue results more effective at increasing expected profit than the other that prioritizes causal impact on conversion rate. Considering the absolute magnitude of margins corroborates this finding, as marginal variations in $\Delta\mu_{it}(x)$, rather than in $\Delta q_{it}(x)$, result less effective to substantially increase the expected profit.

Furthermore, the conversion-revenue treatment policy combines the two previous policies in a more restrictive way. It targets a user only if the pricing treatment positively impacts both conversion probability and expected revenue relative to the control counterfactual. Estimates show that such pricing policy has the same return as the revenue-targeted policy (+39%), indicating that positive CATEs at the extensive margin are strongly associated to positive CATEs at the intensive margin. Finally, the optimal treatment policy is the profit-maximizing policy (+40%). It accounts for heterogeneity of the treatment effects at the extensive margin ($\Delta\mu_{it}(x)$) and intensive margin ($\Delta q_{it}(x)$), combined to the baseline response variation of the controlled group ($\mu_{iC}(x)$ and $q_{iC}(x)$) for discrete levels of treatment. Unlike all the alternative

policies, the optimal treatment function also takes into consideration the expected cost of discount ($\mu_{it}(x)\gamma_{it}(x)q_{it}(x)d_t$). The optimal treatment condition specifies a user-specific treatment policy under the condition that the causal impact of treatment t on the expected revenue dominates the expected cost of treatment t. Therefore, the results show that accounting for individual-specific demand heterogeneity, evaluating treatment-specific variation in the response of extensive and intensive margins and controlling for policy-specific expected costs, generates to the most profit-maximizing, personalized promotional treatment policy.

The reduced-form analysis and causal inference estimates in Section 3 support the hypothesis that demand heterogeneity among consumers is associated with different purchase histories. Loyal customers exhibit a higher propensity to purchase and to spend more, relative to new users. Moreover, estimated average treatment effects highlight how new and returning users respond differently to RCT treatments. On average, the demand of returning individuals at the extensive margin is more elastic to different level of discount d_t compared to new clients. In contrast, the variation in minimum purchase threshold MPT_t mainly affect response of first-time users at the intensive margin, while returning users are less affected by the spending constraint. The information about users' purchase records and session-specific behaviors is fundamental to assess the effectiveness of treatments and, consequently, the causal impact on targeted outcomes to improve algorithms of personalized promotions.

Table 6 replicates the counterfactual analysis of Table 5 and presents the estimates of counterfactual simulation on expected profit of two different segments of customers: new and returning users. Compared to inexperienced users, the loyal customers are more responsive to treatments on average. The optimal treatment policy raises the profit of returning users by 44%, while new users exhibit an 29% increase. The aggregate demand's estimate for the equivalent policy (40% in Table 5) falls between these two estimates, but is closer to 44% (returning users) rather than 29% (new users). This is consistent with the fact that returning users weight more in conditioning the profit as they are 46% of the e-commerce demand, despite representing only 26% of the entire RCT sample. The policy estimates of new users' segment are in line with the corresponding estimates at the aggregate level in Table 5. This comes at no surprise as these results are derived from policy simulation based on 74% of the dataset.

Remarkably, the optimal treatment policy fails to maximize the expected profit for the group of returning customers. The conversion-targeted policy leads to a 48% increase in profit of returning users, which generates an additional 11% profits' gain compared to the optimal treatment counterfactual. The economic intuition lies in the heterogeneity of treatments elasticity for extensive and intensive demand margins across behavior-specific users' groups. Because the segment of returning users is characterized by a very responsive extensive margin, the conversion-targeted policy performs better in addressing the heterogeneity of these customers. It uniquely targets the treatment effect on conversion rate and disregards the intensive margin effect. Conversely, since loyal consumers on average exhibit an inelastic demand at the intensive margin, any treatment policy based on the causal effect on e-commerce spending, does not optimally assign promotional treatments, resulting in a under-performing policy in terms of increasing expected profit.

This is a crucial discovery because it empirically corroborates the theoretical result that it is profitable for the e-commerce platform to adopt different promotional policies depending on whether a user

Table 6: Estimates of counterfactual simulations of expected profit for new and returning users

New users

	Semi-Treatment		Full-Treatment		Combined Treatments	
Policy	$\hat{\pi}^{IPW}_{f_{iST}(x)}$	$\widehat{\Delta\pi}_{f_{iST}(x)}$	$\hat{\pi}^{IPW}_{f_{iFT}(x)}$	$\widehat{\Delta\pi}_{f_{iFT}(x)}$	$\hat{\pi}^{IPW}_{f_{i\{ST,FT\}}(x)}$	$\widehat{\Delta\pi}_{f_{i\{ST,FT\}}(x)}$
Uniform	0.0113	–	0.0057	–	0.0069	–
Non-targeted	0.0113	0.00%	0.0057	0.00%	0.0069	0.00%
		(1.5504)		(1.1980)		(2.3773)
Conversion-targeted	0.0103	-8.20%	0.0065	+13.17%	0.0079	+15.03%
		(1.3008)		(1.1941)		(2.1248)
Revenue-targeted	0.0121	+7.86%	0.0063	+9.32%	0.0088	+28.45%
		(1.4779)		(1.1797)		(2.3059)
Conversion-revenue	0.0122	+8.01%	0.0065	+13.10%	0.0086	+25.54%
		(1.4787)		(1.1942)		(2.3084)
Optimal treatment	0.0126	+11.62%	0.0064	+10.73%	0.0089	+29.83%
		(1.4961)		(1.1883)		(2.3176)

Returning users

	Semi-Treatment		Full-Treatment		Combined Treatments	
Policy	$\hat{\pi}^{IPW}_{f_{iST}(x)}$	$\widehat{\Delta\pi}_{f_{iST}(x)}$	$\hat{\pi}^{IPW}_{f_{iFT}(x)}$	$\widehat{\Delta\pi}_{f_{iFT}(x)}$	$\hat{\pi}^{IPW}_{f_{i\{ST,FT\}}(x)}$	$\widehat{\Delta\pi}_{f_{i\{ST,FT\}}(x)}$
Uniform	0.0228	–	0.0113	–	0.0134	–
Non-targeted	0.0086	–62.46%	0.0113	0.00%	0.0184	+37.22%
		(2.1011)		(3.2380)		(6.2570)
Conversion-targeted	0.0288	+26.47%	0.012	+6.09%	0.0198	+48.33%
		(3.9209)		(3.1911)		(6.0011)
Revenue-targeted	0.0268	+17.43%	0.0123	+9.21%	0.0191	+42.44%
		(3.9530)		(3.2057)		(6.1085)
Conversion-revenue	0.0288	+26.44%	0.0123	+9.28%	0.0196	+46.17%
		(3.9210)		(3.2057)		(6.1085)
Optimal treatment	0.0277	+21.75%	0.0125	+11.13%	0.0192	+43.81%
		(3.9724)		(3.2117)		(6.1167)

Note: The estimates of the counterfactual analysis of alternative policies and optimal treatment policy conditional on new users' sessions and on returning users' sessions. First and second column consider the treatment space $\mathcal{T} = \{C, ST\}$ and $\mathcal{T} = \{C, FT\}$, respectively. The third column simulates the pricing algorithm described in (5.11) and implemented by Algorithm 6 where the treatment space $\mathcal{T} = \{C, ST, FT\}$ combines $t = ST$ and $t = FT$. The first section presents the estimates of expected profit's simulation analysis of (5.18) of new users per counterfactual policies and optimal treatment policy. The second section shows the results of expected profit's simulation analysis of (5.18) of returning users per counterfactual policies and optimal treatment policy. Comparison between alternative policies is framed in terms of percentage variation defined as $\widehat{\Delta\pi}_{f_{it}(x)} = \left(\frac{\widehat{\pi}^{IPW}_{f_{it}(x)} - \widehat{\pi}^{IPW}_{uniform}}{\widehat{\pi}^{IPW}_{uniform}}\right)_{\%}$, where the non-treatment, uniform policy stands for the benchmark case. For new users, the simulation dataset includes only sessions where the treatment was received and the testing dataset is set to 20%. Semi-treatment testing dataset features 5,517 observations, while full-treatment dataset has 8,353 observations. The combined treatments testing dataset counts 13,870 sessions. For returning users, semi-treatment testing dataset features 3,146 observations, while full-treatment dataset has 5,073 observations. The combined treatments testing dataset counts 8,219 sessions. Standard errors in parenthesis.

belongs to a certain segment, defined according to observable, purchase history-specific characteristics. In the next section, this intuition is formalized through the definition of behavior-based promotional discrimination, and introduced to study the welfare implications of BBPD.

5.4 Behavior-based Promotion Discrimination

Definition 1. *For any treatment policy f and for new users i and z, and for returning user j, behavior-based promotion discrimination arises if*

$$f_{it}(x|\, i) = f_{zt}(x|\, z) \quad \textit{and} \quad f_{it}(x|\, i) \neq f_{jt}(x|\, j).$$

Definition 1 specifies the necessary condition for BBPD to occur in the treatment policy model. Behavior-based promotion discrimination takes place when firm F adopts different treatment policies, depending on the history of users' past purchases. In this specific application, the behavior-based demand segmentation considers whether the user is new to the e-commerce digital environment or whether they have browsed the e-commerce before and they are returning to the platform. Based on results in Table 6, the new individuals are assigned to the optimal treatment policy, while the returning customers to the conversion-targeted policy. It is incentive compatible for the e-commerce company to behavior-based promotion discriminate at the following condition.

Proposition 3. *For new user i and for returning user j, behavior-based promotion discrimination is profitable for firm F only if there exist two treatment policies f and g such that*

$$\Pi_{f_{it}}(x|\, i) > \Pi_{g_{it}}(x|\, i) \quad \textit{and} \quad \Pi_{g_{jt}}(x|\, j) > \Pi_{f_{jt}}(x|\, j).$$

Proposition 3 provides the sufficient condition for behavior-based promotion discrimination to be profit-maximizing. From estimates of Table 6, BBPD generates an additional 11% increase in profit from loyal users' segment, compared to the optimal treatment policy that maximizes expected profit at the aggregate level. Moreover, behavior-based promotion discrimination complements personalized promotional treatments. With respect to the benchmark case of uniform policy, 89% of the extra profit derives from personalization of discounts and coupons, while 11% generates from behavior-based discrimination in treatment policies.

Because it has the opportunity to observe the users' history, the e-commerce firm leverages its informational advantage and exploits the heterogeneity of demand's elasticity between these customers' segments in order to increase the profit. Counterfactual analysis shows that discriminating between treatment policies is profitable. Behavior-based price discrimination complements personalized pricing. Relative to uniform pricing counterfactual, 89% of the extra profit derives from personalization of treatments, while 11% comes from behavior-based discrimination in treatment policies. The last chapter studies the welfare effects of BBPD on the consumers' surplus and e-commerce profit. The analysis combines the estimates of demand's primitives and the simulated promotional treatment algorithms of policy evaluation. The aim is to measure the percentage changes in consumers' and producer's welfare under different promotional treatment policies, as prescribed by Proposition 3.

6 Welfare Analysis

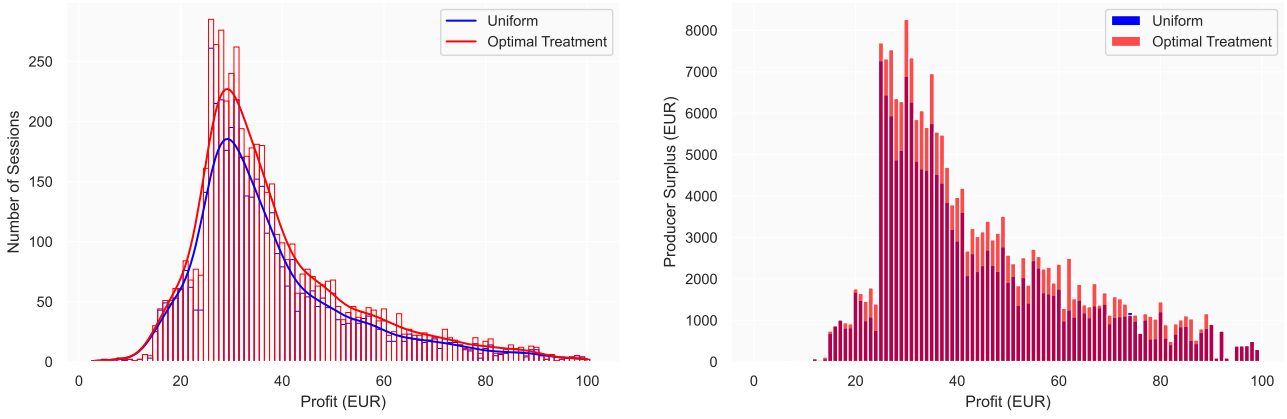
Does the optimal treatment policy increase the producer surplus relative to the uniform non-treatment policy at the aggregate level? How does behavior-based promotion discrimination interact with personalized promotions in terms of welfare implications? If so, is there a segment of users that benefits more from BBPD? By combining the demand estimates in Section 4.4 and the policy simulations of the counterfactual analysis in Section 5.2, this last chapter answers these research questions and concludes the paper.

6.1 Welfare Analysis with Promotion Discrimination

Given the simulated demand and the estimates of treatment policy's primitives derived in Section 5.2, for the treatment space $\mathcal{T} = \{C, ST, FT\}$, the producer surplus (PS) associated to the policy function $f_{it}(x)$ corresponds to the firm F's profit:

$$PS_{f_{it}(x)} = \sum_{i=1}^{I} \Pi\Big(q_{it}^{l}(x), q_{it}^{h}(x), f_{it}(x)\Big) = \sum_{i=1}^{I} (1 - \mathbb{1}_{\{f_{it}(x)>0\}} d_t)(q_{it}^{l}(x) + q_{it}^{h}(x)). \tag{6.1}$$

Figure 5a depicts the aggregate profit distributions associated to the uniform, non-treatment policy and the optimal treatment policy, for the aggregate demand below €101. The red curve dominates the blue one, indicating that the optimal treatment policy has a larger causal impact on both demand extensive and intensive margin. Intuitively, the optimal treatment algorithm targets better those customers whose demand is sufficiently elastic at the margin, so that the positive effect on revenue is strictly greater than the opportunity cost incurred by the firm in offering promotional treatments. At the aggregate level, the difference between the blue and red profit distributions reflects how the causal effects on revenues net to expected cost vary. At the peak of the profit distribution it is concentrated the mass



(a) Number of converted sessions by profit

(b) Producer surplus by profit

Figure 5: Aggregate producer surplus

Note: Panel (a) shows the distribution of aggregate profit. Panel (b) displays the distribution of aggregate producer surplus by profit levels. Sessions with gross order greater than €100 are not depicted. Blue plot refers to uniform policy (5,859 sessions with €28.38 mean), red plot to optimal treatment policy (5,859 sessions with €35.21 mean). Figures D.10a and D.10b provide additional details.

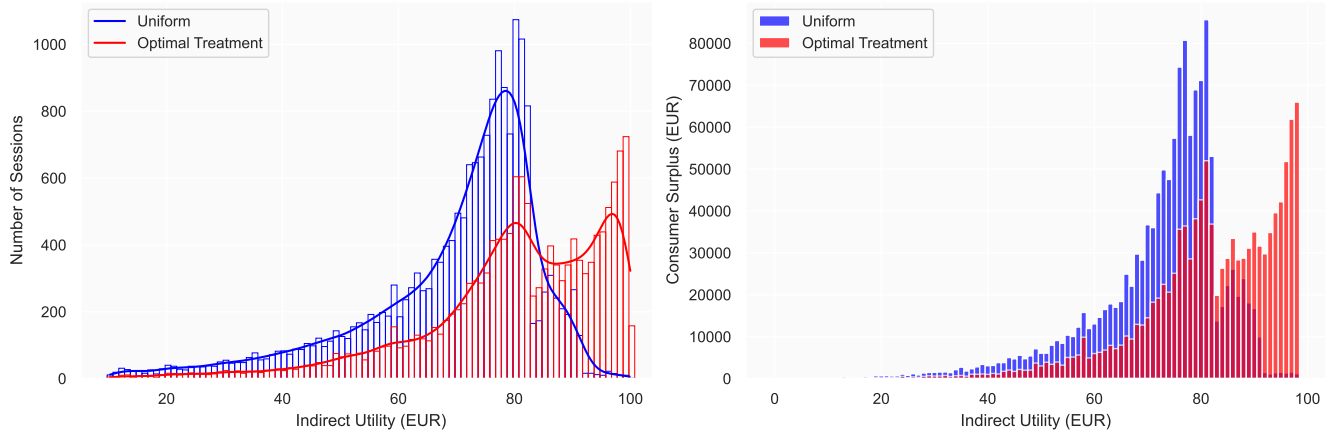
of users characterized by an elastic demand to discounts, which translates into greater gains in profits. Consequently, personalized discounts increase producer surplus compared to a uniform pricing scheme where no price promotions are customized by the digital platform. In Figure 5b profits are aggregated at the spending level and the bars measure the producer surplus. Since the blue area is completely covered by the red one, the producer surplus with uniform prices is lower than the producer surplus with customized promotions. According to the simulated distributions, the magnitude of producer welfare associated to the uniform and optimal treatment policy, are €166,251.4 and €206,311.4, respectively. The optimal treatment policy increases producer surplus by 24%, relative to the non-treatment policy:

$$\Delta PS = \left(\frac{PS_{optimal} - PS_{uniform}}{PS_{uniform}}\right)_{\%} = +24.16\%.$$

The consumer surplus (CS) is directly related to the definition of utility function in (4.10). Given the estimates of aggregate demand's parameters in Table 2, the simulated demand's extensive and intensive margins, and the estimates of treatment policy derived in Section 5.2, the consumer surplus is the sum of consumers' indirect utility for the policy function $f_{it}(x)$:

$$CS_{f_{it}(x)} = \sum_{i=1}^{I} \mathcal{U}^*\left(q_{it}^l(x), q_{it}^h(x), f_{it}(x)\right) = \tag{6.2}$$
$$\sum_{i=1}^{I} \frac{\alpha_l e^{\epsilon_l}}{\gamma_l} \ln(\gamma_l q_{it}^l(x) + 1) + \frac{\alpha_h e^{\epsilon_h}}{\gamma_h} \ln(\gamma_h q_{it}^h(x) + 1) + \alpha_0\left(M_i - (1 - \mathbb{1}_{\{f_{it}(x)>0\}} d_t)(q_{it}^l(x) + q_{it}^h(x))\right).$$

Figure 6a compares the uniform policy indirect utility distribution with that of optimal treatment policy. The uniform treatment policy is characterized by two demand equilibria: No Conversion and Low Conversion. For non-treatment pricing, the blue plot illustrates the welfare generated by the simulated demand of good l and the outside option x. The left, low-utility side of the distribution captures the share



(a) Number of converted sessions by indirect utility

(b) Consumer surplus by indirect utility

Figure 6: Aggregate consumer surplus

Note: Panel (a) shows the distribution of aggregate indirect utility. Panel (b) displays the distribution of aggregate consumer surplus by indirect utility levels. Sessions with gross order greater than €100 are not depicted. Blue plot refers to uniform policy (74,830 sessions with €69.57 mean), red plot to optimal treatment policy (74,830 sessions with €80.13 mean). Figures D.10c and D.10d provide additional details.

of users who do not purchase on e-commerce and prefer the outside option x. Whereas, the high-utility peak of the blue distribution represents all the individuals who buy on the e-commerce at full price, revealing their preference for good l. The optimal treatment policy expands the pricing menu by adding a third demand equilibrium, High Conversion. By combining the demand of good l and good h, the red curve features two peaks, depicting how consumer surplus is generated with personalized promotions. The consumer surplus by consumers purchasing good l below MPT_t at full price concentrates on the left red peak, while the demand of good h with discount d_t gravitates toward the right red peak of the consumer surplus distribution. The optimal treatment policy shifts the indirect utility distribution rightward. While the uniform policy sets a unit price for the e-commerce good, personalized promotional treatments expands the pricing menu through promotion discrimination, through the use of targeted discounts and coupons. Some individuals access the interior good at a discounted price, whereas other are incentivized to marginally increase their spending to benefit from the price reduction. Whether the optimal treatment policy also increases consumer welfare relative to the uniform policy, depends on how the simulated demands translate into indirect utility estimates. The bars of Figure 6b stack the consumer surplus for different levels of indirect utility. Graphically, the comparison between the red and blue cumulative distributions measures how much personalized treatments change consumer welfare. With non-treatment policy the absolute consumer surplus equal to €6,959,958.3 and optimal treatment consumer surplus €7,247,376.7, promotion discrimination through the optimal treatment policy increases the consumer welfare by 4%:

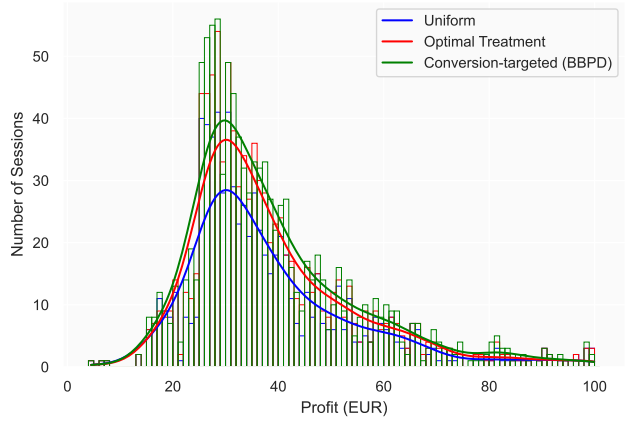
$$\Delta CS = \left(\frac{CS_{optimal} - CS_{uniform}}{CS_{uniform}} \right)_{\%} = +4.15\%.$$

The final section studies how behavior-based promotion discrimination influences the welfare outcomes of promotion discrimination.

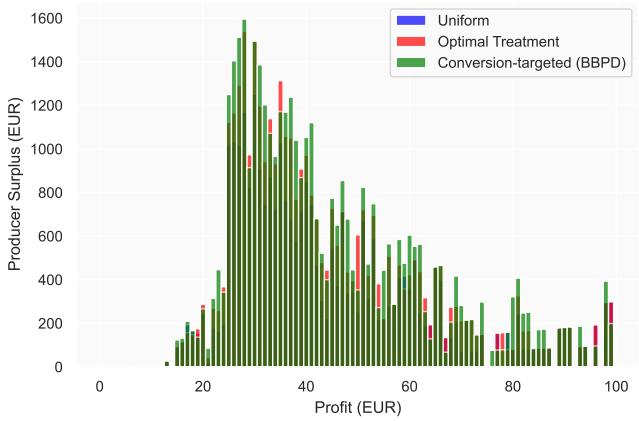
6.2 Welfare Analysis with Behavior-based Promotion Discrimination

According to Definition 1, behavior-based promotion discrimination involves adopting a shared treatment policy for all the returning users, and a different, common policy for all the new customers. The results of the last column in Table 6 support the hypothesis that implementing behavior-based promotional discrimination, particularly where loyal individuals are assigned to the conversion-targeted policy and new customers to the optimal treatment policy, can generate higher profits. For existing customers, the conversion-targeted policy maximizes the expected profit's estimator as also shown by the profit distributions in Figure 7a, where the green distribution corresponds to the deviation from the aggregate optimal treatment policy that defines BBPD. Figure 7b reports that the producer surplus associated to the non-treatment and the optimal treatment policies is €29,553 and €37,127.5 respectively, and the conversion-targeted policy generates a producer surplus equal to €41,290.[28] Therefore, the optimal treatment policy increases producer surplus by 25% compared to the uniform scenario, while behavior-

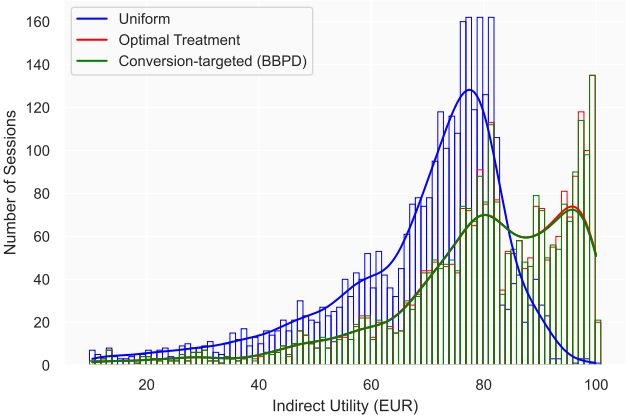
[28] In the Online Appendix, Figures D.11a and D.11b replicate the same analysis of producer surplus for new users. The producer surplus generated by the uniform and the optimal treatment policies is €44,935 and €53,160, resulting in an 18.30% increase in producer surplus.



