

Business analytics meets artificial intelligence: Assessing the demand effects of discounts on Swiss train tickets

Martin Huber ^{a,1,2}, Jonas Meier ^{b,3,2}, Hannes Wallimann ^{c,*}

^a University of Fribourg, Department of Economics, Switzerland

^b University of Amsterdam, Faculty of Economics and Business, Netherlands

^c University of Applied Sciences and Arts Lucerne, Competence Center for Mobility, Switzerland

ARTICLE INFO

JEL classification:

C21

R41

R48

Keywords:

Causal machine learning

Double machine learning

Treatment effect

Business analytics

Causal forest

Public transportation

ABSTRACT

We assess the demand effects of discounts on train tickets issued by the Swiss Federal Railways, the so-called ‘supersaver tickets’, based on machine learning, a subfield of artificial intelligence. Considering a survey-based sample of buyers of supersaver tickets, we use causal machine learning to assess the impact of the discount rate on rescheduling a trip, which seems relevant in the light of capacity constraints at rush hours. Assuming that (i) the discount rate is quasi-random conditional on our rich set of characteristics and (ii) the buying decision increases weakly monotonically in the discount rate, we identify the discount rate’s effect among ‘always buyers’, who would have traveled even without a discount, based on our survey that asks about customer behavior in the absence of discounts. We find that on average, increasing the discount rate by one percentage point increases the share of rescheduled trips by 0.16 percentage points among always buyers. Investigating effect heterogeneity across observables suggests that the effects are higher for leisure travelers and during peak hours when controlling several other characteristics.

1. Introduction

Organizing public transport involves a well-known trade-off between consumer welfare and provider revenue. Typically, consumers value frequency, reliability, space, and low fares (Redman et al., 2013) while suppliers aim at operating with a minimum number of vehicles to maximize profits. In general, the allocation can be improved as providers do not account for the positive externalities on consumers (Mohring, 1972). In particular, service frequency reduces travelers’ access and waiting costs. This so-called ‘Mohring-effect’ leads to economies of scale, implying the need for subsidies to achieve the first-best solution in terms of welfare. Consequently, it may be socially optimal to subsidize railway companies to reduce fares (Parry and Small, 2009). To assess such a measure’s effectiveness on demand, policy-makers would need to know how individuals respond to lower fares. However, it is generally challenging to identify causal effects of discounts on train tickets (or goods and services in general) due to confounding or selection. For instance, discounts might typically be provided for dates or hours with low train utilization such that connections with and without discount are not comparable in terms of baseline demand. A naive comparison of sold tickets with and without discount would therefore mix the influence of the discount with that of baseline demand. In this context, we apply machine learning (a

* Correspondence to: University of Applied Sciences and Arts Lucerne, Rösslmatte 48, 6002 Luzern, Switzerland.

E-mail addresses: martin.huber@unifr.ch (M. Huber), j.c.meier@uva.nl (J. Meier), hannes.wallimann@hslu.ch (H. Wallimann).

¹ University of Fribourg, Bd. de Pérolles 90, 1700 Fribourg, Switzerland.

² All authors have contributed in multiple roles to the paper.

³ University of Amsterdam, Roeterstraat 11, 1018WB Amsterdam, Netherlands.

<https://doi.org/10.1016/j.trb.2022.06.006>

Received 27 July 2021; Received in revised form 17 May 2022; Accepted 15 June 2022

Available online 6