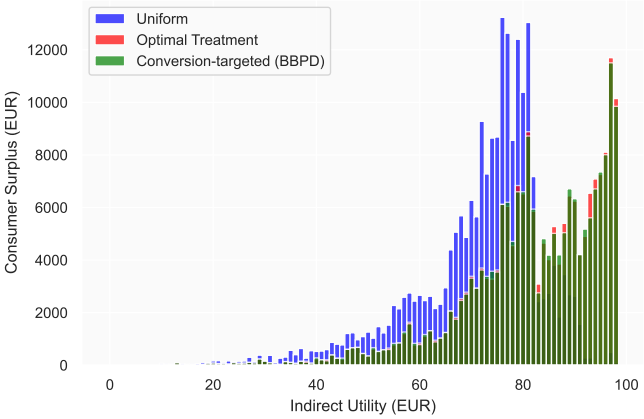
(a) Number of converted sessions by profit



(b) Producer surplus by profit



(c) Number of converted sessions by indirect utility



(d) Consumer surplus by indirect utility

Figure 7: Welfare analysis of behavior-based price discrimination

Note: Panel (a) shows the distribution of profit of returning users. Panel (b) displays the distribution of producer surplus of returning users by profit levels. Panel (c) illustrates the distribution of indirect utility of returning users. Panel (d) reports the distribution of consumer surplus of returning users by indirect utility levels. For returning users, 10,991 sessions with gross order below to €100 are depicted. Blue plot refers to uniform policy, red plot to optimal treatment policy and green plot to conversion-targeted policy. Figures D.12a and D.12b provide additional details for producer surplus. Figures D.12c and D.12d offer further graphical intuition for consumer surplus.

based promotion discrimination raises profit by an additional 11%, as follows

$$\Delta\widetilde{PS} = \left(\frac{PS_{BBPD} - PS_{optimal}}{PS_{optimal}}\right)_{\%} = +11.21\%. \tag{6.3}$$

As far as the consumers' welfare is concerned, the optimal treatment policy increases the consumer surplus of returning users (4.89%) more than the average, aggregate demand change (4.15%).[29] The indirect utility distributions for both the optimal treatment policy and the conversion-targeted policy are almost identical. The economic intuition is that, unlike the non-treatment policy, both policies offer the same pricing menu and provide access to good l. Moreover, the targeting mechanism of these two

[29]In the Online Appendix, Figures D.11c and D.11d reproduce the same analysis of consumer surplus for new users. The consumer surplus generated by the uniform and the optimal treatment policies is €993,651.5 and €1,029,956.5, resulting in an 3.65% increase in the consumer surplus.

policies differs in considering the expected cost of treatments, which is not captured in the specification of the utility function, therefore it is not reflected in different indirect utility distributions. Figures 7c and 7d in fact underline how consumer surplus is not influenced by the implementation of BBPD by the e-commerce, that is the firm switching from optimal treatment policy to conversion-targeted policy, as follows

$$\Delta\widetilde{CS} = \left(\frac{CS_{BBPD} - CS_{optimal}}{CS_{optimal}}\right)_{\%} = +0.01\%. \tag{6.4}$$

7 Conclusion

This paper investigates the welfare effects of behavior-based promotion discrimination. By developing a personalized promotions' model that allows for counterfactual simulations, it is shown that behavior-based personalized discounts lead to a welfare improvement for both the e-commerce retailer and its customers.

Methodologically, the empirical analysis exploits field experimental data on a cosmetics e-commerce website. Instead of employing observational data, the main advantage of using an RCT lies in the causal estimation of heterogeneous treatment effects. By leveraging the exogenous treatments' variation induced by the randomized trials, the structural demand addresses demand's heterogeneity and estimates the consumers' elasticity. Consequently, randomly assigned treatments are endogenized according to a pricing policy model. The supply model develops an optimal treatment algorithm, which is first trained via supervised machine learning, and then used to simulate alternative pricing policies. At last, the counterfactual policy analysis shows how the consumers' welfare and monopolist profitability vary under different promotional discrimination algorithms. Personalized promotion discrimination increases the e-commerce profit by 24% and the consumer surplus by 4%, compared to a uniform, non-treatment benchmark. Furthermore, the past purchase history is informative for the monopolistic firm, because implementing behavior-based promotion discrimination improves the producer welfare by an additional 11%. While the consumer surplus rises with personalized discounts, promotional discrimination based on consumer purchase histories seems not to harm loyal clients.

This paper argues that a digital company with monopolistic access to users' online history, adopts behavior-based promotion discrimination as a demand segregation mechanism. To improve the efficiency of personalized pricing algorithms, the e-commerce assigns history-specific promotional treatment policies to different consumers' segments. The results show that committing to a unique pricing policy is suboptimal for a monopolist whose users' purchase records are available. In case of first-degree imperfect price discrimination, the set of users' information maps into discount treatments, for a exogenously defined rule. Instead, third-degree behavior-based promotion discrimination improves the profit by inverting this relationship: different market segments defined according to purchase history, are assigned to different treatments' policies, which compute the profit-maximizing promotional content for each individual given user's covarites.

The policy implications are significant. The public debate concerning fairness of tailored pricing policies, fails to consider to what extent consumers benefit from personalized discounts, in addition to producer welfare consequences. Moreover, limiting the information that the e-commerce retailers

have access to for personalized price-related purposes, could result detrimental to consumer welfare. Compared to uniform price strategies, both sides of the market benefit from behavior-based promotion discrimination, as it increases total welfare, particularly for loyal customers. In conclusion, this paper extends the empirical marketing research on behavior-based promotion discrimination. By presenting a novel definition behavior-based promotional discrimination and contributing to the quantitative literature, this paper provides an original perspective in the current public policy debate regarding e-commerce pricing strategies.

References

Angrist, J. D., & Imbens, G. W. (1995). Two-stage least squares estimation of average causal effects in models with variable treatment intensity. *Journal of the American statistical Association*, *90*(430), 431–442.

Arora, N., Dreze, X., Ghose, A., Hess, J. D., Iyengar, R., Jing, B., ... others (2008). Putting one-to-one marketing to work: Personalization, customization, and choice. *Marketing Letters*, *19*, 305–321.

Athey, S., & Imbens, G. W. (2017). The econometrics of randomized experiments. In *Handbook of economic field experiments* (Vol. 1, pp. 73–140). Elsevier.

Baik, A., Simon Anderson, & Larson, N. (2023). Price discrimination in the information age: Prices, poaching, and privacy with personalized targeted discounts. *The Review of Economic Studies*, *90*(5), 2085–2115.

Belleflamme, P., & Peitz, M. (2015). *Industrial organization: markets and strategies*. Cambridge University Press.

Bergemann, D., Brooks, B., & Morris, S. (2015). The limits of price discrimination. *American Economic Review*, *105*(3), 921–957.

Bergemann, D., Castro, F., & Weintraub, G. (2022). Third-degree price discrimination versus uniform pricing. *Games and Economic Behavior*, *131*, 275–291.

Berger, J. (2013). *Statistical decision theory: foundations, concepts, and methods*. Springer Science & Business Media.

Bhat, C. R. (2008). The multiple discrete-continuous extreme value (mdcev) model: role of utility function parameters, identification considerations, and model extensions. *Transportation Research Part B: Methodological*, *42*(3), 274–303.

Breiman, L. (1996a). *Arcing classifiers* (Tech. Rep.). Technical report, University of California, Department of Statistics.

Breiman, L. (1996b). Bagging predictors. *Machine learning*, *24*, 123–140.

Brokesova, Z., Deck, C., & Peliova, J. (2014). Experimenting with purchase history based price discrimination. *International Journal of Industrial Organization*, *37*, 229–237.

Caillaud, B., & De Nijs, R. (2014). Strategic loyalty reward in dynamic price discrimination. *Marketing Science*, *33*(5), 725–742.

Carroni, E. (2018). Behaviour-based price discrimination with cross-group externalities. *Journal of Economics*, *125*(2), 137–157.

Chandra, S., Verma, S., Lim, W. M., Kumar, S., & Donthu, N. (2022). Personalization in personalized marketing: Trends and ways forward. *Psychology & Marketing*, *39*(8), 1529–1562.

Chen, Y., & Pearcy, J. (2010). Dynamic pricing: when to entice brand switching and when to reward consumer loyalty. *The RAND Journal of Economics*, *41*(4), 674–685.

Chernozhukov, V., Demirer, M., Duflo, E., & Fernandez-Val, I. (2018). *Generic machine learning inference on heterogeneous treatment effects in randomized experiments, with an application to immunization in india* (Tech. Rep.). National Bureau of Economic Research.

Choe, C., King, S., & Matsushima, N. (2018). Pricing with cookies: Behavior-based price discrimination

and spatial competition. *Management Science*, *64*(12), 5669–5687.

Chu, Y., Yang, H.-K., & Peng, W.-C. (2019). Predicting online user purchase behavior based on browsing history. In *2019 ieee 35th international conference on data engineering workshops (icdew)* (pp. 185–192).

Colombo, S. (2016). Imperfect behavior-based price discrimination. *Journal of Economics & Management Strategy*, *25*(3), 563–583.

Colombo, S. (2018). Behavior-and characteristic-based price discrimination. *Journal of Economics & Management Strategy*, *27*(2), 237–250.

Cowan, S. (2012). Third-degree price discrimination and consumer surplus. *The Journal of Industrial Economics*, *60*(2), 333–345.

De Cornière, A., Mantovani, A., & Shekhar, S. (2023). Third-degree price discrimination in two-sided markets.

Dubé, J.-P. (2019). Microeconometric models of consumer demand. In *Handbook of the economics of marketing* (Vol. 1, pp. 1–68). Elsevier.

Dubé, J.-P., & Misra, S. (2023). Personalized pricing and consumer welfare. *Journal of Political Economy*, *131*(1), 131–189.

Esteves, R. B. (2009). *A survey on the economics of behaviour-based price discrimination* (Tech. Rep.). NIPE-Universidade do Minho.

Esteves, R.-B. (2010). Pricing with customer recognition. *International Journal of Industrial Organization*, *28*(6), 669–681.

Esteves, R.-B. (2014). Price discrimination with private and imperfect information. *The Scandinavian Journal of Economics*, *116*(3), 766–796.

Esteves, R.-B., Liu, Q., & Shuai, J. (2022). Behavior-based price discrimination with nonuniform distribution of consumer preferences. *Journal of Economics & Management Strategy*, *31*(2), 324–355.

Esteves, R.-B., & Reggiani, C. (2014). Elasticity of demand and behaviour-based price discrimination. *International Journal of Industrial Organization*, *32*, 46–56.

Esteves, R. B., & Shuai, J. (2023). Behavior-based price discrimination with a general demand. *Available at SSRN 4403885*.

Freund, Y., & Schapire, R. E. (1997). A decision-theoretic generalization of on-line learning and an application to boosting. *Journal of computer and system sciences*, *55*(1), 119–139.

Friedman, J. H. (2002). Stochastic gradient boosting. *Computational statistics & data analysis*, *38*(4), 367–378.

Fudenberg, D., & Tirole, J. (2000). Customer poaching and brand switching. *RAND Journal of Economics*, 634–657.

Fudenberg, D., & Villas-Boas, J. M. (2006). Behavior-based price discrimination and customer recognition. *Handbook on economics and information systems*, *1*, 377–436.

Fudenberg, D., & Villas-Boas, J. M. (2012). In the digital economy. *The Oxford handbook of the digital economy*, 254.

Gehrig, T., Shy, O., & Stenbacka, R. (2012). A welfare evaluation of history-based price discrimination. *Journal of Industry, Competition and Trade*, *12*, 373–393.

Gordon, B. R., Zettelmeyer, F., Bhargava, N., & Chapsky, D. (2019). A comparison of approaches to advertising measurement: Evidence from big field experiments at facebook. *Marketing Science*,

38(2), 193–225.

Gosnell, G. K., List, J. A., & Metcalfe, R. D. (2020). The impact of management practices on employee productivity: A field experiment with airline captains. *Journal of Political Economy*, *128*(4), 1195–1233.

Han, S., & Kim, H. (2019). On the optimal size of candidate feature set in random forest. *Applied Sciences*, *9*(5), 898.

Hanemann, W. M. (1984). Discrete/continuous models of consumer demand. *Econometrica: Journal of the Econometric Society*, 541–561.

Hitsch, G. J., Hortacsu, A., & Lin, X. (2021). Prices and promotions in us retail markets. *Quantitative Marketing and Economics*, *19*(3), 289–368.

Hitsch, G. J., Misra, S., & Zhang, W. (2023). Heterogeneous treatment effects and optimal targeting policy evaluation. *Available at SSRN 3111957*.

Ho, T. K. (1995). Random decision forests. In *Proceedings of 3rd international conference on document analysis and recognition* (Vol. 1, pp. 278–282).

Horvitz, D. G., & Thompson, D. J. (1952). A generalization of sampling without replacement from a finite universe. *Journal of the American statistical Association*, *47*(260), 663–685.

Hospedales, T., Antoniou, A., Micaelli, P., & Storkey, A. (2021). Meta-learning in neural networks: A survey. *IEEE transactions on pattern analysis and machine intelligence*, *44*(9), 5149–5169.

Huber, M. (2014). Identifying causal mechanisms (primarily) based on inverse probability weighting. *Journal of Applied Econometrics*, *29*(6), 920–943.

Kallus, N., & Zhou, A. (2021). Fairness, welfare, and equity in personalized pricing. In *Proceedings of the 2021 acm conference on fairness, accountability, and transparency* (pp. 296–314).

Kim, C., Smith, A. N., Kim, J., & Allenby, G. M. (2023). Outside good utility and substitution patterns in direct utility models. *Journal of choice modelling*, *49*, 100447.

Kotsiantis, S. B. (2014). Bagging and boosting variants for handling classifications problems: a survey. *The Knowledge Engineering Review*, *29*(1), 78–100.

Lee, S., & Allenby, G. M. (2014). Modeling indivisible demand. *Marketing Science*, *33*(3), 364–381.

Liu, X. (2023). Dynamic coupon targeting using batch deep reinforcement learning: An application to livestream shopping. *Marketing Science*, *42*(4), 637–658.

Mahmood, A. (2014). How do customer characteristics impact behavior-based price discrimination? an experimental investigation. *Journal of Strategic Marketing*, *22*(6), 530–547.

Mienye, I. D., & Sun, Y. (2022). A survey of ensemble learning: Concepts, algorithms, applications, and prospects. *IEEE Access*, *10*, 99129–99149.

Miller, A., & Hosanagar, K. (2020). Personalized discount targeting with causal machine learning.

Opitz, D., & Maclin, R. (1999). Popular ensemble methods: An empirical study. *Journal of artificial intelligence research*, *11*, 169–198.

Oza, N. C., & Tumer, K. (2008). Classifier ensembles: Select real-world applications. *Information fusion*, *9*(1), 4–20.

Pazgal, A., & Soberman, D. (2008). Behavior-based discrimination: Is it a winning play, and if so, when? *Marketing Science*, *27*(6), 977–994.

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

Rubin, D. B. (2005). Causal inference using potential outcomes: Design, modeling, decisions. *Journal of the American Statistical Association*, *100*(469), 322–331.

Sagi, O., & Rokach, L. (2018). Ensemble learning: A survey. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, *8*(4), e1249.

Schonlau, M., & Zou, R. Y. (2020). The random forest algorithm for statistical learning. *The Stata Journal*, *20*(1), 3–29.

Seiler, S., & Yao, S. (2017). The impact of advertising along the conversion funnel. *Quantitative Marketing and Economics*, *15*(3), 241–278.

Shiller, B. R. (2020). Approximating purchase propensities and reservation prices from broad consumer tracking. *International Economic Review*, *61*(2), 847–870.

Shiller, B. R., et al. (2013). *First degree price discrimination using big data*. Brandeis Univ., Department of Economics.

Shin, J., & Sudhir, K. (2010). A customer management dilemma: When is it profitable to reward one's own customers? *Marketing Science*, *29*(4), 671–689.

Shobana, G., & Umamaheswari, K. (2021). Forecasting by machine learning techniques and econometrics: A review. In *2021 6th international conference on inventive computation technologies (icict)* (pp. 1010–1016).

Simester, D., Timoshenko, A., & Zoumpoulis, S. I. (2020). Efficiently evaluating targeting policies: Improving on champion vs. challenger experiments. *Management Science*, *66*(8), 3412–3424.

Smith, A. N., Seiler, S., & Aggarwal, I. (2023). Optimal price targeting. *Marketing Science*, *42*(3), 476–499.

Streitfeld, D. (2000). On the web price tags blur: What you pay could depend on who you are. *The Washington Post*(September 27).

Taunton, E. (2023). Mcdonald's accused of 'price-gouging' loyal customers. *Stuff* (August 5).

Tkáč, M., & Verner, R. (2016). Artificial neural networks in business: Two decades of research. *Applied Soft Computing*, *38*, 788–804.

Villas-Boas, J. M. (1999). Dynamic competition with customer recognition. *The Rand Journal of Economics*, 604–631.

Wang, C., Xu, S., & Yang, J. (2021). Adaboost algorithm in artificial intelligence for optimizing the iri prediction accuracy of asphalt concrete pavement. *Sensors*, *21*(17), 5682.

Wang, F., Li, Z., He, F., Wang, R., Yu, W., & Nie, F. (2019). Feature learning viewpoint of adaboost and a new algorithm. *IEEE Access*, *7*, 149890–149899.

Yang, Y., Lv, H., & Chen, N. (2023). A survey on ensemble learning under the era of deep learning. *Artificial Intelligence Review*, *56*(6), 5545–5589.

Yoganarasimhan, H., Barzegary, E., & Pani, A. (2020). Design and evaluation of personalized free trials. *arXiv preprint arXiv:2006.13420*.

Part

Online Appendix

Table of Contents

Table of Contents i

A Technical Appendix ii

A.1 Proof of Proposition 1 . . . ii

A.2 Derivation of log-Likelihood Functions . . . iii

A.3 Derivation of Demand Elasticity . . . iv

A.4 Proof of Proposition 2 . . . vi

B Data Description viii

B.1 Description of the Variables . . . viii

B.2 Summary Statistics . . . xi

C Omitted Tables xvii

D Omitted Figures xxiii

E Machine Learning and Algorithms xxxvii

E.1 Ensemble Machine Learning Methods . . . xxxvii

E.2 Ensemble Learning: Bagging, Random Forest, Adaptive Boosting, Gradient Boosting and Neural Network . . . xxxix

E.3 Estimation of Primitives of the Policy Model . . . xli

A Technical Appendix

A.1 Proof of Proposition 1

Demand equilibrium $\{q_i^{l^*}, q_i^{h^*}\}$ *is defined in the interval* $(\underline{B}_i^m, \overline{B}_i^m)$ *of the support of error* ϵ_m *distribution* $\forall m \in \{l, h\}$ *where:*

$$\underline{B}_i^m = \ln\left(\frac{\alpha_0(1-\mathbb{1}_{\{q_i^{m^*}\geq MPT_t\}}d_t)\gamma_m}{\alpha_m}\right) - \ln\left(\ln\left(\frac{\gamma_m q_i^{m^*}+1}{\gamma_m(q_i^{m^*}-\Delta)+1}\right)\right),$$
$$\overline{B}_i^m = \ln\left(\frac{\alpha_0(1-\mathbb{1}_{\{q_i^{m^*}\geq MPT_t\}}d_t)\gamma_m}{\alpha_m}\right) - \ln\left(\ln\left(\frac{\gamma_m(q_i^{m^*}+\Delta)+1}{\gamma_m q_i^{m^*}+1}\right)\right).$$

Proof: From (4.12) the following two expressions are derived:

$$\boldsymbol{\mathcal{U}}^*(q_i^{m^*}, q_i^{-m^*}) > \boldsymbol{\mathcal{U}}^*(q_i^{m^*}-\Delta, q_i^{-m^*}),$$
$$\frac{\alpha_m e^{\epsilon_m}}{\gamma_m}\ln(\gamma_m q_i^{m^*}+1) + \frac{\alpha_{-m}e^{\epsilon_{-m}}}{\gamma_{-m}}\ln(\gamma_{-m}q_i^{-m^*}+1) +$$
$$\alpha_0\Big(M_i - (1-\mathbb{1}_{\{q_i^{m^*}\geq MPT_t\}}d_t)q_i^{m^*} - (1-\mathbb{1}_{\{q_i^{-m^*}\geq MPT_t\}}d_t)q_i^{-m^*}\Big) >$$
$$\frac{\alpha_m e^{\epsilon_m}}{\gamma_m}\ln(\gamma_m(q_i^{m^*}-\Delta)+1) + \frac{\alpha_{-m}e^{\epsilon_{-m}}}{\gamma_{-m}}\ln(\gamma_{-m}q_i^{-m^*}+1) +$$
$$\alpha_0\Big(M_i - (1-\mathbb{1}_{\{q_i^{m^*}\geq MPT_t\}}d_t)(q_i^{m^*}-\Delta) - (1-\mathbb{1}_{\{q_i^{-m^*}\geq MPT_t\}}d_t)q_i^{-m^*}\Big),$$
$$\frac{\alpha_m e^{\epsilon_m}}{\gamma_m}\Big(\ln(\gamma_m q_i^{m^*}+1) - \ln(\gamma_m(q_i^{m^*}-\Delta)+1)\Big) > \alpha_0(1-\mathbb{1}_{\{q_i^{m^*}\geq MPT_t\}}d_t)\Delta,$$
$$e^{\epsilon_m}\ln\left(\frac{\gamma_m q_i^{m^*}+1}{\gamma_m(q_i^{m^*}-\Delta)+1}\right) > \frac{\alpha_0(1-\mathbb{1}_{\{q_i^{m^*}\geq MPT_t\}}d_t)\gamma_m\Delta}{\alpha_m},$$
$$\epsilon_m + \ln\left(\ln\left(\frac{\gamma_m q_i^{m^*}+1}{\gamma_m(q_i^{m^*}-\Delta)+1}\right)\right) > \ln\left(\frac{\alpha_0(1-\mathbb{1}_{\{q_i^{m^*}\geq MPT_t\}}d_t)\gamma_m\Delta}{\alpha_m}\right),$$
$$\epsilon_m > \ln\left(\frac{\alpha_0(1-\mathbb{1}_{\{q_i^{m^*}\geq MPT_t\}}d_t)\gamma_m\Delta}{\alpha_m}\right) - \ln\left(\ln\left(\frac{\gamma_m q_i^{m^*}+1}{\gamma_m(q_i^{m^*}-\Delta)+1}\right)\right).$$

$$\mathcal{U}^*(q_i^{m^*}, q_i^{-m^*}) > \mathcal{U}^*(q_i^{m^*} + \Delta, q_i^{-m^*}),$$

$$\frac{\alpha_m e^{\epsilon_m}}{\gamma_m} \ln(\gamma_m q_i^{m^*} + 1) + \frac{\alpha_{-m} e^{\epsilon_{-m}}}{\gamma_{-m}} \ln(\gamma_{-m} q_i^{-m^*} + 1) +$$
$$\alpha_0 \Big(M_i - (1 - \mathbb{1}_{\{q_i^{m^*} \geq MPT_t\}} d_t) q_i^{m^*} - (1 - \mathbb{1}_{\{q_i^{-m^*} \geq MPT_t\}} d_t) q_i^{-m^*} \Big) >$$
$$\frac{\alpha_m e^{\epsilon_m}}{\gamma_m} \ln(\gamma_m (q_i^{m^*} + \Delta) + 1) + \frac{\alpha_{-m} e^{\epsilon_{-m}}}{\gamma_{-m}} \ln(\gamma_{-m} q_i^{-m^*} + 1) +$$
$$\alpha_0 \Big(M_i - (1 - \mathbb{1}_{\{q_i^{m^*} \geq MPT_t\}} d_t)(q_i^{m^*} + \Delta) - (1 - \mathbb{1}_{\{q_i^{-m^*} \geq MPT_t\}} d_t) q_i^{-m^*} \Big),$$

$$\frac{\alpha_m e^{\epsilon_m}}{\gamma_m} \Big(\ln(\gamma_m q_i^{m^*} + 1) - \ln(\gamma_m (q_i^{m^*} + \Delta) + 1) \Big) > -\alpha_0 (1 - \mathbb{1}_{\{q_i^{m^*} \geq MPT_t\}} d_t) \Delta,$$

$$e^{\epsilon_m} \ln \left(\frac{\gamma_m (q_i^{m^*} + \Delta) + 1}{\gamma_m q_i^{m^*} + 1} \right) < \frac{\alpha_0 (1 - \mathbb{1}_{\{q_i^{m^*} \geq MPT_t\}} d_t) \gamma_m \Delta}{\alpha_m},$$

$$\epsilon_m + \ln \left(\ln \left(\frac{\gamma_m (q_i^{m^*} + \Delta) + 1}{\gamma_m q_i^{m^*} + 1} \right) \right) < \ln \left(\frac{\alpha_0 (1 - \mathbb{1}_{\{q_i^{m^*} \geq MPT_t\}} d_t) \gamma_m \Delta}{\alpha_m} \right),$$

$$\epsilon_m < \ln \left(\frac{\alpha_0 (1 - \mathbb{1}_{\{q_i^{m^*} \geq MPT_t\}} d_t) \gamma_m \Delta}{\alpha_m} \right) - \ln \left(\ln \left(\frac{\gamma_m (q_i^{m^*} + \Delta) + 1}{\gamma_m q_i^{m^*} + 1} \right) \right)$$

Therefore, $\underline{B}_i^m = \ln \left(\frac{\alpha_0 (1 - \mathbb{1}_{\{q_i^{m^*} \geq MPT_t\}} d_t) \gamma_m \Delta}{\alpha_m} \right) - \ln \left(\ln \left(\frac{\gamma_m q_i^{m^*} + 1}{\gamma_m (q_i^{m^*} - \Delta) + 1} \right) \right)$ with $\epsilon_m > \underline{B}_i^m$ and $\overline{B}_i^m = \ln \left(\frac{\alpha_0 (1 - \mathbb{1}_{\{q_i^{m^*} \geq MPT_t\}} d_t) \gamma_m \Delta}{\alpha_m} \right) - \ln \left(\ln \left(\frac{\gamma_m (q_i^{m^*} + \Delta) + 1}{\gamma_m q_i^{m^*} + 1} \right) \right)$ with $\epsilon_m < \overline{B}_i^m$. At last, it is proven that:

$$\underline{B}_i^m < \epsilon_m < \overline{B}_i^m . \blacksquare$$

A.2 Derivation of log-Likelihood Functions

Through logarithmic transformation of $\mathcal{L}(\Theta)$ in Equation 4.14, multinomial log-likelihood function $\ell(\Theta)$ is derived and specified for each combinations of (d_t, MPT_t), namely:

$$\ell_C(\Theta | q_i^{l^*}, d_t = 0, MPT_t = 0) = \sum_{i=1}^{I} \Bigg\{ \mathbb{1}_{q_i^{l^*}=0} \underbrace{\ln \left[\Phi \Big(\frac{\overline{B}_i^l |_{q_i^{l^*}=0} - \mu_l}{\sigma_l} \Big) \right]}_{\text{No Convertion}} +$$
$$\mathbb{1}_{q_i^{l^*}>0} \underbrace{\ln \left[\Phi \Big(\frac{\overline{B}_i^l |_{q_i^{l^*}>0} - \mu_l}{\sigma_l} \Big) - \Phi \Big(\frac{\underline{B}_i^l |_{q_i^{l^*}>0} - \mu_l}{\sigma_l} \Big) \right]}_{\text{Low Convertion}} \Bigg\}, \tag{A.2}$$

$$\ell_{ST}(\Theta|q_i^{h^*}, d_t > 0, MPT_t = 0) = \sum_{i=1}^{I}\Bigg\{\mathbb{1}_{q_i^{h^*}=0}\underbrace{\ln\left[\Phi\left(\frac{\overline{B}_i^h|_{q_i^{h^*}=0}-\mu_h}{\sigma_h}\right)\right]}_{\text{No Convertion}}+$$
$$\mathbb{1}_{q_i^{h^*}>0}\underbrace{\ln\left[\Phi\left(\frac{\overline{B}_i^h|_{q_i^{h^*}>0}-\mu_h}{\sigma_h}\right)-\Phi\left(\frac{\underline{B}_i^h|_{q_i^{h^*}>0}-\mu_h}{\sigma_h}\right)\right]}_{\text{High Convertion}}\Bigg\}, \tag{A.3}$$

$$\ell_{FT}(\Theta|q_i^{l^*}, q_i^{h^*}, d_t > 0, MPT_t > 0) = \sum_{i=1}^{I}\Bigg\{\mathbb{1}_{q_i^{l^*}=0}\underbrace{\ln\left[\Phi\left(\frac{\overline{B}_i^l|_{q_i^{l^*}=0}-\mu_l}{\sigma_l}\right)\right]+\mathbb{1}_{q_i^{h^*}0}\ln\left[\Phi\left(\frac{\overline{B}_i^h|_{q_i^{h^*}=0}-\mu_h}{\sigma_h}\right)\right]}_{\text{No Convertion}}+$$
$$\mathbb{1}_{q_i^{l^*}>0}\underbrace{\ln\left[\Phi\left(\frac{\overline{B}_i^l|_{q_i^{l^*}>0}-\mu_l}{\sigma_l}\right)-\Phi\left(\frac{\underline{B}_i^l|_{q_i^{l^*}>0}-\mu_l}{\sigma_l}\right)\right]}_{\text{Low Convertion}}+$$
$$\mathbb{1}_{q_i^{h^*}>0}\underbrace{\ln\left[\Phi\left(\frac{\overline{B}_i^h|_{q_i^{h^*}>0}-\mu_h}{\sigma_h}\right)-\Phi\left(\frac{\underline{B}_i^h|_{q_i^{h^*}>0}-\mu_h}{\sigma_h}\right)\right]}_{\text{High Convertion}}\Bigg\}, \tag{A.4}$$

where the function $\ell(\Theta)$ is computed over the sessions of I users. Equation A.2 is log-likelihood function of control treatment $t = C$ (Figure D.7a), Equation A.3 of semi-treatment $t = ST$ (Figure D.7b) and Equation A.4 of full-treatment $t = FT$ (Figure D.7c).

A.3 Derivation of Demand Elasticity

Proof: Marginal rate of substitution of utility function in (4.3) is

$$MRS_{q_i^{l(h)}x_i^0} = \frac{\frac{\partial \mathcal{U}(q_i^l, q_i^h, x_i^0)}{\partial q_i^{l(h)}}}{\frac{\partial \mathcal{U}(q_i^l, q_i^h, x_i^0)}{\partial x_i^0}} = \frac{\alpha_{l(h)} e^{\epsilon_{l(h)}}}{(\gamma_{l(h)} q_i^{l(h)} + 1)\alpha_0},$$

where $\frac{\partial \mathcal{U}(q_i^l, q_i^h, x_i^0)}{\partial q_i^{l(h)}} = \frac{\alpha_{l(h)} e^{\epsilon_{l(h)}}}{\gamma_{l(h)} q_i^{l(h)}+1}$ and $\frac{\partial \mathcal{U}(q_i^l, q_i^h, x_i^0)}{\partial x_i^0} = \alpha_0$. By Walras's law the optimal demand lies on the budget constraint line. At the point of tangency, the marginal rate of substitution between the two goods

$q_i^{l(h)}$ and x_i^0 is equal to their price ratio:

$$MRS_{q_i^{l(h)}x_i^0} = 1 - \mathbb{1}_{\{q_i^{l(h)} \geq MPT_t\}} d_t$$
$$\frac{\alpha_{l(h)} e^{\epsilon_{l(h)}}}{(\gamma_{l(h)} q_i^{l(h)} + 1)\alpha_0} = 1 - \mathbb{1}_{\{q_i^{l(h)} \geq MPT_t\}} d_t$$
$$q_i^{l(h)} = \frac{\alpha_{l(h)} e^{\epsilon_{l(h)}} - (1 - \mathbb{1}_{\{q_i^{l(h)} \geq MPT_t\}} d_t)\alpha_0}{(1 - \mathbb{1}_{\{q_i^{l(h)} \geq MPT_t\}} d_t)\alpha_0 \gamma_{l(h)}}.$$

From the budget constraint $(1 - \mathbb{1}_{\{q_i^{l(h)} \geq MPT_t\}} d_t) q_i^{l(h)} + x_i^0 = M_i$ it follows that:

$$x_i^0 = M_i - (1 - \mathbb{1}_{\{q_i^{l(h)} \geq MPT_t\}} d_t) q_i^{l(h)}$$
$$x_i^0 = M_i - (1 - \mathbb{1}_{\{q_i^{l(h)} \geq MPT_t\}} d_t) \frac{\alpha_{l(h)} e^{\epsilon_{l(h)}} - (1 - \mathbb{1}_{\{q_i^{l(h)} \geq MPT_t\}} d_t)\alpha_0}{(1 - \mathbb{1}_{\{q_i^{l(h)} \geq MPT_t\}} d_t)\alpha_0 \gamma_{l(h)}}.$$

Price elasticity $\varepsilon_{l(h)}$ of good $l(h)$ is derived:

$$\varepsilon_{l(h)} = \frac{\frac{\partial q_i^{l(h)}}{q_i^{l(h)}}}{\frac{\partial 1 - \mathbb{1}_{\{q_i^{l(h)} \geq MPT_t\}} d_t}{1 - \mathbb{1}_{\{q_i^{l(h)} \geq MPT_t\}} d_t}}$$
$$= \frac{\partial q_i^{l(h)}}{\partial 1 - \mathbb{1}_{\{q_i^{l(h)} \geq MPT_t\}} d_t} \frac{1 - \mathbb{1}_{\{q_i^{l(h)} \geq MPT_t\}} d_t}{q_i^{l(h)}}$$
$$= \frac{-\alpha_0 (1 - \mathbb{1}_{\{q_i^{l(h)} \geq MPT_t\}} d_t)\alpha_0 \gamma_{l(h)} - (\alpha_{l(h)} e^{\epsilon_{l(h)}} - (1 - \mathbb{1}_{\{q_i^{l(h)} \geq MPT_t\}} d_t)\alpha_0)\alpha_0 \gamma_{l(h)}}{((1 - \mathbb{1}_{\{q_i^{l(h)} \geq MPT_t\}} d_t)\alpha_0 \gamma_{l(h)})^2} \frac{1 - \mathbb{1}_{\{q_i^{l(h)} \geq MPT_t\}} d_t}{\frac{\alpha_{l(h)} e^{\epsilon_{l(h)}} - (1 - \mathbb{1}_{\{q_i^{l(h)} \geq MPT_t\}} d_t)\alpha_0}{(1 - \mathbb{1}_{\{q_i^{l(h)} \geq MPT_t\}} d_t)\alpha_0 \gamma_{l(h)}}}$$

$$= \frac{-\alpha_{l(h)} e^{\epsilon_{l(h)}} \alpha_0 \gamma_{l(h)}}{((1 - \mathbb{1}_{\{q_i^{l(h)} \geq MPT_t\}} d_t)\alpha_0 \gamma_{l(h)})^2} \frac{(1 - \mathbb{1}_{\{q_i^{l(h)} \geq MPT_t\}} d_t)^2 \alpha_0 \gamma_{l(h)}}{\alpha_{l(h)} e^{\epsilon_{l(h)}} - (1 - \mathbb{1}_{\{q_i^{l(h)} \geq MPT_t\}} d_t)\alpha_0}$$
$$= -\frac{\alpha_{l(h)} e^{\epsilon_{l(h)}}}{\alpha_{l(h)} e^{\epsilon_{l(h)}} - (1 - \mathbb{1}_{\{q_i^{l(h)} \geq MPT_t\}} d_t)\alpha_0}.$$

Cross-price elasticity ε_0 of outside good with respect to good $l(h)$ is derived:

$$\begin{aligned}
\varepsilon_0 &= \frac{\frac{\partial x_i^0}{x_i^0}}{\frac{\partial 1-\mathbb{1}_{\{q_i^{l(h)}\geq MPT_t\}}d_t}{1-\mathbb{1}_{\{q_i^{l(h)}\geq MPT_t\}}d_t}} \\
&= \frac{\partial x_i^0}{\partial 1-\mathbb{1}_{\{q_i^{l(h)}\geq MPT_t\}}d_t}\frac{1-\mathbb{1}_{\{q_i^{l(h)}\geq MPT_t\}}d_t}{x_i^0} \\
&= \frac{1}{\gamma_{l(h)}}\frac{1-\mathbb{1}_{\{q_i^{l(h)}\geq MPT_t\}}d_t}{M_i-(1-\mathbb{1}_{\{q_i^{l(h)}\geq MPT_t\}}d_t)\frac{\alpha_{l(h)}e^{\epsilon_{l(h)}}-(1-\mathbb{1}_{\{q_i^{l(h)}\geq MPT_t\}}d_t)\alpha_0}{(1-\mathbb{1}_{\{q_i^{l(h)}\geq MPT_t\}}d_t)\alpha_0\gamma_{l(h)}}} \\
&= \frac{(1-\mathbb{1}_{\{q_i^{l(h)}\geq MPT_t\}}d_t)\alpha_0}{M_i\alpha_0\gamma_{l(h)}-\alpha_{l(h)}e^{\epsilon_{l(h)}}+(1-\mathbb{1}_{\{q_i^{l(h)}\geq MPT_t\}}d_t)\alpha_0}.
\end{aligned}$$

Therefore, price elasticity and cross-price elasticity are respectively:

$$\varepsilon_{l(h)} = -\frac{\alpha_{l(h)}e^{\epsilon_{l(h)}}}{\alpha_{l(h)}e^{\epsilon_{l(h)}}-(1-\mathbb{1}_{\{q_i^{l(h)}\geq MPT_t\}}d_t)\alpha_0}, \tag{A.5}$$

$$\varepsilon_0 = \frac{(1-\mathbb{1}_{\{q_i^{l(h)}\geq MPT_t\}}d_t)\alpha_0}{M_i\alpha_0\gamma_{l(h)}-\alpha_{l(h)}e^{\epsilon_{l(h)}}+(1-\mathbb{1}_{\{q_i^{l(h)}\geq MPT_t\}}d_t)\alpha_0}. \tag{A.6}$$

A.4 Proof of Proposition 2

Optimal treatment policy is defined as

$$f_{it}(x) = \Delta\mu_{it}(x)\Delta q_{it}(x) + q_{iC}(x)\Delta\mu_{it}(x) + \mu_{iC}(x)\Delta q_{it}(x) - \mu_{it}(x)\gamma_{it}(x)q_{it}(x)d_t.$$

Proof: From (5.7) the following expression is derived:

$$\mathbb{E}\Big[\Pi_i \mid T_i = t, X_i = x\Big] > \mathbb{E}\Big[\Pi_i \mid T_i = C, X_i = x\Big],$$

$$\mu_{it}(x)\gamma_{it}(x)\Big(q_{it}(x)(1-d_t)-c\Big) + \mu_{it}(x)\Big(1-\gamma_{it}(x)\Big)\Big(q_{it}(x)-c\Big) > \mu_{iC}(x)q_{iC}(x).$$

Marginal cost c is assumed equal to zero. With treatment set $\mathcal{T} = \{C, SF, FT\}$, two elements, i.e. $t = ST$ and $t = FT$, belong to the the set $\mathcal{T}_{-\{C\}}$ of Equation 5.10. By definition $\gamma_{iC}(x) = \gamma_{iST}(x) = 1$. Thus, two cases follow:

(i) Semi-treatment:

$$\mu_{iST}(x)q_{iST}(x)(1-d_{ST}) > \mu_{iC}(x)q_{iC}(x),$$

$$\mu_{iST}(x)q_{iST}(x) - \mu_{iC}(x)q_{iC}(x) > \mu_{iST}(x)q_{iST}(x)d_{ST}.$$

By adding $\Big(-\mu_{iST}(x)q_{iC}(x) + \mu_{iST}(x)q_{iC}(x)\Big)$ and $\Big(+\mu_{iC}(x)q_{iST}(x) - \mu_{iC}(x)q_{iST}(x)\Big)$ and $\Big(-\mu_{iC}(x)q_{iC}(x) +$

$\mu_{iC}(x)q_{iC}(x)\Big)$ on the left-hand side, inequality can be rewritten as:

$$\mu_{iST}(x)q_{iST}(x) + \Big(- \mu_{iST}(x)q_{iC}(x) + \mu_{iST}(x)q_{iC}(x)\Big) - \mu_{iC}(x)q_{iC}(x) + \Big(\mu_{iC}(x)q_{iST}(x)$$
$$-\mu_{iC}(x)q_{iST}(x)\Big) + \Big(- \mu_{iC}(x)q_{iC}(x) + \mu_{iC}(x)q_{iC}(x)\Big)$$
$$> \mu_{iST}(x)q_{iST}(x)d_{ST},$$
$$\mu_{iST}(x)\Big(q_{iST}(x) - q_{iC}(x)\Big) - \mu_{iC}(x)\Big(q_{iST}(x) - q_{iC}(x)\Big) - 2\mu_{iC}(x)q_{iC}(x) + \mu_{iC}(x)q_{iST}(x) + \mu_{iST}(x)q_{iC}(x)$$
$$> \mu_{iST}(x)q_{iST}(x)d_{ST},$$
$$\mu_{iST}(x)\Delta q_{iST}(x) - \mu_{iC}(x)\Delta q_{iST}(x) - 2\mu_{iC}(x)q_{iC}(x) + \mu_{iC}(x)q_{iST}(x) + \mu_{iST}(x)q_{iC}(x)$$
$$> \mu_{iST}(x)q_{iST}(x)d_{ST},$$
$$\Big(\mu_{iST}(x) - \mu_{iC}(x)\Big)\Delta q_{iST}(x) + \mu_{iC}(x)\Big(q_{iST}(x) - q_{iC}(x)\Big) + \Big(\mu_{iST}(x) - \mu_{iC}(x)\Big)q_{iC}(x)$$
$$> \mu_{iST}(x)q_{iST}(x)d_{ST},$$
$$\Delta\mu_{iST}(x)\Delta q_{iST}(x) + \mu_{iC}(x)\Delta q_{iST}(x) + \Delta\mu_{iST}(x)q_{iC}(x)$$
$$> \mu_{iST}(x)q_{iST}(x)d_{ST},$$
$$\Delta\mu_{iST}(x)\Delta q_{iST}(x) + \mu_{iC}(x)\Delta q_{iST}(x) + q_{iC}(x)\Delta\mu_{iST}(x) - \mu_{iST}(x)q_{iST}(x)d_{ST} > 0.$$

Therefore, for $t = ST$, the optimal treatment function is

$$f_{iST}(x) = \Delta\mu_{iST}(x)\Delta q_{iST}(x) + \mu_{iC}(x)\Delta q_{iST}(x) + q_{iC}(x)\Delta\mu_{iST}(x) - \mu_{iST}(x)q_{iST}(x)d_{ST}.$$

(ii) Full-treatment:

$$\mu_{iFT}(x)\gamma_{iFT}(x)q_{iFT}(x)(1 - d_{FT}) + \mu_{iFT}(x)(1 - \gamma_{iFT}(x))q_{iFT}(x) > \mu_{iC}(x)q_{iC}(x),$$

$$\mu_{iFT}(x)\gamma_{iFT}(x)q_{iFT}(x) - \mu_{iFT}(x)\gamma_{iFT}(x)q_{iFT}(x)d_{FT} + \mu_{iFT}(x)q_{iFT}(x) - \mu_{iFT}(x)\gamma_{iFT}(x)q_{iFT}(x) > \mu_{iC}(x)q_{iC}(x),$$
$$\mu_{iFT}(x)q_{iFT}(x) - \mu_{iC}(x)q_{iC}(x) > \mu_{iFT}(x)\gamma_{iFT}(x)q_{iFT}(x)d_{FT}.$$

The remaining part of the proof resembles the case of $t = ST$. Therefore, for $t = FT$, the optimal treatment function is

$$f_{iFT}(x) = \Delta\mu_{iFT}(x)\Delta q_{iFT}(x) + \mu_{iC}(x)\Delta q_{iFT}(x) + q_{iC}(x)\Delta\mu_{iFT}(x) - \mu_{iFT}(x)\gamma_{iFT}(x)q_{iFT}(x)d_{FT}.\blacksquare$$

B Data Description

The analysis in Sections 3.1, 3.2, 4.1, 4.2, 5.2, 6.1, 6.2 and in Online Appendix explicitly refers to selected variables for clarity of explanation and due to space limitations. The variables are separated into four groups based on their content and are reported with their group names due to space limitations: (i) main variables, (ii) technology controls, (iii) time controls, and (iv) location controls.

B.1 Description of the Variables

1. **Main variables**: These variables are at the core of investigation of this paper. They are employed in the reduced-form analysis of Sections 3.1 and 4.1, and also in the causal analysis of Section 3.2, and in the estimation of the structural demand model of Section 4.2 and of nonparametric policy model in Section 5.2. Depending on the specification used, the variables of interest are directly mentioned. Summary statistics are reported in Table B.1 at the aggregate level, and for control, semi-treatment, and full-treatment groups. Table B.2 reports the correlations between cross-correlation coefficients of continuous, main variables at the aggregate level. This list follows the chronological order of the conversion funnel: (i) characteristics of session, (ii) characteristics of treatment, (iii) characteristics of cart, (iv) characteristics of order, and (v) characteristics of user history.

   (a) *Time to session* measures the length of each session in minutes. It is computed by comparing the day and hour reported by *Session start* and *Session end* for the 5,087,685 observations. Mean is 13 minutes with standard deviation of 17.64 minutes.

   (b) *Number of pages* counts the number of e-commerce website pages visualized by the user in each session. This variable is a proxy of user's engagement in the conversion funnel. It considers all the 5,087,685 sessions with an average of 14.3 pages viewed per session, ranging from a minimum of 1 to a maximum of 120 pages.

   (c) The dummy *Treatment* is equal to 1 if the session is assigned to the treatment group (either semi-treatment or full-treatment), 0 otherwise. 4,246,413 sessions are assigned to the treatment group (2,490,092 to semi-treatment, 1,756,351 to full-treatment), 841,242 to the control group. In relative terms, propensity scores of control, semi-treatment, and full-treatment are 0.15, 0.5 and 0.35, respectively.

   (d) The variable d reports the RCT percentage discount assigned to the treated session which is visualized by the user if the cart page is attained. Depending on the dummy *Treatment*, $d \in \{0, 5, 10, 15, 20\}$.

   (e) The variable *MPT* reports the RCT minimum purchase threshold on gross order q to benefit the discount, which it is visualized by the user, combined with the discount percentage d, if the cart page is attained. Unit of measure is €. For all the sessions reporting different currencies, a conversion is made using the average exchange rate of the day the session occurred. Depending on the dummy *Treatment*, $MPT \in \{0, 15, 25, 35, 45\}$.

(f) The dummy *Treatment compliance* is equal to 1 if gross order q is equal to or greater than minimum purchase threshold *MPT*. *Treatment compliance* is 1 for the sessions assigned to control and semi-treatment groups.

(g) The dummy *Attained cart* is equal to 1 if the user visualizes the cart page, 0 otherwise. On average, between 26% and 30% of all sessions reach the cart page.

(h) *Time to cart* counts the minutes each user takes to visualize the cart for the first time. This variable is a proxy of user's interest in the conversion funnel. It records 1,456,608 sessions, ranging from a minimum of 1 to a maximum of 338 minutes (5 hours and 38 minutes). Mean is 10 minutes with standard deviation of 11.90 minutes.

(i) The dummy *Converted order* is equal to 1 if the user converts the session, 0 otherwise. Conversion rate of the sample is 5%.

(j) *Time to order* counts the minutes each user needs to attain the checkout order page. This variable is a proxy of user's commitment in the conversion funnel. It records 271,918 sessions, ranging from a minimum of 3 to a maximum of 446 minutes (7 hours and 26 minutes). Mean is 30 minutes with standard deviation of 24.83 minutes.

(k) *Number of products* counts the number of items purchased in each converted session. This variable records 271,918 sessions with mean of 5 and standard deviation of 3.43.

(l) q measures the gross order of each converted session. Unit of measure is €. For all the sessions reporting different currencies, a conversion is made using the average exchange rate of the day the session occurred. This variable counts 271,918 sessions with an average of €45 per converted session.

(m) The dummy *New user* is equal to 1 if the user never visited the e-commerce website before, 0 otherwise. On average, between 73% and 77% of all users are new users.

(n) The dummy *Purchase history* is equal to 1 if the user has already purchased at least once from the e-commerce website, 0 otherwise. Rate of users with previous purchase records is 3%.

(o) *Number of sessions* counts the number of sessions per user key in the RCT experiment time. This variable measures the frequency of sessions per user and it is a proxy of users' interaction on the e-commerce website. It considers all the 5,087,685 sessions with an average of 1.7 sessions per user, ranging from 1.16 in the control group, to 1.96 and 1.78 in the semi-treatment group and full-treatment group.

(p) *Number of prior visits* measures the number of sessions per user before the RCT session. This variable is a proxy of user purchase history, and it considers 1,923,779 sessions.

(q) *Number of prior purchases* measures the number of converted sessions per user before the RCT session. This variable is a proxy of user purchase history and it considers 1,923,779 sessions.

(r) *Number of prior abandoned carts* measures the number of abandoned carts per user before the RCT session. This variable is a proxy of user purchase history and it considers 638,710 sessions.

2. **Technology controls**: These variables report information related to device typologies, browsing technologies and operative systems used in each session. They are described in Table B.3. Technology variables include the following controls to check for e-commerce heterogeneity:

   (a) Device (decreasing frequency order): Mobile, Computer, and Tablet. . Dummies for device: *Mobile*, *Computer*, and *Tablet* (2 variables).

   (b) Browser (decreasing frequency order with at least 1,000 sessions): Chrome Mobile, Mobile Safari, Chrome, Apple WebKit, Safari, Firefox, Firefox Mobile, Microsoft Edge Mobile, Firefox Mobile iOS, and Mozilla. Dummies for browser: *Chrome Mobile*, *Mobile Safari*, *Chrome*, *Apple WebKit*, *Safari*, *Firefox*, *Firefox Mobile*, and *Others* (7 variables).

   (c) Operative system (decreasing frequency order with at least 1,000 sessions): Mac OS X (iPhone), Android 1.x, Windows 10, Android Mobile, Mac OS X, Mac OS X (iPad), Windows 7, Android, Chrome OS, Windows 8.1, Linux, Android 6.x, Android 5.x, and Windows 8. Dummies for operative system: *Mac OS X (iPhone)*, *Android 1.x*, *Windows 10*, *Android Mobile*, *Mac OS X*, *Mac OS X (iPad)*, *Windows 7*, and *Others* (7 variables).

3. **Time controls**: This is a group of additional time covariates which are inferred by comparing timing information from both *Session start* and *Session end*. These variables are related to month, day and hour of each session. In the case where a session overlaps, so that *Session start* and *Session end* would belong to different time controls, then the end of the session is the reference term of month, day, and hour. They are described in Table B.4. Time variables include the following controls to address time trends and seasonality:

   (a) Control of the month in which the session occurred. Dummies for month: *January*, *February*, *March*, *April*, *May*, *June*, *July*, *August*, *September*, *October*, *November*, and *December* (11 variables).

   (b) Control of the time of the month in which the session took place. Dummies for time of the month: *Beginning* (1st-10th), *Halfway* (11th-20th), and *End* (21st-28th/30th/31st) (2 variables).

   (c) Control of the time of the day in which the session ends. Dummies for time of the day: *Night* (12:00AM-6:00AM), *Morning* (6:00AM-12:00PM), *Afternoon* (12:00PM-6:00PM), and *Evening* (6:00PM-12:00AM) (3 variables).

4. **Location controls**: These variables include the country and the region in which the sessions took place, and their related currency. They are described in Table B.5. Location variables include the following controls to account for nationality and currency heterogeneity:

   (a) Country (decreasing frequency order with at least 10,000 sessions): France, Italy, Spain, Portugal, United Kingdom, Denmark, Netherlands, Belgium, Ukraine, United States, Brazil, Egypt, Poland, Romania, Switzerland, Morocco, Thailand, Ireland, Russia, Greece, Croatia, Canada, Lithuania, South Africa, Algeria, and Latvia. Dummies for countries: *France*, *Italy*, *Spain*, *Portugal*, *United Kingdom*, *Denmark*, *Netherlands*, *Belgium*, *Ukraine*, *United States*, and *Others* (10 variables).

(b) Region: Since each country has its own specific administrative division into states and/or regions (either numerical, nominal or by acronym), it results challenging to code all the dummy variables to account for such a plethora. Thus, region is an information available, but it is not considered in the analysis.

(c) Currency: Euro (EUR), British pound sterling (GBP), and United States dollar (USD). For all sessions recorded in a country where the official currency is neither Euro, nor British pound sterling, nor United States dollar, the local currency is converted into United States dollar at the exchange rate of the day the session took place. Dummies for currency: *EUR*, *GBP*, and *USD* (2 variables).

B.2 Summary Statistics

Table B.1: Summary statistics of main variables

Variable	Aggregate			Control			Semi-Treatment			Full-Treatment		
	Obs.	Mean	St.Dev.	Obs.	Mean	St.Dev.	Obs.	Mean	St.Dev.	Obs.	Mean	St.Dev.
Characteristics of session												
Time to session (minute)	5,087,685	12.96	17.64	841,242	12.04	17.64	2,490,092	12.39	17.02	1,756,351	14.19	18.42
Number of pages	5,087,685	14.31	15.95	841,242	12.59	15.09	2,490,092	15.06	15.66	1,756,351	14.07	16.66
Characteristics of treatment												
Treatment (0/1)	5,087,685	0.85	0.52	841,242	0.00	0.00	2,490,092	1.00	0.00	1,756,351	1.00	0.00
d (%)	5,087,685	9.59	5.01	841,242	0.00	0.00	2,490,092	12.46	3.36	1,756,351	10.13	0.91
MPT (€)	5,087,685	10.78	15.54	841,242	0.00	0.00	2,490,092	0.00	0.00	1,756,351	31.21	7.82
Treatment compliance (0/1)	5,087,685	0.05	0.21	841,242	1.00	0.00	2,490,092	1.00	0.00	1,756,351	0.04	0.19
Characteristics of cart												
Attained cart (0/1)	5,087,685	0.29	0.45	841,242	0.26	0.44	2,490,092	0.30	0.46	1,756,351	0.28	0.45
Time to cart (minute)	1,456,608	9.83	11.90	218,753	9.64	12.49	749,093	8.74	10.82	488,762	11.59	12.94
Characteristics of order												
Converted order (0/1)	5,087,685	0.05	0.22	841,242	0.05	0.22	2,490,092	0.06	0.23	1,756,351	0.05	0.22
Time to order (minute)	271,918	29.89	24.83	41,359	30.43	26.05	140,046	27.36	23.25	90,513	33.56	26.10
Number of products	271,918	5.13	3.43	41,359	5.01	3.32	140,046	4.69	3.28	90,513	5.86	3.59
q (€)	271,918	45.42	65.41	41,359	46.49	67.25	140,046	46.43	77.69	90,513	43.35	37.96
Characteristics of user history												
New user (0/1)	5,087,685	0.74	0.48	841,242	0.77	0.49	2,490,092	0.73	0.48	1,756,351	0.75	0.49
Purchase history (0/1)	5,087,685	0.03	0.17	841,242	0.03	0.18	2,490,092	0.03	0.18	1,756,351	0.02	0.15
Number of sessions	5,087,685	1.77	3.36	841,242	1.16	0.73	2,490,092	1.96	4.13	1,756,351	1.78	2.82
Number of prior visits	1,923,779	8.89	22.12	330,597	9.11	23.25	928,402	8.64	22.16	664,780	9.12	21.46
Number of prior purchases	1,923,779	0.18	1.25	330,597	0.18	1.20	928,402	0.20	1.38	664,780	0.14	1.08
Number of prior abandoned carts	638,710	3.78	8.51	110,368	3.92	9.00	312,792	3.90	9.02	215,550	3.53	7.43

Table B.2: Correlation matrix for continuous variables

	Time to session	Number of pages	d	MPT	Time to cart	Time to order	Number of products	q	Number of sessions	Number of prior visits	Number of prior purchases	Number of prior abandoned carts
Time to session	1.000											
Number of pages	0.591	1.000										
d	0.005	0.054	1.000									
MPT	0.000	-0.055	0.077	1.000								
Time to cart	0.493	0.429	0.006	0.013	1.000							
Time to order	0.386	0.360	0.009	0.002	0.327	1.000						
Number of products	0.245	0.299	0.015	0.006	0.251	0.744	1.000					
q	0.172	0.197	0.021	-0.012	0.155	0.483	0.611	1.000				
Number of sessions	0.053	0.032	0.107	0.015	0.009	0.018	0.009	0.173	1.000			
Number of prior visits	0.045	0.014	0.010	0.000	0.010	0.010	0.013	0.027	0.457	1.000		
Number of prior purchases	0.041	0.033	0.021	-0.014	0.032	0.052	0.067	0.352	0.101	0.431	1.000	
Number of prior abandoned carts	0.054	0.041	0.013	-0.014	0.037	0.027	0.032	0.035	0.311	0.650	0.416	1.000

Table B.3: Summary statistics of technology controls

	Aggregate			Control			Semi-Treatment			Full-Treatment		
	Obs.	Mean	St.Dev.	Obs.	Mean	St.Dev.	Obs.	Mean	St.Dev.	Obs.	Mean	St.Dev.
Controls of device												
Mobile (0/1)	5,087,685	0.87	0.34	841,242	0.79	0.41	2,490,092	0.89	0.32	1,756,351	0.87	0.33
Computer (0/1)	5,087,685	0.13	0.33	841,242	0.21	0.40	2,490,092	0.11	0.31	1,756,351	0.12	0.32
Tablet (0/1)	5,087,685	0.01	0.08	841,242	0.01	0.08	2,490,092	0.01	0.08	1,756,351	0.01	0.09
Controls of browser												
Chrome Mobile (0/1)	5,087,685	0.44	0.50	841,242	0.40	0.49	2,490,092	0.42	0.49	1,756,351	0.48	0.50
Mobile Safari (0/1)	5,087,685	0.36	0.48	841,242	0.33	0.47	2,490,092	0.41	0.49	1,756,351	0.31	0.46
Chrome (0/1)	5,087,685	0.10	0.29	841,242	0.15	0.36	2,490,092	0.08	0.27	1,756,351	0.09	0.29
Apple WebKit (0/1)	5,087,685	0.06	0.25	841,242	0.05	0.22	2,490,092	0.06	0.23	1,756,351	0.08	0.27
Safari (0/1)	5,087,685	0.03	0.17	841,242	0.05	0.21	2,490,092	0.02	0.15	1,756,351	0.03	0.16
Firefox (0/1)	5,087,685	0.01	0.09	841,242	0.01	0.11	2,490,092	0.01	0.08	1,756,351	0.01	0.08
Firefox Mobile (0/1)	5,087,685	0.00	0.03	841,242	0.00	0.03	2,490,092	0.00	0.03	1,756,351	0.00	0.03
Other (0/1)	5,087,685	0.00	0.03	841,242	0.00	0.03	2,490,092	0.00	0.03	1,756,351	0.00	0.03
Controls of operative system												
Mac OS X (iPhone) (0/1)	5,087,685	0.44	0.50	841,242	0.40	0.49	2,490,092	0.48	0.50	1,756,351	0.41	0.49
Android 1.x (0/1)	5,087,685	0.37	0.48	841,242	0.34	0.48	2,490,092	0.36	0.48	1,756,351	0.40	0.49
Windows 10 (0/1)	5,087,685	0.08	0.27	841,242	0.13	0.33	2,490,092	0.06	0.25	1,756,351	0.07	0.26
Android Mobile (0/1)	5,087,685	0.05	0.21	841,242	0.04	0.20	2,490,092	0.04	0.19	1,756,351	0.07	0.25
Mac OS X (0/1)	5,087,685	0.04	0.20	841,242	0.06	0.25	2,490,092	0.03	0.18	1,756,351	0.04	0.19
Mac OS X (iPad) (0/1)	5,087,685	0.01	0.08	841,242	0.01	0.07	2,490,092	0.01	0.08	1,756,351	0.01	0.08
Windows 7 (0/1)	5,087,685	0.00	0.06	841,242	0.01	0.08	2,490,092	0.00	0.05	1,756,351	0.00	0.00
Other (0/1)	5,087,685	0.01	0.10	841,242	0.01	0.12	2,490,092	0.01	0.09	1,756,351	0.01	0.10

Table B.4: Summary statistics of time controls

	Aggregate			Control			Semi-Treatment			Full-Treatment		
	Obs.	Mean	St.Dev.	Obs.	Mean	St.Dev.	Obs.	Mean	St.Dev.	Obs.	Mean	St.Dev.
Controls of month												
January (0/1)	5,087,685	0.22	0.41	841,242	0.19	0.39	2,490,092	0.19	0.39	1,756,351	0.26	0.44
February (0/1)	5,087,685	0.08	0.26	841,242	0.06	0.23	2,490,092	0.00	0.00	1,756,351	0.19	0.39
March (0/1)	5,087,685	0.05	0.22	841,242	0.09	0.29	2,490,092	0.00	0.02	1,756,351	0.10	0.30
April (0/1)	5,087,685	0.06	0.24	841,242	0.04	0.20	2,490,092	0.05	0.22	1,756,351	0.09	0.28
May (0/1)	5,087,685	0.05	0.22	841,242	0.03	0.18	2,490,092	0.04	0.19	1,756,351	0.07	0.26
June (0/1)	5,087,685	0.06	0.23	841,242	0.04	0.19	2,490,092	0.06	0.23	1,756,351	0.07	0.25
July (0/1)	5,087,685	0.09	0.29	841,242	0.06	0.24	2,490,092	0.09	0.29	1,756,351	0.12	0.32
August (0/1)	5,087,685	0.08	0.27	841,242	0.05	0.22	2,490,092	0.08	0.28	1,756,351	0.09	0.29
September (0/1)	5,087,685	0.06	0.24	841,242	0.09	0.28	2,490,092	0.08	0.27	1,756,351	0.02	0.12
October (0/1)	5,087,685	0.07	0.25	841,242	0.08	0.27	2,490,092	0.11	0.32	1,756,351	0.00	0.00
November (0/1)	5,087,685	0.06	0.24	841,242	0.13	0.33	2,490,092	0.08	0.27	1,756,351	0.00	0.00
December (0/1)	5,087,685	0.13	0.33	841,242	0.14	0.35	2,490,092	0.21	0.41	1,756,351	0.00	0.00
Controls of time of the month												
Beginning (0/1)	5,087,685	0.36	0.48	841,242	0.34	0.47	2,490,092	0.34	0.47	1,756,351	0.39	0.49
Halfway (0/1)	5,087,685	0.33	0.47	841,242	0.33	0.47	2,490,092	0.33	0.47	1,756,351	0.32	0.46
End (0/1)	5,087,685	0.31	0.46	841,242	0.32	0.47	2,490,092	0.32	0.47	1,756,351	0.30	0.46
Controls of time of the day												
Night (0/1)	5,087,685	0.14	0.34	841,242	0.11	0.31	2,490,092	0.16	0.36	1,756,351	0.12	0.32
Morning (0/1)	5,087,685	0.14	0.35	841,242	0.15	0.36	2,490,092	0.12	0.33	1,756,351	0.16	0.36
Afternoon (0/1)	5,087,685	0.27	0.44	841,242	0.31	0.46	2,490,092	0.23	0.42	1,756,351	0.31	0.46
Evening (0/1)	5,087,685	0.46	0.50	841,242	0.42	0.49	2,490,092	0.50	0.50	1,756,351	0.42	0.49

Table B.5: Summary statistics of location controls

	Aggregate			Control			Semi-Treatment			Full-Treatment		
	Obs.	Mean	St.Dev.	Obs.	Mean	St.Dev.	Obs.	Mean	St.Dev.	Obs.	Mean	St.Dev.
Controls of country												
France (0/1)	5,087,685	0.25	0.43	841,242	0.26	0.44	2,490,092	0.27	0.44	1,756,351	0.22	0.41
Italy (0/1)	5,087,685	0.23	0.42	841,242	0.22	0.42	2,490,092	0.20	0.40	1,756,351	0.27	0.44
Spain (0/1)	5,087,685	0.15	0.36	841,242	0.16	0.37	2,490,092	0.15	0.36	1,756,351	0.15	0.36
Portugal (0/1)	5,087,685	0.07	0.25	841,242	0.06	0.25	2,490,092	0.08	0.27	1,756,351	0.06	0.24
United Kingdom (0/1)	5,087,685	0.05	0.22	841,242	0.05	0.22	2,490,092	0.06	0.25	1,756,351	0.03	0.18
Denmark (0/1)	5,087,685	0.04	0.20	841,242	0.04	0.19	2,490,092	0.04	0.19	1,756,351	0.05	0.21
Netherlands (0/1)	5,087,685	0.04	0.19	841,242	0.04	0.19	2,490,092	0.04	0.19	1,756,351	0.04	0.20
Belgium (0/1)	5,087,685	0.02	0.13	841,242	0.02	0.14	2,490,092	0.02	0.13	1,756,351	0.02	0.14
Ukraine (0/1)	5,087,685	0.02	0.13	841,242	0.02	0.12	2,490,092	0.01	0.10	1,756,351	0.03	0.16
United States (0/1)	5,087,685	0.01	0.12	841,242	0.02	0.12	2,490,092	0.01	0.12	1,756,351	0.01	0.12
Others (0/1)	5,087,685	0.12	0.32	841,242	0.12	0.32	2,490,092	0.12	0.33	1,756,351	0.12	0.32
Controls of currency												
EUR (0/1)	5,087,685	0.99	0.08	841,242	0.99	0.08	2,490,092	0.99	0.09	1,756,351	0.99	0.07
GBR (0/1)	5,087,685	0.00	0.07	841,242	0.00	0.06	2,490,092	0.01	0.08	1,756,351	0.00	0.06
USD (0/1)	5,087,685	0.00	0.04	841,242	0.00	0.04	2,490,092	0.00	0.04	1,756,351	0.00	0.04

C Omitted Tables

Table C.1: Likelihood of conversion with new user dummy and purchase history dummy

Likelihood of conversion	Control	Semi-Treatment	Full-Treatment
Intersection (γ_0)	-2.396***	-2.328***	-2.058***
	(0.279)	(0.134)	(0.174)
New user dummy (γ_U)	-0.121***[11]	-0.101***[10]	-0.141***[13]
	(0.013)	(0.007)	(0.009)
Purchase history dummy (γ_H)	0.472***[60]	0.657***[93]	0.608***[84]
	(0.030)	(0.014)	(0.017)
Main variables	✓	✓	✓
Technology controls	✓	✓	✓
Time controls	✓	✓	✓
Location controls	✓	✓	✓
Pseudo R-squared	0.095	0.077	0.100
Observations	218,753	749,093	488,762

Note: * $-p < 0.1$; ** $-p < 0.05$; *** $-p < 0.01$. Logit regression of the likelihood of conversion on new user dummy, past purchase dummy, and covariates as described in (3.1). The control variables are defined in Section B.1. Standard errors in parenthesis. Odds ratio absolute value percentage variations in brackets.

Table C.2: Gross order with new user dummy and purchase history dummy

Gross order	Control	Semi-Treatment	Full-Treatment
Intersection (δ_0)	3.940***	3.739***	3.822***
	(0.116)	(0.054)	(0.065)
New user dummy (δ_U)	-0.025***	-0.028***	-0.047***
	(0.005)	(0.003)	(0.003)
Purchase history dummy (δ_H)	0.083***	0.104***	0.014***
	(0.010)	(0.006)	(0.007)
Main variables	✓	✓	✓
Technology controls	✓	✓	✓
Time controls	✓	✓	✓
Location controls	✓	✓	✓
Adj. R-squared	0.166	0.182	0.136
Observations	41,359	140,046	90,513

Note: * $-p < 0.1$; ** $-p < 0.05$; *** $-p < 0.01$. OLS regression of the gross order on new user dummy, past purchase dummy, and covariates as described in (3.2). The control variables are defined in Section B.1. Standard errors in parenthesis.

Table C.3: Estimates of ITT and LATE for new and returning users

New users						
	Combined Treatments		Semi-Treatment		Full-Treatment	
Likelihood of conversion	ITT	LATE	ITT	LATE	ITT	LATE
Intersection (β_0)	-3.143***	-2.361***	-3.143***	-2.466***	-3.143***	-2.161***
	(0.007)	(0.038)	(0.007)	(0.035)	(0.007)	(0.052)
ITT (β_{ITT})	0.135***[14]		0.131***[13]		0.142***[15]	
	(0.008)		(0.008)		(0.008)	
α_0^I		-1.210***		-1.210***		-1.210***
		(0.003)		(0.003)		(0.003)
α_{IV}^I		0.210***		0.2332***		0.1752***
		(0.004)		(0.004)		(0.004)
LATE (β_{LATE})		0.646***[91]		0.560***[75]		0.812***[125]
		(0.036)		(0.034)		(0.048)
Pseudo R-squared	0.001	0.001	0.001	0.001	0.001	0.001
Observations	3,163,906	3,163,906	2,072,335	2,072,335	1,602,216	1,602,216
Returning users						
	Combined Treatments		Semi-Treatment		Full-Treatment	
Likelihood of conversion	ITT	LATE	ITT	LATE	ITT	LATE
Intersection (β_0)	-2.730***	-2.099***	-2.730***	-1.856***	-2.730***	-3.899***
	(0.007)	(0.052)	(0.007)	(0.035)	(0.007)	(0.199)
ITT (β_{ITT})	0.086***[9]		0.174***[19]		-0.051***[5]	
	(0.008)		(0.008)		(0.009)	
α_0^I		-0.948***		-0.948***		-0.948***
		(0.004)		(0.004)		(0.004)
α_{IV}^I		0.128***		0.189***		0.041***
		(0.004)		(0.004)		(0.005)
LATE (β_{LATE})		0.666***[95]		0.922***[151]		-1.233***[7]
		(0.062)		(0.044)		(0.216)
Pseudo R-squared	0.001	0.001	0.001	0.001	0.001	0.001
Observations	1,923,779	1,923,779	1,258,999	1,258,999	995,377	995,377

Note: $^{*} - p < 0.1$; $^{**} - p < 0.05$; $^{***} - p < 0.01$. The results of causal inference analysis for combined treatments, semi-treatment group and full-treatment group as described in (3.5) and (3.6). Standard errors in parenthesis. Odds ratio absolute value percentage variation in brackets.

Table C.4: Likelihood of conversion with discount percentage and minimum purchase threshold

Likelihood of conversion	Control	Semi-Treatment	Full-Treatment
Intersection (γ_0)	-2.452***	-2.419***	-2.260***
	(0.279)	(0.135)	(0.177)
Discount percentage (γ_d)		0.019***[2]	0.028***[3]
		(0.001)	(0.002)
Minimum purchase threshold (γ_{MPT})			-0.003***[0]
			(0.001)
Main variables	✓	✓	✓
Technology controls	✓	✓	✓
Time controls	✓	✓	✓
Location controls	✓	✓	✓
Pseudo R-squared	0.095	0.081	0.102
Observations	218,753	749,093	488,762

Note: $^{*} - p < 0.1$; $^{**} - p < 0.05$; $^{***} - p < 0.01$. Logit regression of the likelihood of conversion on discount percentage, minimum purchase threshold and covariates as described in (4.1). The control variables are defined in Section B.1. Standard errors in parenthesis. Odds ratio absolute value percentage variations in brackets.

Table C.5: Gross order with discount percentage and minimum purchase threshold

Gross order	Control	Semi-Treatment	Full-Treatment
Intersection (δ_0)	3.940***	3.460***	3.893 ***
	(0.116)	(0.053)	(0.066)
Discount percentage (δ_d)		0.021***	-0.011***
		(0.000)	(0.001)
Minimum purchase threshold (δ_{MPT})			0.002***
			(0.000)
Main variables	✓	✓	✓
Technology controls	✓	✓	✓
Time controls	✓	✓	✓
Location controls	✓	✓	✓
Adj. R-squared	0.166	0.194	0.137
Observations	41,359	140,046	90,513

Note: $^{*} - p < 0.1$; $^{**} - p < 0.05$; $^{***} - p < 0.01$. OLS regression of the gross order on discount percentage, minimum purchase threshold and covariates as described in (4.2). The control variables are defined in Section B.1. Standard errors in parenthesis.

Table C.6: Estimates of treatment elasticity

	Control	Semi-Treatment	Full-Treatment		Aggregate Demand	
	Low (l)	High (h)	Low (l)	High (h)	Low (l)	High (h)
			Discount 5%			
$\varepsilon_{l(h)}$	-1.294	-1.256	-1.592	-2.285	-2.178	-1.222
	[-1.359,-1.127]	[-1.431,-1.209]	[-1.655,-1.492]	[-2.374,-2.229]	[-2.240,-2.029]	[-1.335,-1.135]
ε_0	0.021	0.022	0.015	0.011	0.011	0.018
	[-0.124,0.087]	[-0.136,0.111]	[-0.025,0.128]	[-0.084,0.061]	[-0.106,0.104]	[-0.053,0.143]
			Discount 10%			
$\varepsilon_{l(h)}$	-1.312	-1.272	-1.639	-2.434	-2.307	-1.235
	[-1.468,-1.221]	[-1.381,-1.147]	[-1.791,-1.633]	[-2.536,-2.384]	[-2.407,-2.186]	[-1.305,-1.116]
ε_0	0.022	0.023	0.015	0.012	0.011	0.018
	[-0.047,0.204]	[-0.178,0.028]	[-0.046,0.108]	[-0.043,0.092]	[-0.113,0.097]	[-0.029,0.178]
			Discount 15%			
$\varepsilon_{l(h)}$	-1.331	-1.288	-1.688	-2.604	-2.453	-1.249
	[-1.450,-1.196]	[-1.449,-1.226]	[-1.835,-1.680]	[-2.729,-2.576]	[-2.617,-2.390]	[-1.266,-1.042]
ε_0	0.023	0.024	0.016	0.012	0.012	0.019
	[-0.011,0.229]	[-0.134,0.107]	[-0.039,0.102]	[0.028,0.165]	[-0.095,0.125]	[-0.166,0.061]
			Discount 20%			
$\varepsilon_{l(h)}$	-1.351	-1.304	-1.740	-2.799	-2.618	-1.262
	[-1.399,-1.165]	[-1.339,-1.114]	[-1.827,-1.685]	[-2.901,-2.751]	[-2.765,-2.559]	[-1.378,-1.167]
ε_0	0.024	0.025	0.017	0.013	0.012	0.020
	[-0.177,0.042]	[-0.051,0.162]	[-0.047,0.092]	[-0.081,0.064]	[-0.114,0.098]	[-0.203,0.006]

Note: Refer to Section A.3 in the Online Appendix. Endowment M_i is set to €101. 95% bootstrap confidence intervals are in square brackets (501 repetitions). The dataset includes sessions in the control group and only sessions where the treatment was received with a order gross equal to or lower than €100.

Table C.7: Performance of classification ensemble methods

Accuracy $a(\mu_{it}(x), \widehat{\mu}_{it}(x))$	Control	Semi-Treatment	Full-Treatment
BAG (`BaggingClassifier`)	0.768	0.778	0.777
RF (`RandomForestClassifier`)	0.813	0.815	0.817
AB (`AdaBoostClassifier`)	0.813	0.815	0.817
GBDT (`GradientBoostingClassifier`)	0.814	0.817	0.819
NN (`MLPClassifier`)	0.814	0.817	0.818
Observations	218,753	749,093	488,762

Note: Classification methods' performance through the accuracy $a(\mu_{it}(x), \widehat{\mu}_{it}(x)) = \frac{1}{N}\sum_{i=0}^{N-1} \mathbb{1}_{\{\widehat{\mu}_{it}(x)=\mu_{it}(x)\}}$. Accuracy score computes the fraction (or the count) of correct predictions in the binary classification. The dataset includes only sessions where the treatment was received (1,456,608 observations).

Table C.8: Performance of regression ensemble methods

Coefficient of determination $R^2(q_{it}(x), \widehat{q}_{it}(x))$	Control	Semi-Treatment	Full-Treatment
BAG (`BaggingRegressor`)	0.012	0.340	-0.119
RF (`RandomForestRegressor`)	0.015	0.393	-0.067
AB (`AdaBoostRegressor`)	0.014	-0.589	-0.402
GBDT (`GradientBoostingRegressor`)	0.176	0.404	0.062
NN (`MLPRegressor`)	0.096	0.380	0.057
Observations	218,753	749,093	488,762

Note: Regression methods' performance through the coefficient of determination $R^2(q_{it}(x), \widehat{q}_{it}(x)) = 1 - \frac{\sum_{i=1}^{N}(q_{it}(x)-\widehat{q}_{it}(x))^2}{\sum_{i=1}^{N}(q_{it}(x)-\bar{q}_{it}(x))^2}$. It represents the proportion of variance of $q_{it}(x)$ that has been explained by the independent variables x in the model. It provides an indication of goodness of fit and thus a measure of how well unseen samples are likely to be predicted by the model, through the proportion of explained variance. Best possible score is 1.0 and it can be negative, because the model can be arbitrarily worse. The dataset includes only sessions where the treatment was received (1,456,608 observations).

Table C.9: Estimates of counterfactual simulations of conversion rate

	Semi-Treatment		Full-Treatment		Combined Treatments	
Policy	$\widehat{\mu}^{IPW}_{f_{iST}(x)}$	$\widehat{\Delta\mu}_{f_{iST}(x)}$	$\widehat{\mu}^{IPW}_{f_{iFT}(x)}$	$\widehat{\Delta\mu}_{f_{iFT}(x)}$	$\widehat{\mu}^{IPW}_{f_{i\{ST,FT\}}(x)}$	$\widehat{\Delta\mu}_{f_{i\{ST,FT\}}(x)}$
Uniform	1.3036	–	1.0077	–	2.3064	–
Non-targeted	0.4601	–64.71%	1.0077	0.00%	2.6982	+16.99%
		(0.0044)		(0.0080)		(0.0227)
Conversion-targeted	1.3055	+0.14%	1.1188	+11.03%	2.6775	+16.09%
		(0.0114)		(0.0074)		(0.0190)
Revenue-targeted	1.4962	+14.77%	1.1457	+13.69%	2.8457	+23.38%
		(0.0125)		(0.0075)		(0.0200)
Conversion-revenue	1.4973	+14.86%	1.1461	+13.73%	2.8466	+23.42%
		(0.0125)		(0.0075)		(0.0200)
Optimal treatment	1.5501	+18.91%	1.1601	+15.13%	2.8642	+24.18%
		(0.0128)		(0.0076)		(0.0200)

Note: The results of simulation analysis of conversion rate of alternative policies and optimal treatment policy. First and second column consider the treatment space $\mathcal{T} = \{C, ST\}$ and $\mathcal{T} = \{C, FT\}$, respectively. The third column simulates the pricing algorithm with the combined treatment space $\mathcal{T} = \{C, ST, FT\}$. Comparison between alternative policies is framed in terms of percentage variation defined as $\widehat{\Delta\mu}_{f_{it}(x)} = \left(\frac{\widehat{\mu}^{IPW}_{f_{it}(x)} - \widehat{\mu}^{IPW}_{uniform}}{\widehat{\mu}^{IPW}_{uniform}}\right)_{\%}$, where the non-treatment, uniform policy stands for the benchmark case. The estimates of conversion rate of (5.17) per counterfactual policies and optimal treatment policy uses dataset sessions where the treatment was received. The testing dataset is equal to 20%. Semi-treatment testing dataset features 19,908 observations, full-treatment dataset has 36,748 sessions, and the combined treatments testing dataset counts 56,656 data points. Standard errors in parenthesis.

Table C.10: Estimates of counterfactual simulations of conversion rate for new and returning users

New users

	Semi-Treatment		Full-Treatment		Combined Treatments	
Policy	$\hat{\mu}^{IPW}_{f_{iST}(x)}$	$\widehat{\Delta\mu}_{f_{iST}(x)}$	$\hat{\mu}^{IPW}_{f_{iFT}(x)}$	$\widehat{\Delta\mu}_{f_{iFT}(x)}$	$\hat{\mu}^{IPW}_{f_{i\{ST,FT\}}(x)}$	$\widehat{\Delta\mu}_{f_{i\{ST,FT\}}(x)}$
Uniform	1.4390	–	1.0874	–	2.5461	–
Non-targeted	1.4390	0.00%	1.0874	0.00%	2.2065	-13.33%
		(0.0275)		(0.0171)		(0.0422)
Conversion-targeted	1.3902	-3.39%	1.2542	+15.33%	2.7037	+6.20%
		(0.0227)		(0.0166)		(0.0386)
Revenue-targeted	1.5844	+10.10%	1.2068	+10.97%	2.9324	+15.18%
		(0.0247)		(0.0162)		(0.0394)
Conversion-revenue	1.5861	+10.21%	1.2528	+15.21%	2.8718	+12.80%
		(0.0247)		(0.0166)		(0.0395)
Optimal treatment	1.6403	+13.99%	1.2204	+12.23%	2.9568	+16.14%
		(0.0252)		(0.0163)		(0.0395)

Returning users

	Semi-Treatment		Full-Treatment		Combined Treatments	
Policy	$\hat{\mu}^{IPW}_{f_{iST}(x)}$	$\widehat{\Delta\mu}_{f_{iST}(x)}$	$\hat{\mu}^{IPW}_{f_{iFT}(x)}$	$\widehat{\Delta\mu}_{f_{iFT}(x)}$	$\hat{\mu}^{IPW}_{f_{i\{ST,FT\}}(x)}$	$\widehat{\Delta\mu}_{f_{i\{ST,FT\}}(x)}$
Uniform	1.4746	–	1.1341	–	2.6089	–
Non-targeted	0.4985	–66.19%	1.1341	0.00%	3.0189	+15.71%
		(0.0114)		(0.0222)		(0.0588)
Conversion-targeted	1.8063	+22.49%	1.2406	+9.39%	3.3554	+28.61%
		(0.0338)		(0.0200)		(0.0518)
Revenue-targeted	1.6273	+10.35%	1.2701	+11.99%	3.1621	+21.20%
		(0.0319)		(0.0204)		(0.0514)
Conversion-revenue	1.8054	+22.44%	1.2714	+12.11%	3.2922	+26.19%
		(0.0338)		(0.0204)		(0.0518)
Optimal treatment	1.6916	+14.71%	1.2949	+14.18%	3.1803	+21.90%
		(0.0326)		(0.0206)		(0.0514)

Note: The estimates of the counterfactual analysis of alternative policies and optimal treatment policy conditional on new users' sessions. First and second column consider the treatment space $\mathcal{T} = \{C, ST\}$ and $\mathcal{T} = \{C, FT\}$, respectively. The third column simulates the pricing algorithm described in (5.11) and implemented by Algorithm 6 where the treatment space $\mathcal{T} = \{C, ST, FT\}$ combines $t = ST$ and $t = FT$. The first section presents the estimates of conversion rate's simulation analysis of (5.17) of new users per counterfactual policies and optimal treatment policy. The second section shows the results of conversion rate's simulation analysis of (5.17) of returning users per counterfactual policies and optimal treatment policy. Comparison between alternative policies is framed in terms of percentage variation defined as $\widehat{\Delta\mu}_{f_{it}(x)} = \left(\frac{\hat{\mu}^{IPW}_{f_{it}(x)} - \hat{\mu}^{IPW}_{uniform}}{\hat{\mu}^{IPW}_{uniform}}\right)_{\%}$, where the non-treatment, uniform policy stands for the benchmark case. For new users, the simulation dataset includes only sessions where the treatment was received and the testing dataset is set to 20%. Semi-treatment testing dataset features 5,517 observations, while full-treatment dataset has 8,353 observations. The combined treatments testing dataset counts 13,870 sessions. For returning users, semi-treatment testing dataset features 3,146 observations, while full-treatment dataset has 5,073 observations. The combined treatments testing dataset counts 8,219 sessions. Standard errors in parenthesis.

D Omitted Figures

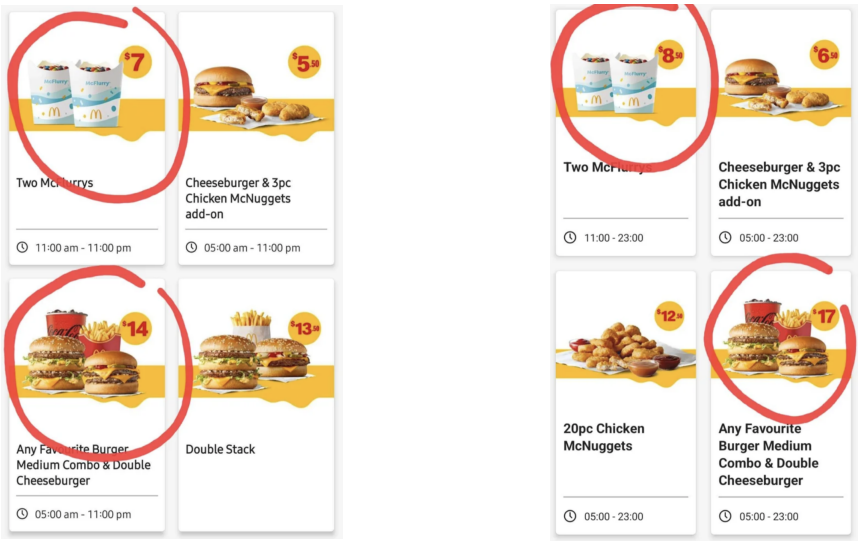


Figure D.1: Screenshots of McDonald's app (Source: Taunton (2023))

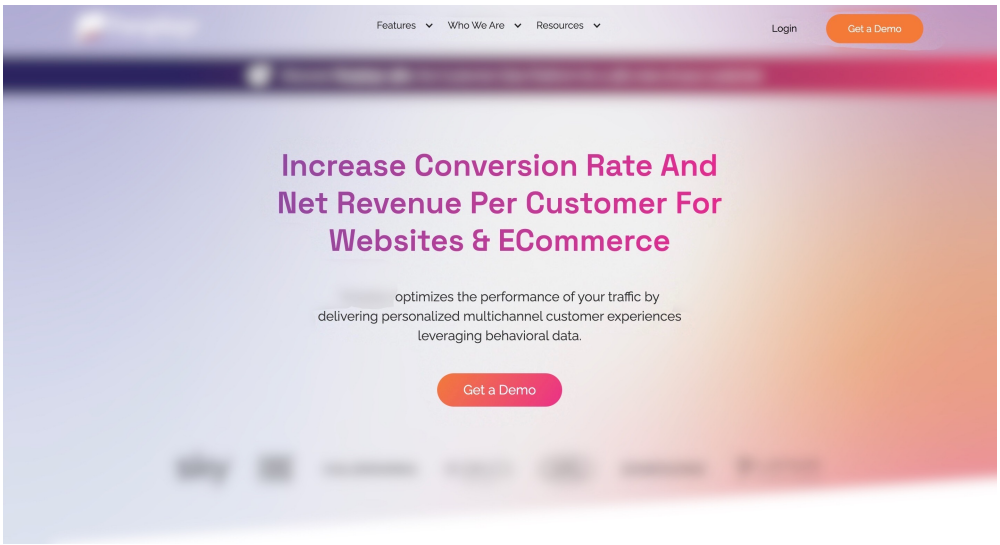
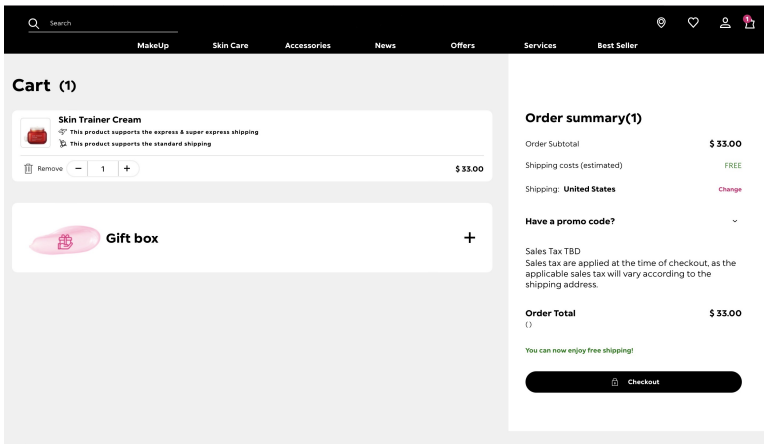
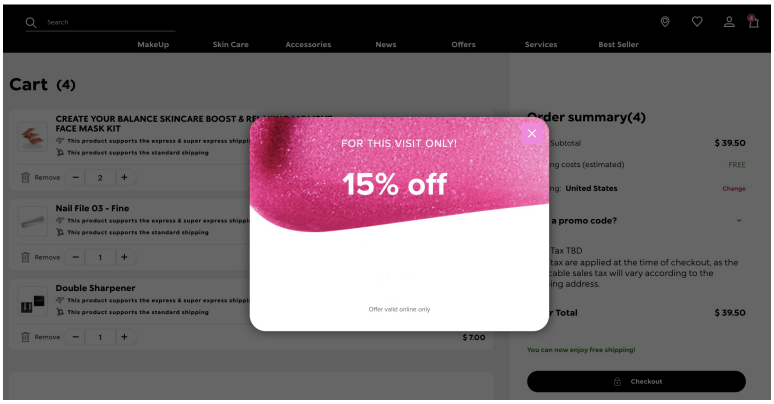


Figure D.2: Homepage of *Alpha* website

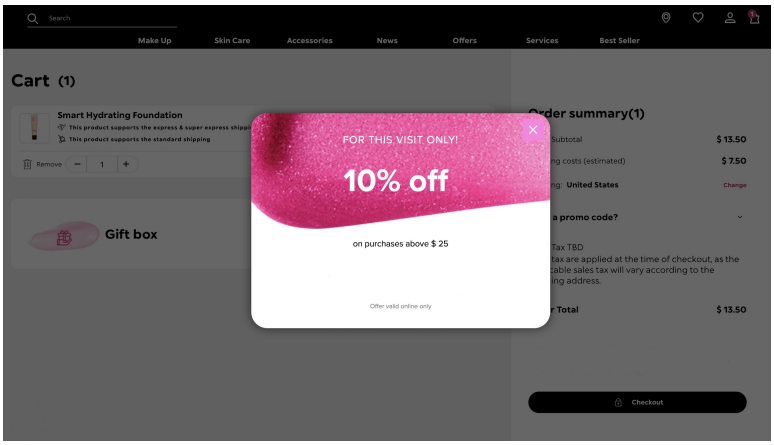
Note: Screenshot of *Alpha* homepage. The text advertises that *Alpha*'s clients would raise the likelihood of conversion and would increase net revenue per user by means of *"personalized multichannel consumer experiences"* which leverage e-commerce behavioral data. Images have been edited to comply with MNDA.



(a) Control group



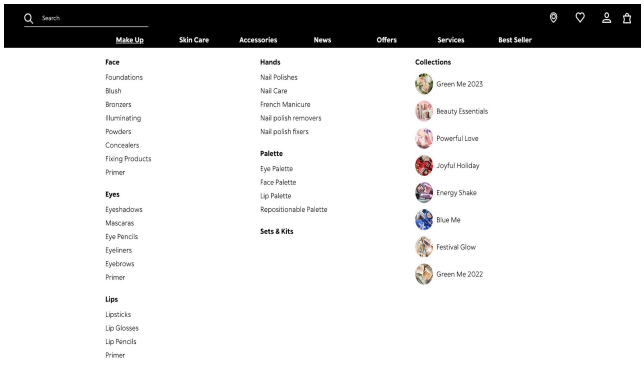
(b) Semi-treatment group



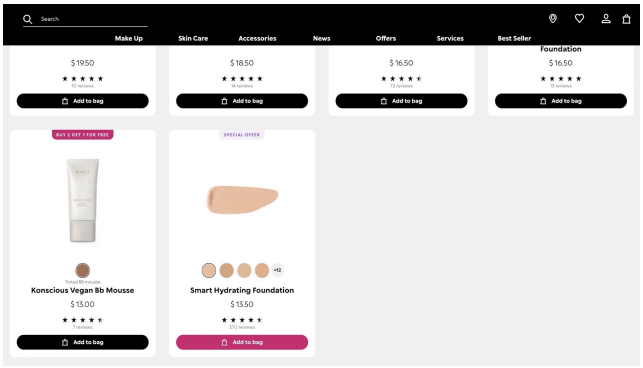
(c) Full-treatment group

Figure D.3: Treatment per control, semi-treatment and full-treatment group

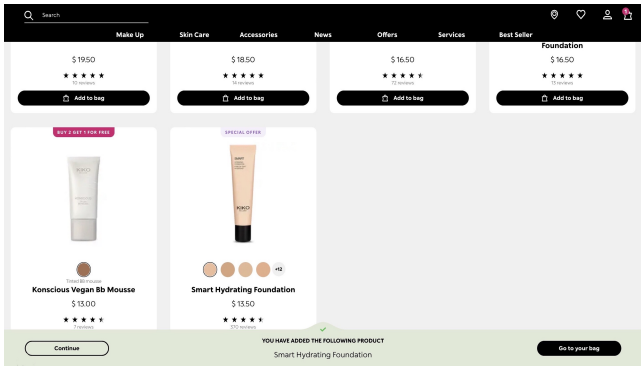
Note: Screenshots of the online experience of receiving the treatment in the form of promotion pop-up based on groups' assignment. (a) Control group C with $d_t = 0$ and $MPT_t = 0$. Promotion pop-up does not appear at the cart page. (b) Semi-treatment group ST with $d_t = 15\%$ and $MPT_t = 0$. The cart is worth \$39.50 and consists of four items and three products' types. The promotion pop-up is received by the time the user reaches the cart page for the first time. (c) Full-treatment group FT with $d_t = 10\%$ and $MPT_t = \$25$. The cart consists of a \$13.50 product and estimated shipping costs equal to \$7.50. Accessed on November 4th, 2022 at 7:35 PM from the San Francisco area using a MacBook Pro laptop with no prior purchases, but some prior visits to the e-commerce website. Images have been edited to comply with MNDA.



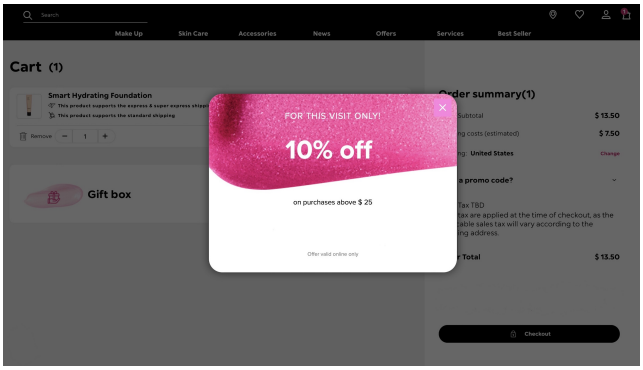
(a) E-commerce homepage



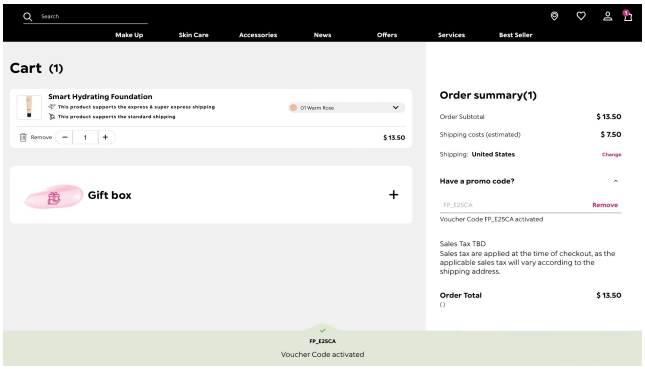
(b) Session started



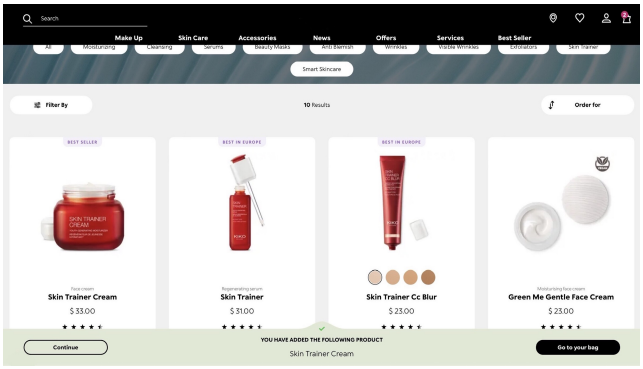
(c) First item added to the cart page



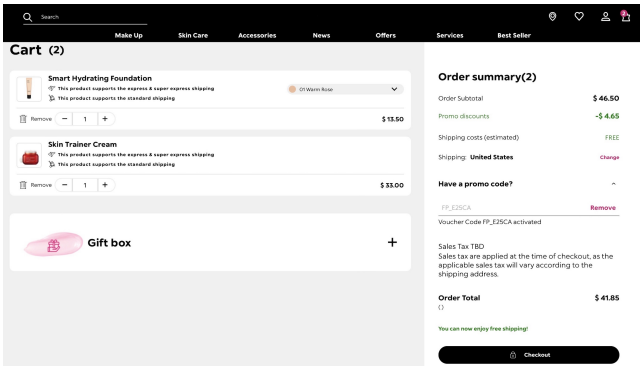
(d) Treatment received



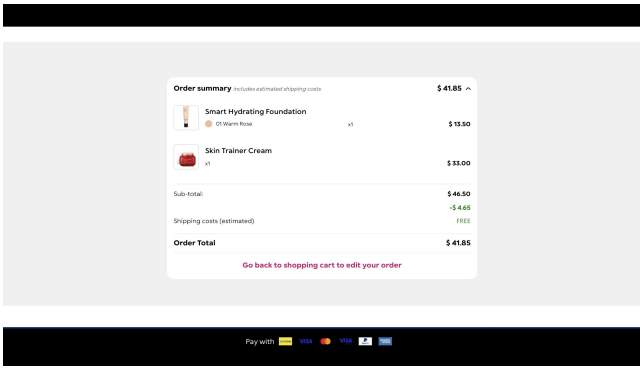
(e) Promotion added to the cart page



(f) Second item added to the cart page



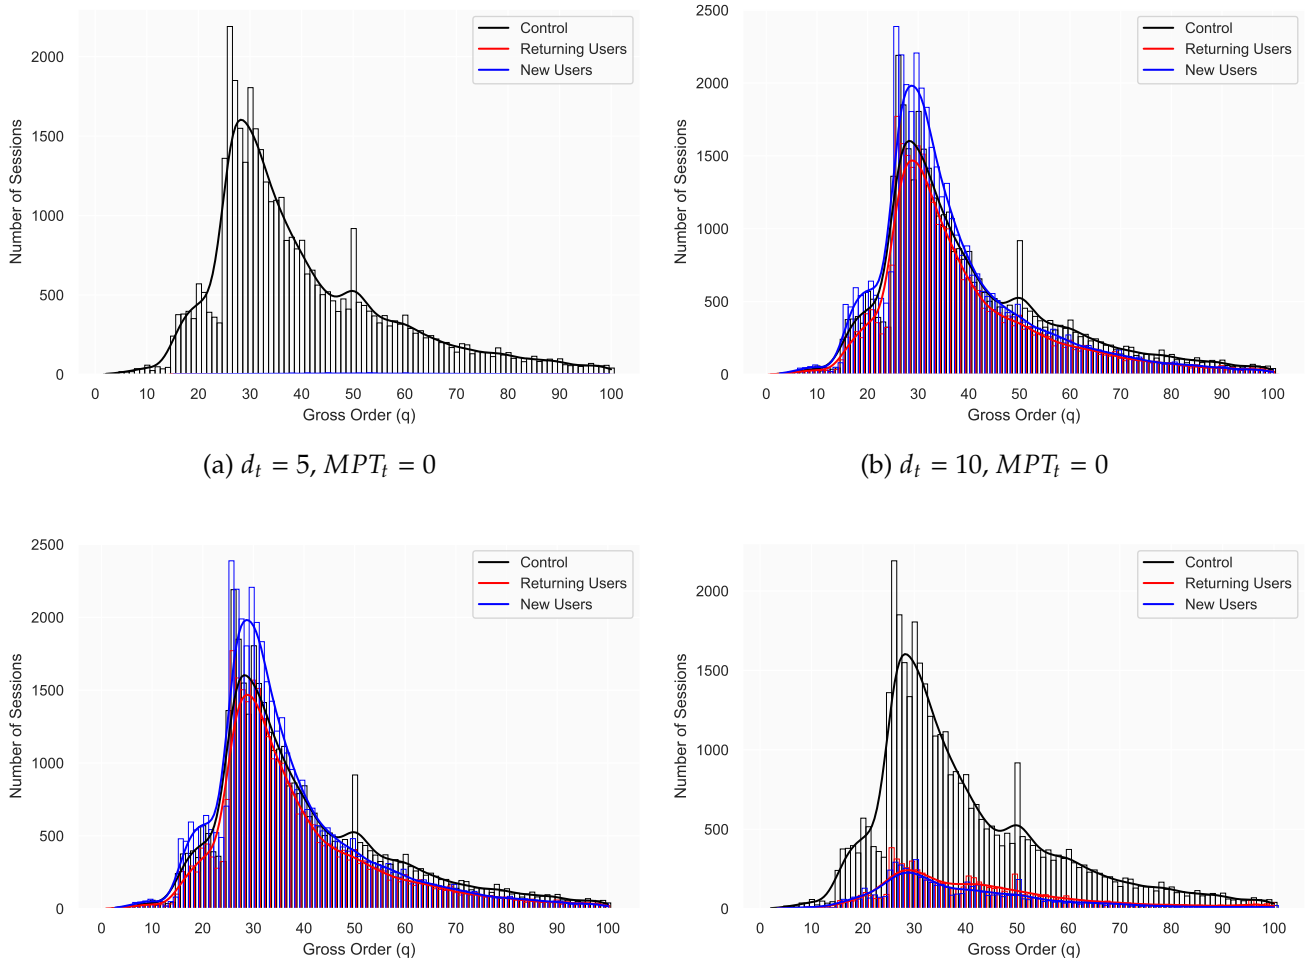
(g) Gross order updated



(h) Order converted

Figure D.4: Example of conversion funnel experience

Note: Images of the conversion funnel on *Alpha*'s client e-commerce website for a session assigned to the full-treatment group. (a) The e-commerce homepage of the website displays the architecture divided by products' typology (i.e., make-up, skin care, accessories) and further features (i.e., offers, services, best seller). Each page is divided accordingly the products' purpose. For instance, the make-up section shows groups divided for face, eyes, lips, hands, palette, and sets and kits' items. (b) The session begins as the user logs-in the e-commerce site, starts browsing the website in search of different products. (c) The first item is selected and a light-green banner at the bottom of the screen notifies that the item has been added to the cart. The notification offers two options: at the bottom-left to continue visiting the products' section or at the bottom-right to visualize the cart page. In this case, the user chooses a Smart Hydrating Foundation at a price of $13.50. (d) The user clicks the bottom-right alternative and the cart page is displayed for the first time during the session. The screen dims as a central promotional pop-up appears. The treatment of the full-treatment group consists of a promotion with two components: the discount percentage and the minimum purchase threshold. This section is assigned to a treatment consisting of a 10% discount on the gross order and a minimum purchase of $25. The promotion is valid only for this online session. (e) On the left side, the cart page summarizes the selected items with their quantities, characteristics and prices. The promotion is automatically added to the order summary located on the right side of the cart page that the user was visiting. In this case, the voucher code FP_E25CA in the promotion section is activated as the green bottom banner confirms the activation. (f) The session continues and a second product is added to the cart page. As a second item, the user picks a Skin Trainer Cream at a price of $33.00. The same green banner appears. (g) Again, the user clicks on the right button "Go to your bag". The order summary shows the number of items in the cart, 2, the order subtotal, $46.50, the promo discount, $4.65, and the order total, $41.85, and the checkout option is located at the bottom. (h) As the user moves to the checkout page, the order summary reports all the useful information in terms of selected products, prices, sub-total, and order total. At last the user needs to add the payment method information, then the order is confirmed and the session ends. Accessed on July 7th, 2022 at 9:55 PM from the San Francisco area using a MacBook Pro laptop with no prior purchases and no prior visits to the e-commerce website. Images have been edited to comply with MNDA.



(a) $d_t = 5$, $MPT_t = 0$

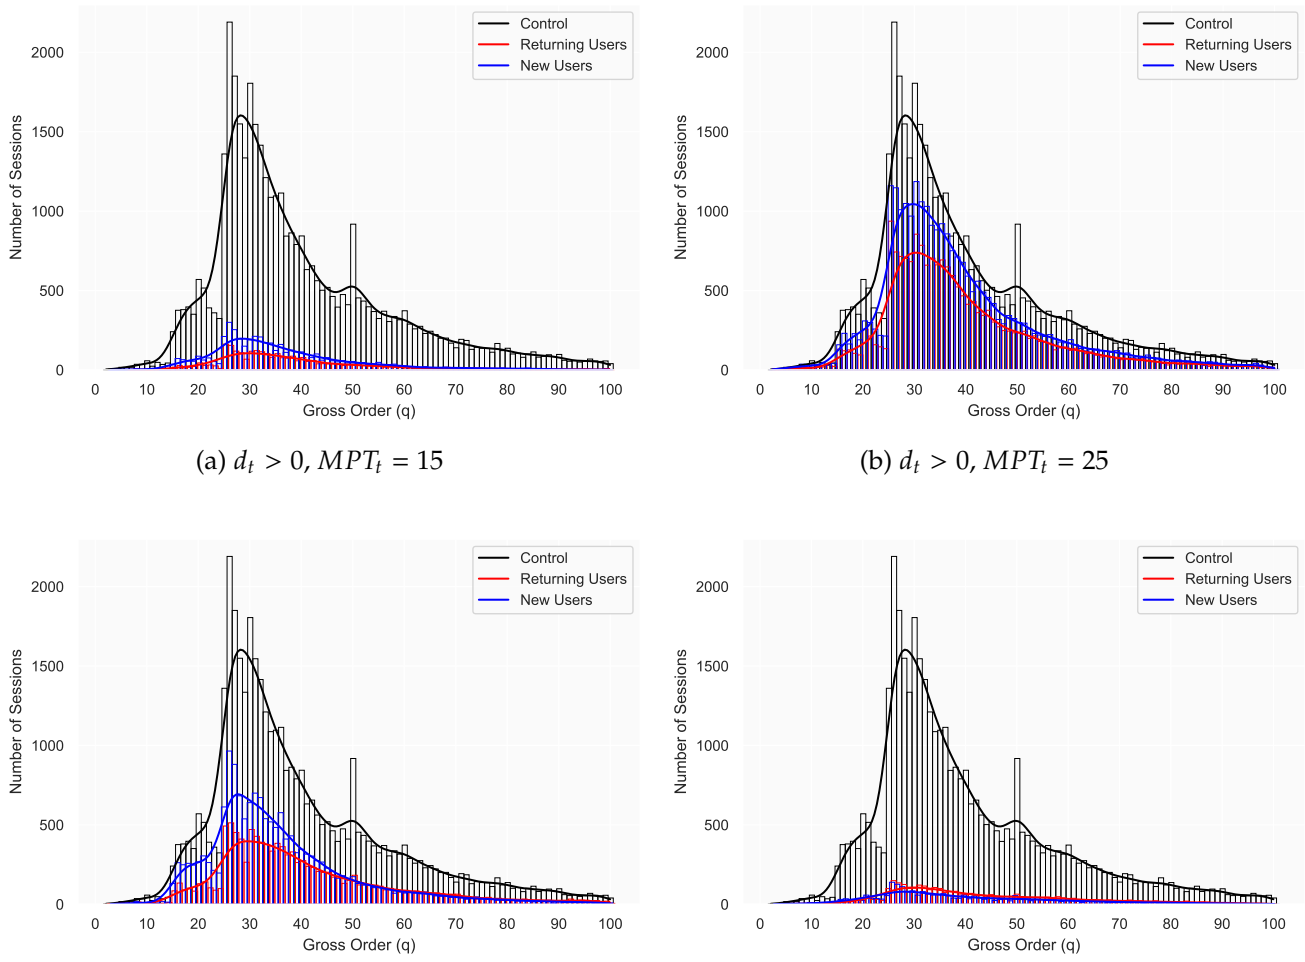
(b) $d_t = 10$, $MPT_t = 0$

(c) $d_t = 15$, $MPT_t = 0$

(d) $d_t = 20$, $MPT_t = 0$

Figure D.5: Number of converted sessions by gross order per semi-treatment discount percentage

Note: Panels show distributions of converted sessions for semi-treatment group per discount percentage treatments, conditionally on returning users' or new users' groups. Panels exclude sessions which were converted with a gross order above €100. Black histograms show the control group demand, which counts 39,552 sessions with €40.00 mean and €17.36 standard deviation. Semi-treatment demand counts 134,274 converted sessions in total, in which 47% are returning users (red plot), while 53% are new users (blue plot). (a) Semi-treatment demand with 5% discount percentage reports 391 converted sessions in total, where 35% are returning users (€54.14 mean, €18.88 st.dev.), while 65% are new users (€53.56 mean, €17.82 st.dev.). (b) Semi-treatment demand with 10% discount percentage reports 73,581 converted sessions in total, where 44% are returning users (€38.34 mean, €15.86 st.dev.), while 56% are new users (€36.70 mean, €15.22 st.dev.). (c) Semi-treatment demand with 15% discount percentage reports 46,677 converted sessions in total, where 51% are returning users (€40.97 mean, €16.48 st.dev.), while 49% are new users (€38.86 mean, €15.82 st.dev.). (d) Semi-treatment demand with 20% discount percentage reports 13,625 converted sessions in total, where 55% are returning users (€42.85 mean, €18.78 st.dev.), while 45% are new users (€39.95 mean, €17.62 st.dev.).



(a) $d_t > 0$, $MPT_t = 15$

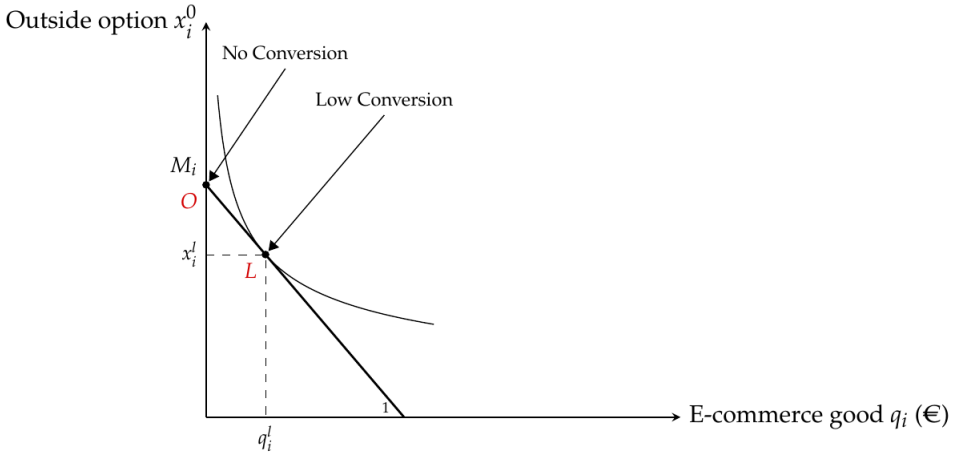
(b) $d_t > 0$, $MPT_t = 25$

(c) $d_t > 0$, $MPT_t = 35$

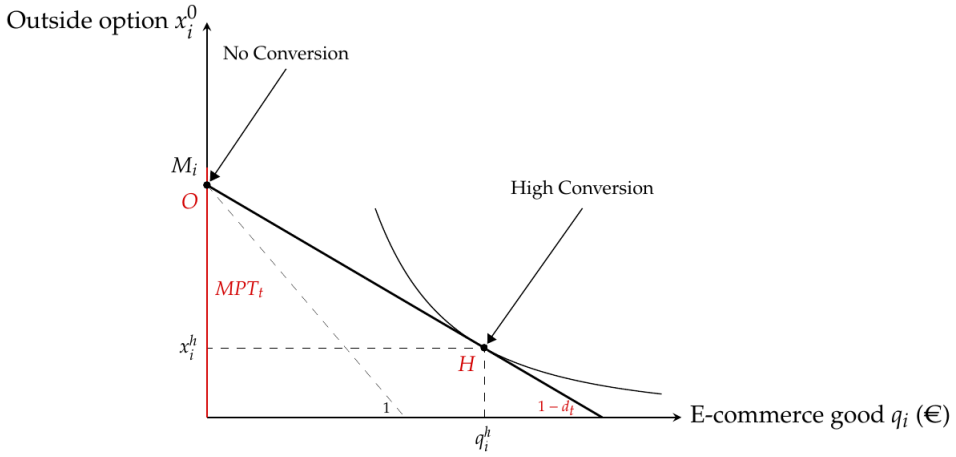
(d) $d_t > 0$, $MPT_t = 45$

Figure D.6: Number of converted sessions by gross order per full-treatment minimum purchase threshold

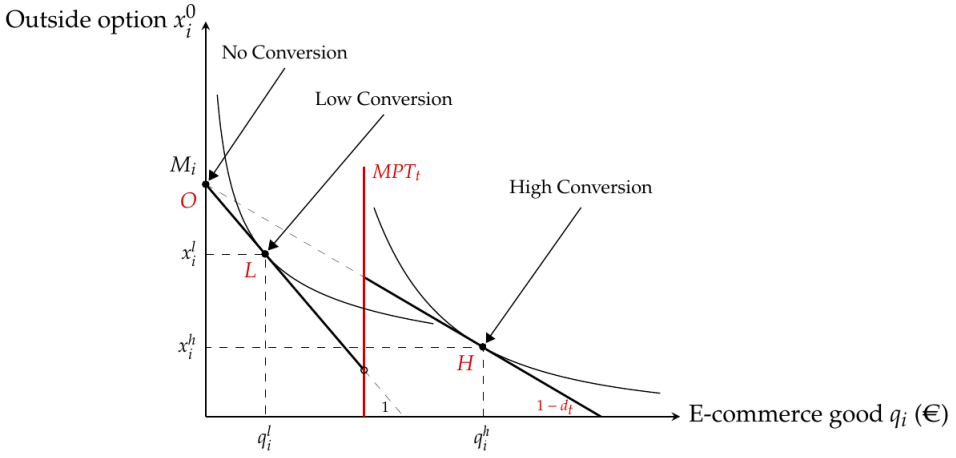
Note: Panels show distributions of converted sessions for semi-treatment group per discount percentage treatments, conditionally on returning users' or new users' groups. Panels exclude sessions which were converted with a gross order above €100. Black histograms show the control group demand, which counts 39,552 sessions with €40.00 mean and €17.36 standard deviation. Full-treatment demand counts 87,331 converted sessions in total, in which 43% are returning users (red plot), while 57% are new users (blue plot). (a) Full-treatment demand with $MPT_t = 15$ reports 7,609 converted sessions in total, where 37% are returning users (€38.85 mean, €15.78 st.dev.), while 63% are new users (€36.51 mean, €14.80 st.dev.). (b) Full-treatment demand with $MPT_t = 25$ reports 45,658 converted sessions in total, where 42% are returning users (€39.91 mean, €15.85 st.dev.), while 58% are new users (€39.14 mean, €15.85 st.dev.). (c) Full-treatment demand with $MPT_t = 35$ reports 28,055 converted sessions in total, where 41% are returning users (€40.77 mean, €17.07 st.dev.), while 59% are new users (€36.24 mean, €14.97 st.dev.). (d) Full-treatment demand with $MPT_t = 45$ reports 6,009 converted sessions in total, where 57% are returning users (€43.21 mean, €18.21 st.dev.), while 43% are new users (€40.72 mean, €18.60 st.dev.).



(a) Control group ($d_t = 0$, $MPT_t = 0$)



(b) Semi-treatment group ($d_t > 0$, $MPT_t = 0$)



(c) Full-treatment group ($d_t > 0$, $MPT_t > 0$)

Figure D.7: Graphical intuition of constrained utility maximization problem

Note: Each panel graphically describes the choice problem in (4.9). Control group has no treatment, the semi-treatment group gets unconditional discounts, and the full-treatment group receives conditional discounts. (a) Control group exhibits two demand equilibria, namely No Conversion and Low Conversion. The slope of budget constraint is 1. (b) Semi-treatment group shows two demand equilibria, that is No Conversion and High Conversion. The slope of budget constraint is equal to $1 - d_t$. (c) Full-treatment group has three demand equilibria: No Conversion, Low Conversion and High Conversion. The slope of budget constraint is 1 if $q_i < MPT_t$ and $1 - d_t$ otherwise.

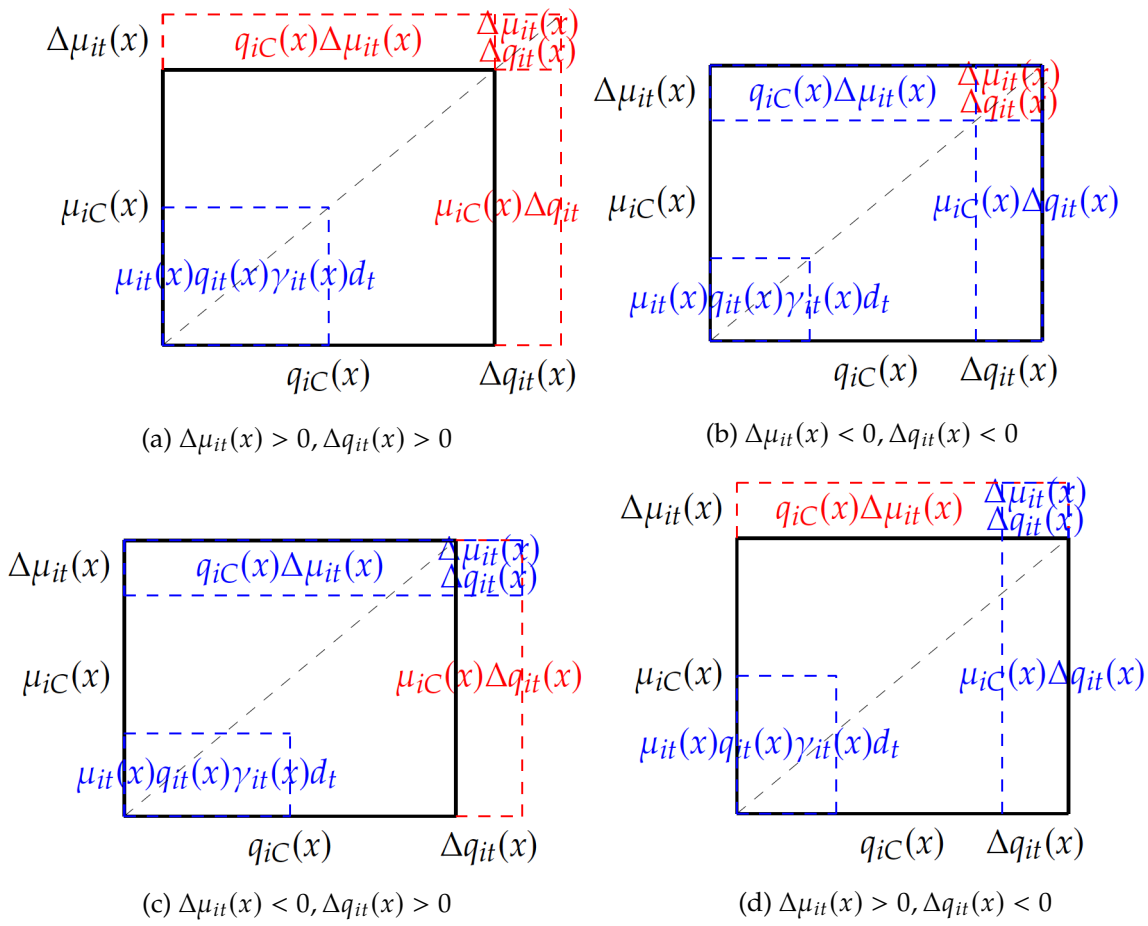
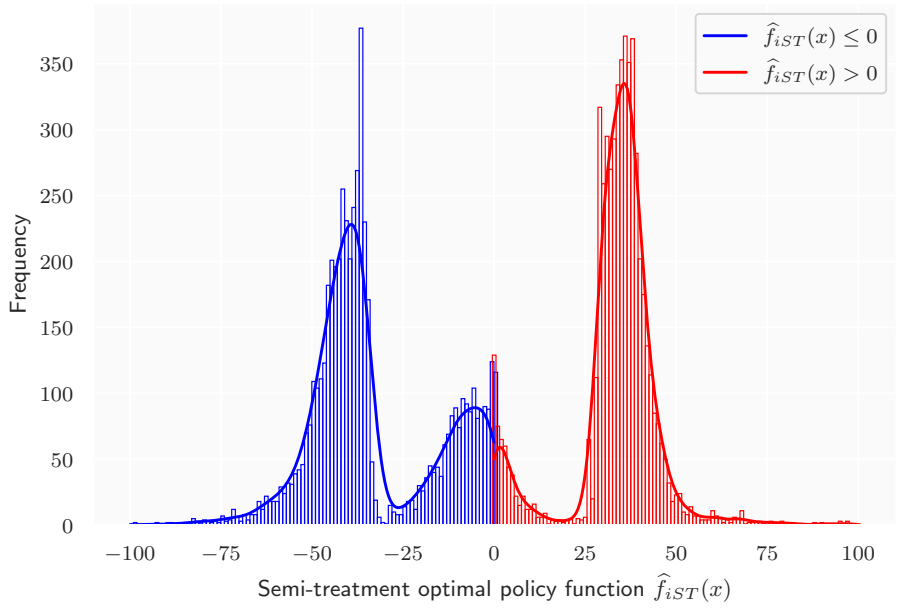
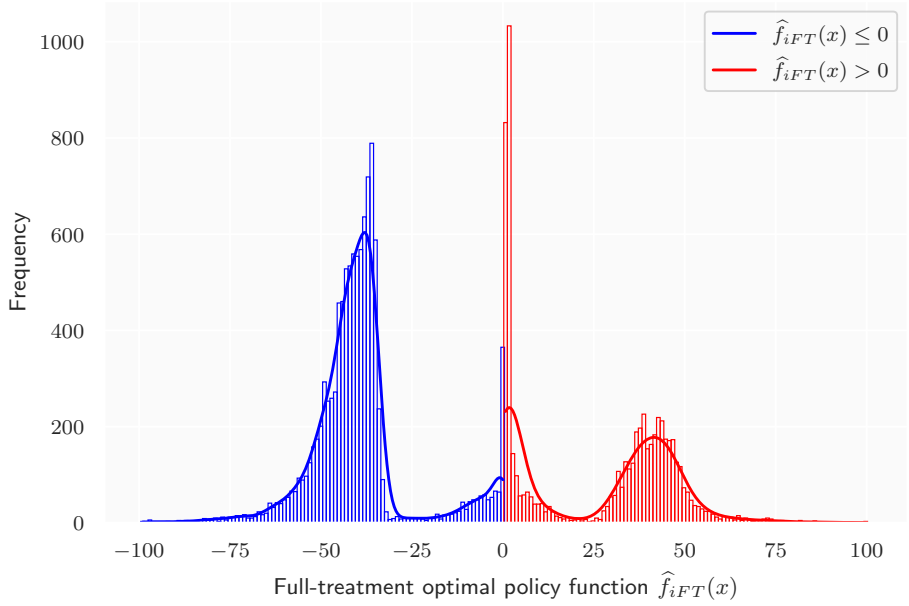


Figure D.8: Graphical analysis of optimal policy marginal variation

Note: Optimal treatment policy function (5.9) varies depending on the sign and the magnitude of CATEs for demand extensive margin $\Delta\mu_{it}(x)$ and demand intensive margin $\Delta q_{it}(x)$. The solid square depicts the expected profit for the control group as a product of the likelihood of conversion $\mu_{iC}(x)$ and the expected gross order $q_{iC}(x)$. Positive and negative CATEs of both dimensions are represented by red and blue dashed lines, respectively, on the right and upper square sides. The dashed parallelogram portrays the expected profit for the treatment group as a product of the likelihood of conversion $\mu_{it}(x) = \mu_{iC}(x) + \Delta\mu_{it}(x)$ and the expected gross order $q_{it}(x) = q_{iC}(x) + \Delta q_{it}(x)$. The blue dashed quadrilateral at the lower left corner stands for the expected cost of treatment t which is equal to the expected profit for the treatment group times the expected discount, namely the discount percentage d_t multiplied by the likelihood of complying to MPT_t $\gamma_{it}(x)$. Four scenarios are represented and described to study marginal changes in optimal treatment policy. In all four cases the expected cost of treatment t is equal to $-\gamma_{it}(x)\mu_{it}(x)q_{it}(x)d_t$ and it is negative (graphically blue).
For cases (a) and (b) the product $\Delta\mu_{it}(x)\Delta q_{it}(x)$ is positive (graphically red), while the control effects $q_{iC}(x)\Delta\mu_{it}(x)$ and $\mu_{iC}(x)\Delta q_{it}(x)$ are either positive or negative, depending on the cases:
(a) $f_{it}(x) = \Big(\Delta\mu_{it}(x)\Big)\Big(\Delta q_{it}(x)\Big) + q_{iC}(x)\Big(\Delta\mu_{it}(x)\Big) + \mu_{iC}(x)\Big(\Delta q_{it}(x)\Big) - \Big(\gamma_{it}(x)\mu_{it}(x)q_{it}(x)\Big)d_t$,
(b) $f_{it}(x) = \Big(\Delta\mu_{it}(x)\Big)\Big(\Delta q_{it}(x)\Big) + q_{iC}(x)\Big(\Delta\mu_{it}(x)\Big) + \mu_{iC}(x)\Big(\Delta q_{it}(x)\Big) - \Big(\gamma_{it}(x)\mu_{it}(x)q_{it}(x)\Big)d_t$.
For cases (c) and (d) the product $\Delta\mu_{it}(x)\Delta q_{it}(x)$ is negative (graphically blue), while the control effects $q_{iC}(x)\Delta\mu_{it}(x)$ and $\mu_{iC}(x)\Delta q_{it}(x)$ are either positive or negative, depending on the cases:
(c) $f_{it}(x) = \Big(\Delta\mu_{it}(x)\Big)\Big(\Delta q_{it}(x)\Big) + q_{iC}(x)\Big(\Delta\mu_{it}(x)\Big) + \mu_{iC}(x)\Big(\Delta q_{it}(x)\Big) - \Big(\gamma_{it}(x)\mu_{it}(x)q_{it}(x)\Big)d_t$,
(d) $f_{it}(x) = \Big(\Delta\mu_{it}(x)\Big)\Big(\Delta q_{it}(x)\Big) + q_{iC}(x)\Big(\Delta\mu_{it}(x)\Big) + \mu_{iC}(x)\Big(\Delta q_{it}(x)\Big) - \Big(\gamma_{it}(x)\mu_{it}(x)q_{it}(x)\Big)d_t$.



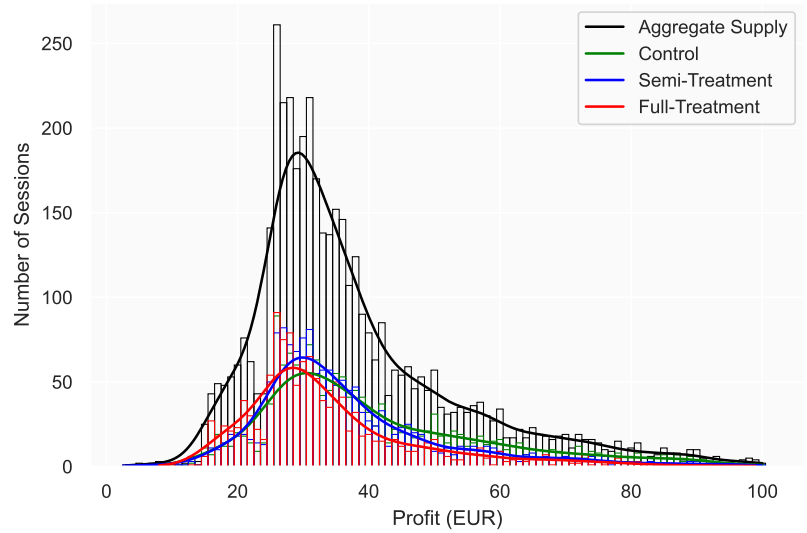
(a) Semi-treatment group



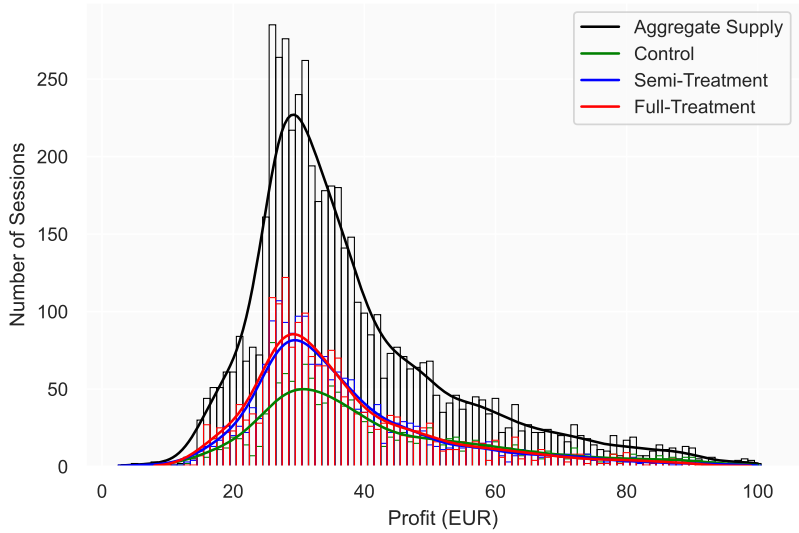
(b) Full-treatment group

Figure D.9: Distributions of optimal policy function $\widehat{f}_{it}(x)$ per treatment

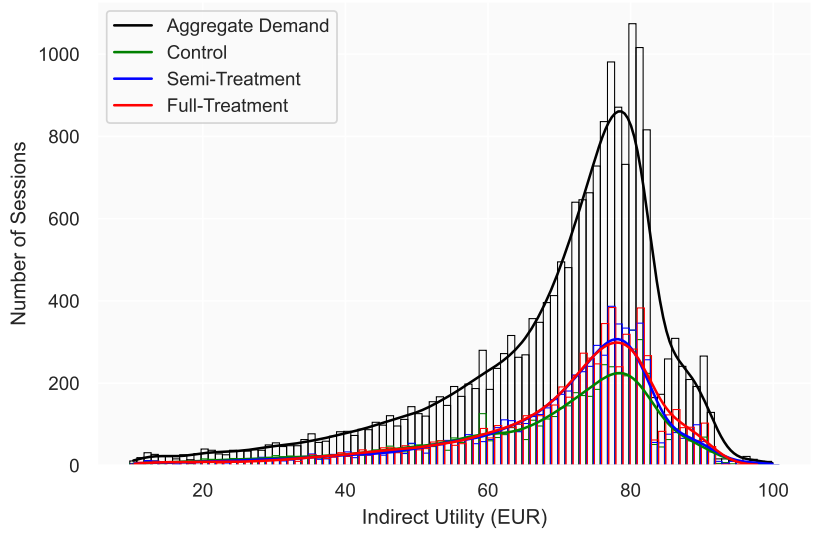
Note: Panels show the distribution of estimated optimal treatment policy function $\widehat{f}_{it}(x)$ for semi-treatment and full-treatment computed though GBDT ensemble method as described in Algorithm 6. The semi-treatment testing sample has 19,908 observations, mean is equal to 0.49, standard deviation 31.43 and median 0.00. Instead, the full-treatment testing sample features 36,748 observations, mean is equal to -7.68, standard deviation 35.03 and median 0.02. Histograms focus on the interval centered at 0 between −100 and +100. Given a testing sample of 19,908 and 36,748 observations of semi-treatment and full-treatment respectively, (a) the semi-treatment plot includes 5,577 negative and 5,415 positive data points, while (b) the full-treatment graph counts 10,570 and 24,857 negative and positive observations, respectively.



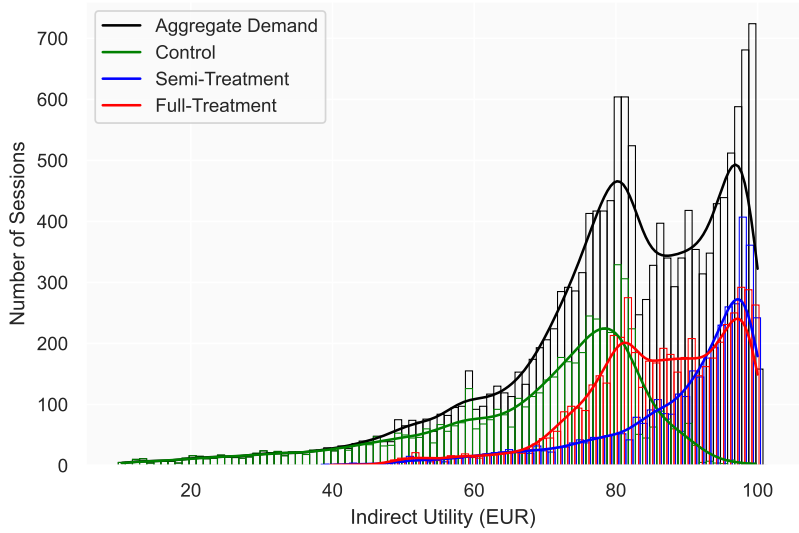
(a) Profit distribution of uniform policy per treatment



(b) Profit distribution of optimal treatment policy per treatment



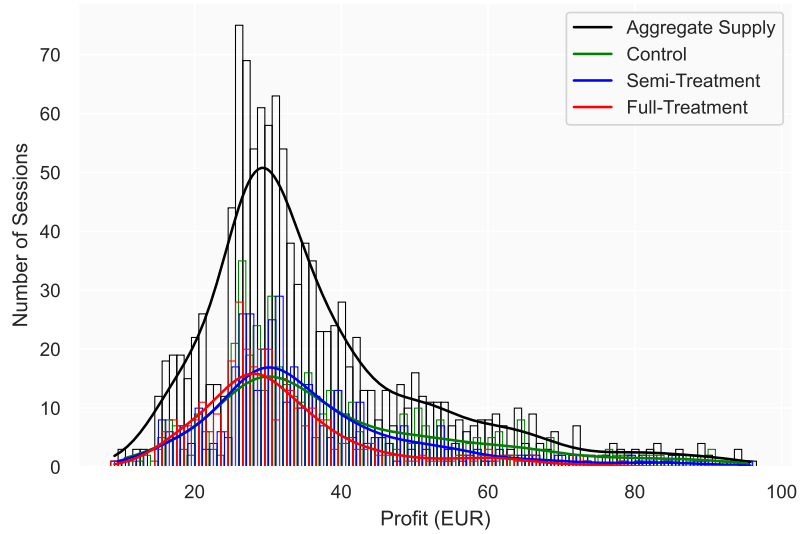
(c) Indirect utility distribution of uniform policy per treatment



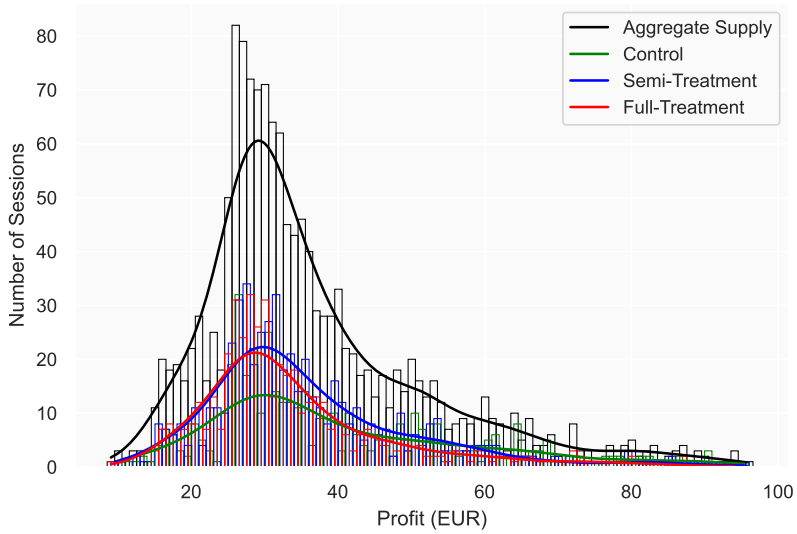
(d) Indirect utility distribution of optimal treatment policy per treatment

Figure D.10: Treatment distributions of profit and indirect utility per policy

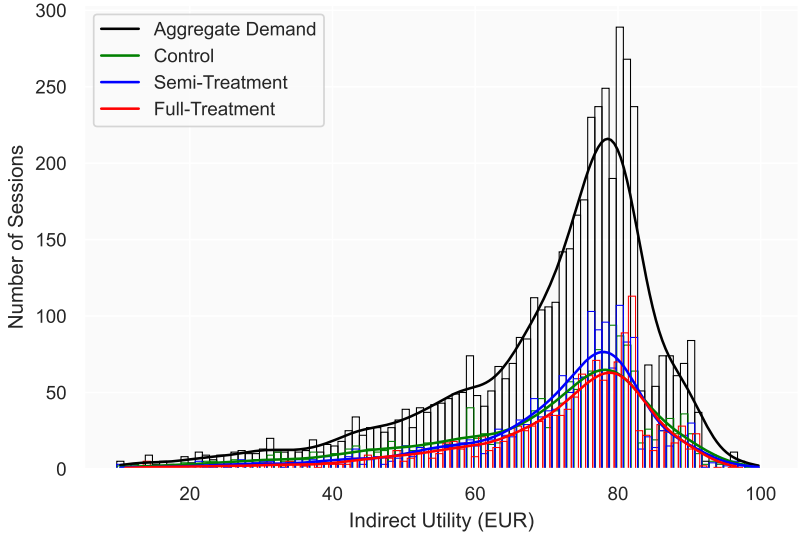
Note: Panels (a) and (b) show the profit distribution of uniform policy and of optimal treatment policy at the aggregate level and at each treatment-group level. Panels (c) and (d) display the indirect utility distribution of uniform policy and of optimal treatment policy at the aggregate level and at each treatment-group level. Sessions with gross order greater than €100 are not depicted. Green plot refers to control group (19,388 observations), blue plot to semi-treatment group (19,388 observations) and red plot to full-treatment group (36,054 observations).



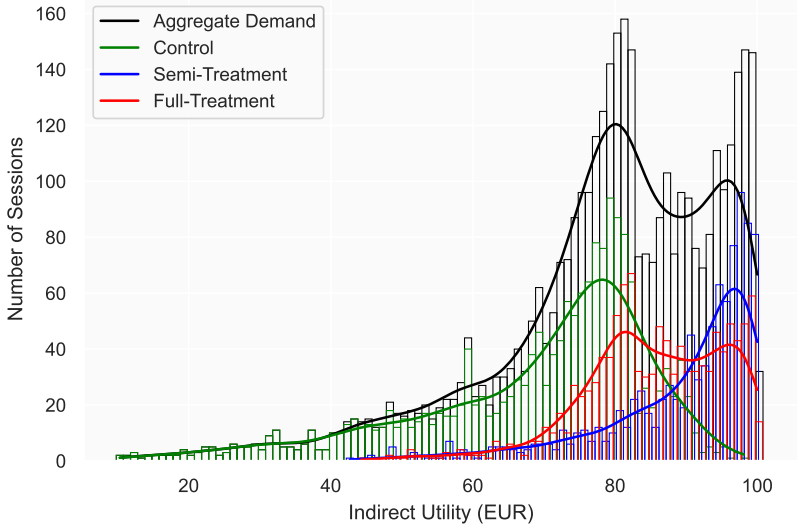
(a) Profit distribution of uniform policy per treatment



(b) Profit distribution of optimal treatment policy per treatment



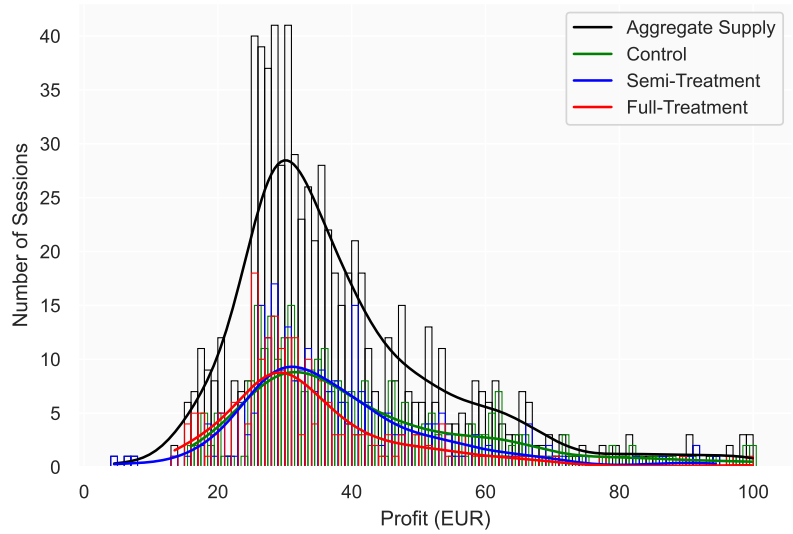
(c) Indirect utility distribution of uniform policy per treatment



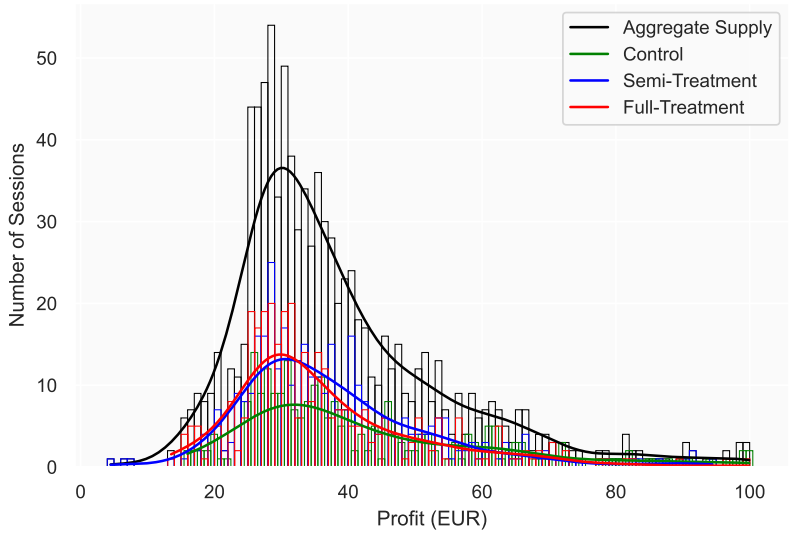
(d) Indirect utility distribution of optimal treatment policy per treatment

Figure D.11: Treatment distributions of profit and indirect utility per policy for new users

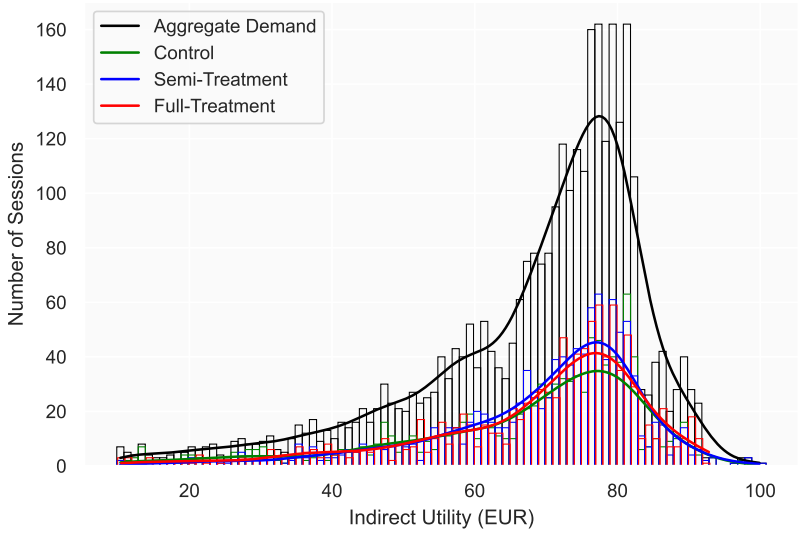
Note: Panels (a) and (b) show the profit distribution of uniform policy and of optimal treatment policy at the aggregate level and at each treatment-group level for new users. Panels (c) and (d) display the indirect utility distribution of uniform policy and of optimal treatment policy at the aggregate level and at each treatment-group level for new users. Sessions with gross order greater than €100 are not depicted. Green plot refers to control group (5,406 observations), blue plot to semi-treatment group (5,406 observations) and red plot to full-treatment group (8,217 observations).



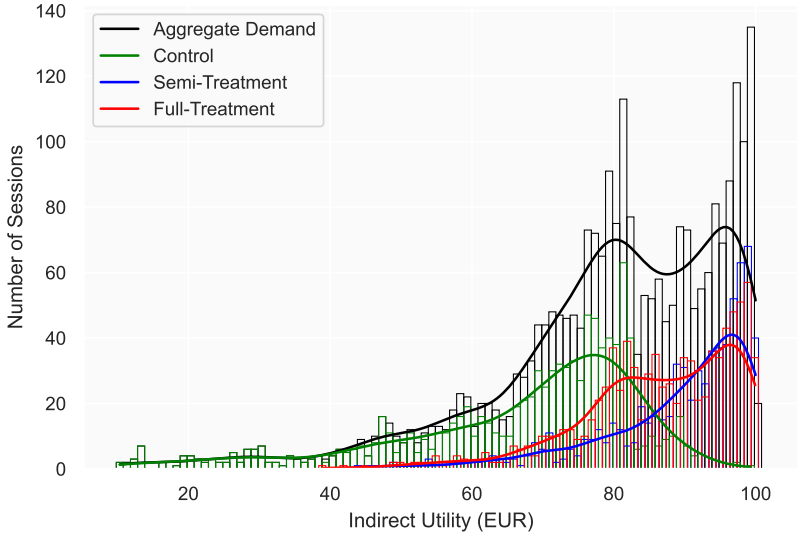
(a) Profit distribution of uniform policy per treatment



(b) Profit distribution of optimal treatment policy per treatment



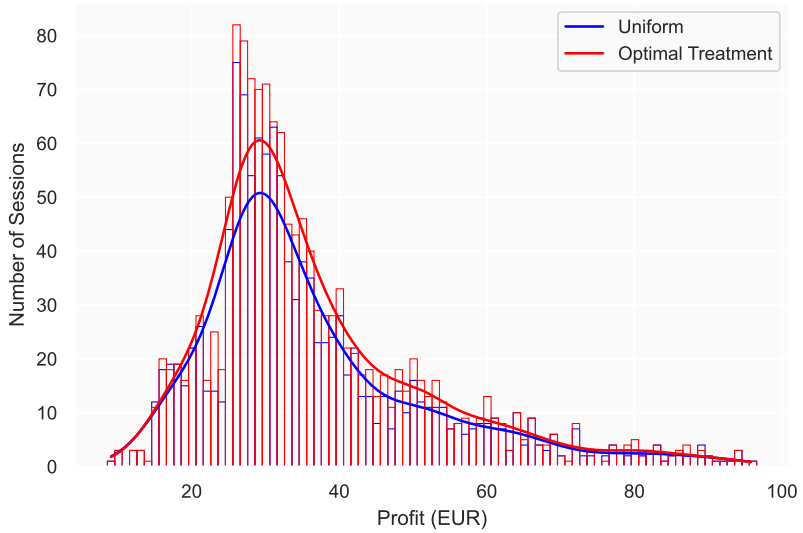
(c) Indirect utility distribution of uniform policy per treatment



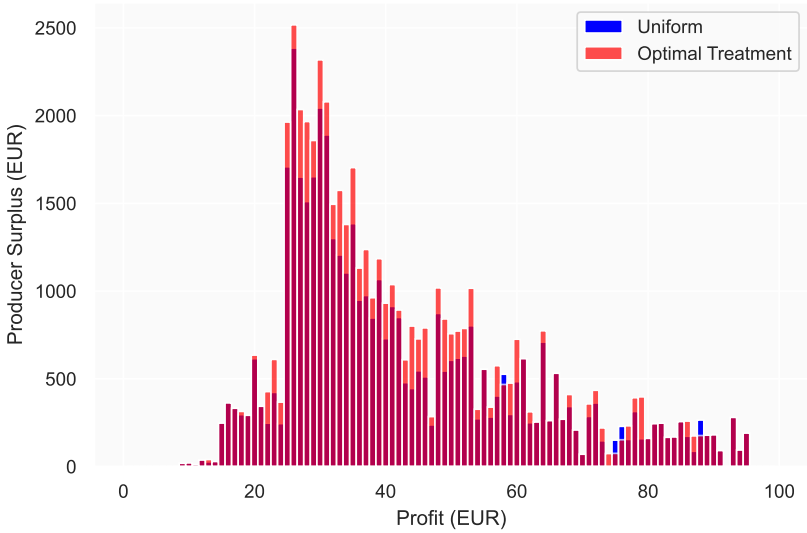
(d) Indirect utility distribution of uniform policy per treatment

Figure D.12: Treatment distributions of profit and indirect utility per policy for returning users

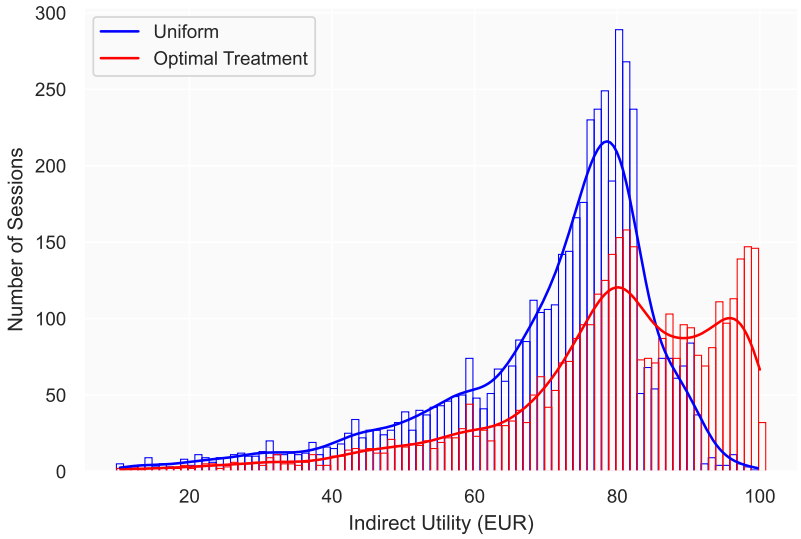
Note: Note: Panels (a) and (b) show the profit distribution of uniform policy and of optimal treatment policy at the aggregate level and at each treatment-group level for returning users. Panels (c) and (d) display the indirect utility distribution of uniform policy and of optimal treatment policy at the aggregate level and at each treatment-group level for returning users. Sessions with gross order greater than €100 are not depicted. Green plot refers to control group (3,030 observations), blue plot to semi-treatment group (3,030 observations) and red plot to full-treatment group (4,931 observations).



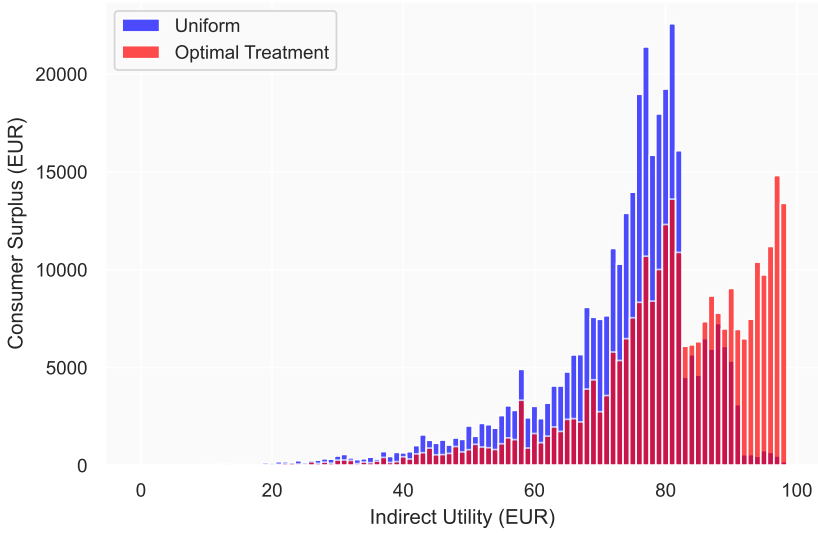
(a) Profit distribution per policy



(b) Producer surplus distribution by profit per policy



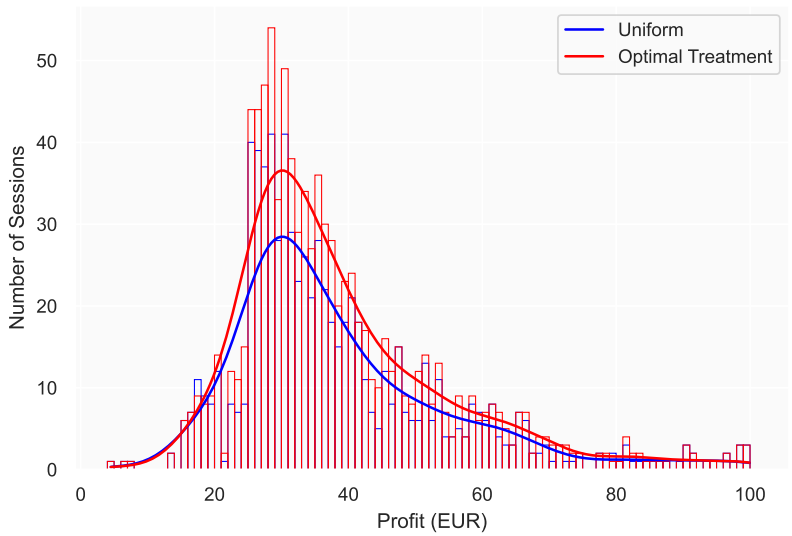
(c) Indirect utility distribution per policy



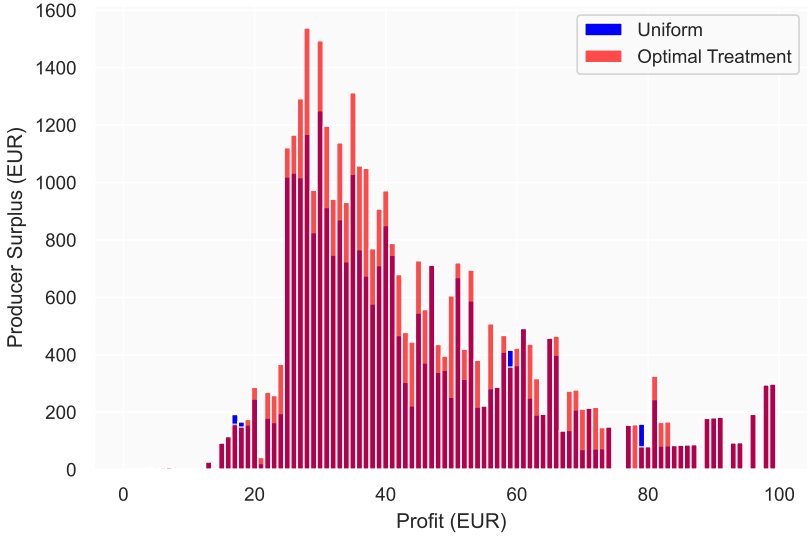
(d) Consumer surplus distribution by indirect utility per policy

Figure D.13: Welfare analysis for new users

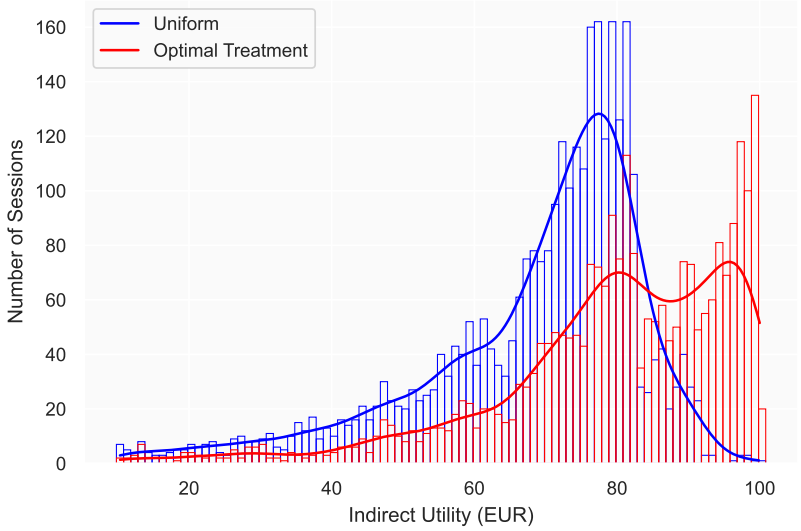
Note: Panel (a) shows the distribution of profit of new users. Panel (b) displays the distribution of producer surplus of new users by profit levels. Panel (c) illustrates the distribution of indirect utility of new users. Panel (d) reports the distribution of consumer surplus of new users by indirect utility levels. For new users, 19,029 sessions with gross order below to €100 are depicted. Blue plot refers to uniform policy and red plot to optimal treatment policy.



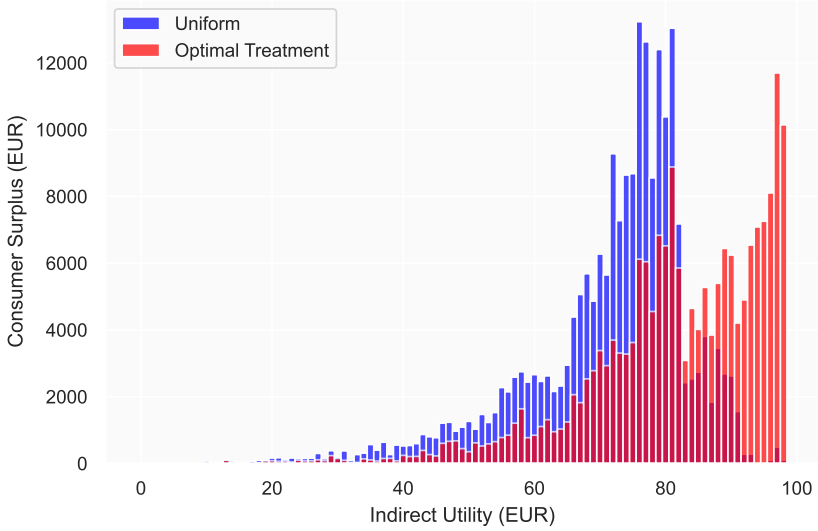
(a) Profit distribution per policy



(b) Producer surplus distribution by profit per policy



(c) Indirect utility distribution per policy



(d) Consumer surplus distribution by indirect utility per policy

Figure D.14: Welfare analysis for returning users

Note: Panel (a) shows the distribution of profit of returning users. Panel (b) displays the distribution of producer surplus of returning users by profit levels. Panel (c) illustrates the distribution of indirect utility of returning users. Panel (d) reports the distribution of consumer surplus of returning users by indirect utility levels. For returning users, 10,991 sessions with gross order below to €100 are depicted. Blue plot refers to uniform policy and red plot to optimal treatment policy.

E Machine Learning and Algorithms

E.1 Ensemble Machine Learning Methods

Ensemble learning is a set of machine learning techniques that uses multiple baseline learners to train an ensemble learner. According to Opitz and Maclin (1999) ensemble learning is divided into ensemble averaging and ensemble boosting. Averaging methods construct several estimators independently and then average their predictions, instead boosting methods built base estimators sequentially and weighted combination of estimators aims at reducing the bias. Bagging (BAG) and random forest (RF) belong to ensemble averaging, while adaptive boosting (AB), gradient boosting (GBDT), and neural network (NN) are examples of ensemble boosting. These methods are discussed and related algorithms are summarized. Mienye and Sun (2022) provides a thorough survey on ensemble learning.

Bagging. Historically the bagging method, literally bootstrap aggregating, was first suggested by Breiman (1996b) to enhance the classification performance of machine learning models by combining the predictions from randomly generated training sets. The BAG estimator is an ensemble meta-algorithm designed to improve both stability and accuracy of the machine learning algorithms used in statistical classification and regression. This ensemble estimator fits base classifiers (regressors) on random subsets of the original dataset and then aggregates their individual predictions to construct a combined outcome. Such a meta-estimator can typically be used in order to reduce the variance of a black-box estimator (e.g., a decision tree or neural network) by means of introducing randomization into its development (Oza & Tumer, 2008). Basically, the bagging method involves splitting the training data for each base learner using random sampling to generate b different subsets used to train b base learners. The b base learners are then combined using majority voting to obtain a strong classifier. The most significant advantage comes from efficiently decreasing the variance without increasing bias (Algorithm 1). Further advantages of BAG include the ability to introduce diversity in the input data because of the bootstrapping approach (Kotsiantis, 2014). Random forests are a standard implementation of the bagging technique.

Random forest. A RF is a meta-estimator that fits a number of decision tree classifiers on different subsamples of the dataset and utilize averaging to improve the predictive accuracy and to control overfitting. In random forests each tree in the ensemble is built from a bootstrap drawn from the training set. For classification, the output of the random forest is the class selected by most trees. For regression, the average prediction of the individual trees is returned. In the pioneering work by Ho (1995), random forests algorithm uses bagging technique to build multiple decision trees by means of bootstrapped samples. The bagging technique generates random samples with replacements from the input data and trains the decision trees from the samples (Algorithm 2). A main advantage of leveraging this algorithm is its power to solve the overfitting issue usual in decision tree models. Furthermore, when input dataset features high dimensional data, prediction performances of random decision forests tend to improve significantly (e.g., Han & Kim, 2019; Schonlau & Zou, 2020).

Adaptive boosting. The adaptive boosting, also AdaBoost, algorithm is a type of boosting algorithm capable of using weak learners to obtain a robust estimator. This boosting method was developed by

Freund and Schapire (1997) and it is among the most efficient ML algorithms. An AB classifier (regressor) is a statistical meta-estimator that begins by fitting a classifier (regressor) on the original dataset. Thus, AdaBoost fits additional copies of the classifier (regressor) on the same dataset, but where the weights of incorrectly classified instances are adjusted such that subsequent estimators focus more on difficult cases. The working principle of the algorithm is to fit a sequence of weak learners on repeatedly modified versions of the data. All predictions are then combined through a weighting rule to produce the aggregate prediction. The data modifications at each boosting iteration consists of applying weights $w_1, \ldots, w_N$ to each of the training samples. Initially, those weights are all set to $w_i = 1/N$, so that the first step simply trains a weak learner on the original data. For each successive iteration, the sample weights are individually modified and the learning algorithm is reapplied to the reweighted data. At a certain step, those training examples that were incorrectly predicted by the boosted model induced at the previous step have their weights increased, whereas the weights are decreased for those that were predicted correctly. Each subsequent weak learner is thereby forced to concentrate on the examples that are missed by the previous ones in the sequence (Algorithm 3) (F. Wang et al., 2019; C. Wang et al., 2021). The AB is flexible, relatively easy to implement, and accommodates a variety of algorithms as the base learner.

Gradient boosting. The gradient boosting is a ML algorithm that uses the boosting technique to create strong ensembles. It is also called gradient boosted decision tree since it mainly adopts decision trees as the base learner to produce a robust ensemble classifier or regressor. GBDT was first introduced by Breiman (1996a) and then developed by the seminal work of Friedman (2002). Gradient boosting for classification (regression) is an algorithm that builds an additive model in a forward stage-wise fashion. Indeed, it allows for the optimization of arbitrary differentiable loss functions and, at each stage, N classes of trees are fit on the negative gradient of the loss function. Specifically, the learning process of this algorithm involves sequentially training new models to obtain a robust classifier (regressor). It is built in a marginal manner, similarly to other boosting techniques, but its core intuition consists of developing base learners that are highly correlated with the negative gradient of the loss function related to the entire ensemble (Algorithm 4). The major advantage of GBDT is that it can learn complex patterns from the input data since it is trained to correct the errors of the previous model.

Neural network. A multilayer perceptron (MLP) is a feed-forward artificial neural network, consisting of fully connected neurons with a nonlinear kind of activation function, organized in at least three layers, and able to distinguish data that is not linearly separable. Mathematically, multilayer perceptron is a supervised learning algorithm that learns a function $f(\cdot) : \mathbb{R}^n \rightarrow \mathbb{R}^o$ by training on a data set, where n is the number of dimensions for input and o is the number of dimensions for output. Given a set of features $X = x_1, x_2, \ldots, x_n$ and a target y, it can learn a non-linear function approximation for either classification or regression. Between the input and the output layer, there can be one or more non-linear layers, called hidden layers (Algorithm 5). The important advantage of NN consists in learning non-linear models (e.g. Tkáč & Verner, 2016; Hospedales et al., 2021; Shobana & Umamaheswari, 2021).

E.2 Ensemble Learning: Bagging, Random Forest, Adaptive Boosting, Gradient Boosting and Neural Network

Algorithm 1 Bagging

Input: $S = \{(x_1, y_1), \ldots, (x_n, y_n)\}$ training data; L base ML algorithm; T number of base learners
Procedure:
for $t = 1, \ldots, T$ **do**
(i) Generate a bootstrap sample S_j from the training data S
(ii) Fit a base learner h_j using S_j, i.e. $h_j = L(S_j)$
(iIi) Combine the outputs of the base learners $H(x) = \text{mode}(h_1(x), \ldots, h_T(x))$
end for
Output: Return the bagging final learner $H(x)$

Algorithm 2 Random forest

Input: $S = \{(x_1, y_1), \ldots, (x_n, y_n)\}$ training data; p attributes and class variables; T number of trees in the forest; C number of class labels
Procedure:
for $t = 1, \ldots, T$ **do**
(i) Generate a bootstrap sample S_j from the training data S
(ii) Fit a base learner h_j using S_j and, for a given node n,
(a) Randomly select k attributes (usually $k = \sqrt{p}$)
(b) Compute the best split features using the randomly selected feature subset
(c) Split the node using the optimal split features obtained in (b)
(d) Repeat (a), (b) and (c) until the stopping criteria is achieved
(iii) Repeat (i) and (ii) for T times to build a forest of T trees
(iv) For a given test sample x, combine the outputs from the T trees in the final predicted class label, i.e. $H_T(x) = argmax_j \sum_{k=1}^{K} \mathbb{1}_{\{h_k(x)=j\}} \quad \forall j = 1, \ldots, C$
end for
Output: Return the random forest final predicted class label $H_T(x)$

Algorithm 3 Adaptive boosting

Input: $S = \{(x_1, y_1), \ldots, (x_n, y_n)\}$ training data; L base ML algorithm; $\{D, \alpha\}$ set of vector of weights; M number of iterations

Procedure: Initialize the model with initial vector D_0 using $D_0(i) = \frac{1}{n} \quad i = 1, \ldots, n$

for $m = 1, \ldots, M$ **do**

(i) Generate a bootstrap sample S_j from the training data S

(ii) Train the base classifier $h_m(x)$ by minimizing $\epsilon_m = \sum_{i=1}^{N} D_m(i) \mathbb{1}_{\{h_m(x_i) \neq y_i\}}$

(iii) Calculate the weights α_m of the classifier $h_m(x)$ using $\alpha_m = \frac{1}{2} \ln(\frac{1-\epsilon_m}{\epsilon_m})$

(iv) Update the vector of weights $D_{m+1}(i) = \frac{D_m(i)}{z_m} \exp\Big(- \alpha_m y_i h_m(x_i) \Big) \quad i = 1, \ldots, n$

(v) Apply $H(x) = \text{sgn}\Big(\sum_{m=1}^{M} \alpha_m h_m(x) \Big)$

end for

Output: Return the adaptive boosting final classifier $H(x)$

Algorithm 4 Gradient boosting

Input: $S = \{(x_1, y_1), \ldots, (x_n, y_n)\}$ training data; $L(y, F(x))$ specified differential loss function; M number of iterations

Procedure: Initialize the model with constant Γ using $F_0 = argmin_{\Gamma} \sum_{i=1}^{N} L(y_i, \Gamma)$

for $m = 1, \ldots, M$ **do**

(i) Generate a bootstrap sample S_j from the training data S

(ii) Calculate the pseudo-residuals $r_{im} = \left[\frac{\partial L(y_i, F(x_i))}{\partial F(x_i)} \right]_{F(x) = F_{m-1}(x)} \quad \forall i = 1, \ldots, n$

(iii) Train a base learner closed under scaling $h_m(x)$ using the training set $D = \Big(x_i, r_{im} \Big)_{i=1}^{N}$

(iv) Compute multiplier ρ_m by performing the line search optimization $\rho_m h_m(x) = argmin_{\rho, h} \sum_{i=1}^{N} L(y_i, F_{m-1}(x_i) + \rho h_m(x_i))$

(iv) Update the model $F_m(x) = F_{m-1}(x) + \rho_m h_m(x)$

end for

Output: Return the gradient boosting final model $F_m(x)$

Algorithm 5 Neural network

Input: $S = \{(x_1, y_1), \ldots, (x_n, y_n)\}$ training data; $L(y, F(x), W)$ specified differential loss function; W vector of weights for respective layers; $\{b_1, b_2\}$ model parameters, η learning rate parameter; L number of layers; M number of iterations

Procedure: Initialize the model with initial vector W_0 and constant Γ using $F_0 = argmin_{\Gamma, W_0} \sum_{i=1}^{N} L(y_i, \Gamma, W_0)$

for $m = 1, \ldots, M$ **do**

(i) Generate a bootstrap sample S_j from the training data S

(ii) Compute the vector of weights $W_m(x)$ by minimizing $L_m(y_i, F(x_i), W_m)$

(iii) Update the vector of weights $W_m = W - \eta \left[\frac{\partial L(y_i, F(x_i), W)}{\partial W} \right]_{W = W_{m-1}} \quad \forall i = 1, \ldots, n$

(iv) Update the model $F_m(x) = W_m g(W_m^T x + b_1) + b_2$

end for

Output: Return the neural network final model $F_m(x)$

E.3 Estimation of Primitives of the Policy Model

Algorithm 6 Pseudocode for estimation of optimal treatment policy and counterfactual policies

Input:

$\mathcal{S} = \{y_i, c_i, q_i, t_i, x_i\}_{i=1}^{I}$ representative sample of RCTs dataset with observed conversion decision y_i, MPT_t compliance response c_i, realized gross order q_i, assigned treatment t_i, and covariates matrix x_i of I users in the experiment;

$X = \{ts_i, p_i, ac_i, tc_i, u_i, h_i, ns_i, npv_i, npp_i, npac_i\}_{i=1}^{I}$ covariates matrix of RCTs dataset with time to session ts_i, number of pages p_i, attained cart dummy ac_i, time to cart tc_i, new user dummy u_i, purchase history dummy h_i, number of sessions ns_i, number of prior visits npv_i, number of prior purchases npp_i, number of prior abandoned carts $npac_i$;

$\mathcal{T} = \{C, ST, FT\}$ treatment space with control set $t = C \in \{0, 0\}$, semi-treatment set $t = ST \in \mathbb{R}^+ \times \{0\}$, and full-treatment set $t = FT \in \mathbb{R}^+ \times \mathbb{R}^+$;

ξ hyperparameters set for ensemble method calibration

Output:

$\widehat{f}_{it}(x)\ \forall\ t \in \mathcal{T}_{-\{C\}}$, estimated treatment policy functions associated to optimal treatment policy and counterfactual policies

Procedure:

function ESTIMATE RESPONSE FUNCTIONS ($\mathcal{S}, \xi$)
  for Treatment pair $t \in \mathcal{T}$ **do**
    $\widehat{\mu}_{it}(x) \longleftarrow$ train selected supervised ensemble classifier method
    $\widehat{q}_{it}(x) \longleftarrow$ train selected supervised ensemble regressor method
    $\widehat{\gamma}_{iFT}(x) \longleftarrow$ train selected supervised ensemble classifier method for $t = FT$
  end for
  return Estimated response functions $\{\widehat{\mu}_{iC}(x), \widehat{\mu}_{iST}(x), \widehat{\mu}_{iFT}(x), \widehat{q}_{iC}(x), \widehat{q}_{iST}(x), \widehat{q}_{iFT}(x), \widehat{\gamma}_{iFT}(x)\}$
end function

function ESTIMATE CATEs ($\widehat{\mu}_{iC}(x), \widehat{\mu}_{iST}(x), \widehat{\mu}_{iFT}(x), \widehat{q}_{iC}(x), \widehat{q}_{iST}(x), \widehat{q}_{iFT}(x), \widehat{\gamma}_{iFT}(x)$)
  for Treatment pair $t \in \mathcal{T}_{-\{C\}}$ **do**
    $\widehat{\Delta\mu}_{it}(x) \longleftarrow \widehat{\mu}_{it}(x) - \widehat{\mu}_{iC}(x)$
    $\widehat{\Delta q}_{it}(x) \longleftarrow \widehat{q}_{it}(x) - \widehat{q}_{iC}(x)$
  end for
  return Estimated CATE functions $\{\widehat{\Delta\mu}_{iST}(x), \widehat{\Delta q}_{iST}(x), \widehat{\Delta\mu}_{iFT}(x), \widehat{\Delta q}_{iFT}(x)\}$
end function

▹ Pseudocode continues

▷ Pseudocode continues

function ESTIMATE TREATMENT POLICY FUNCTIONS $(\mathcal{S}, \xi)$

$\{\widehat{\mu}_{iC}(x), \widehat{\mu}_{iST}(x), \widehat{\mu}_{iFT}(x), \widehat{q}_{iC}(x), \widehat{q}_{iST}(x), \widehat{q}_{iFT}(x), \widehat{\gamma}_{iFT}(x)\} \longleftarrow$ `Estimate Response Functions` $(\mathcal{S}, \xi)$

$\{\widehat{\Delta\mu}_{iST}(x), \widehat{\Delta q}_{iST}(x), \widehat{\Delta\mu}_{iFT}(x), \widehat{\Delta q}_{iFT}(x)\} \longleftarrow$ `Estimate CATEs` $(\widehat{\mu}_{iC}(x), \widehat{\mu}_{iST}(x), \widehat{\mu}_{iFT}(x), \widehat{q}_{iC}(x), \widehat{q}_{iST}(x), \widehat{q}_{iFT}(x)$

Non-treatment policy

$\widehat{f}_{iST}(x) \longleftarrow -1$

$\widehat{f}_{iFT}(x) \longleftarrow -1$

Non-targeted treatment policy

$\widehat{f}_{iST}(x) \longleftarrow \overline{\Delta\mu}_{iST}(x)\overline{\Delta q}_{iST}(x) + \bar{q}_{iC}(x)\overline{\Delta\mu}_{iST}(x) + \bar{\mu}_{iC}(x)\overline{\Delta q}_{iST}(x) - \bar{\mu}_{iST}(x)\bar{q}_{iST}(x)d_{ST}$

$\widehat{f}_{iFT}(x) \longleftarrow \overline{\Delta\mu}_{iFT}(x)\overline{\Delta q}_{iFT}(x) + \bar{q}_{iC}(x)\overline{\Delta\mu}_{iFT}(x) + \bar{\mu}_{iC}(x)\overline{\Delta q}_{iFT}(x) - \bar{\gamma}_{iFT}(x)\bar{\mu}_{iFT}(x)\bar{q}_{iFT}(x)d_{FT}$

Conversion-targeted treatment policy

$\widehat{f}_{iST}(x) \longleftarrow \widehat{\Delta\mu}_{iST}$

$\widehat{f}_{iFT}(x) \longleftarrow \widehat{\Delta\mu}_{iFT}$

Profit-targeted treatment policy

$\widehat{f}_{iST}(x) \longleftarrow \widehat{\Delta q}_{iST}$

$\widehat{f}_{iFT}(x) \longleftarrow \widehat{\Delta q}_{iFT}$

Double-positive treatment treatment policy

$\widehat{f}_{iST}(x) \longleftarrow \left(\mathbb{1}_{\{\widehat{\Delta\mu}_{it}(x)>0\}}\widehat{\Delta\mu}_{it}(x)\right)\left(\mathbb{1}_{\{\widehat{\Delta q}_{it}(x)>0\}}\widehat{\Delta q}_{it}(x)\right)$

$\widehat{f}_{iFT}(x) \longleftarrow \left(\mathbb{1}_{\{\widehat{\Delta\mu}_{it}(x)>0\}}\widehat{\Delta\mu}_{it}(x)\right)\left(\mathbb{1}_{\{\widehat{\Delta q}_{it}(x)>0\}}\widehat{\Delta q}_{it}(x)\right)\widehat{\gamma}_{it}(x)$

Optimal treatment policy

$\widehat{f}_{iST}(x) \longleftarrow \widehat{\Delta\mu}_{iST}(x)\widehat{\Delta q}_{iST}(x) + \widehat{q}_{iC}(x)\widehat{\Delta\mu}_{iST}(x) + \widehat{\mu}_{iC}(x)\widehat{\Delta q}_{iST}(x) - \widehat{\mu}_{iST}(x)\widehat{q}_{iST}(x)d_{ST}$

$\widehat{f}_{iFT}(x) \longleftarrow \widehat{\Delta\mu}_{iFT}(x)\widehat{\Delta q}_{iFT}(x) + \widehat{q}_{iC}(x)\widehat{\Delta\mu}_{iFT}(x) + \widehat{\mu}_{iC}(x)\widehat{\Delta q}_{iFT}(x) - \widehat{\gamma}_{iFT}(x)\widehat{\mu}_{iFT}(x)\widehat{q}_{iFT}(x)d_{FT}$

return Estimated treatment policy functions $\{\widehat{f}_{iST}(x), \widehat{f}_{iFT}(x)\}$ for optimal treatment policy and counterfactual policies.

end function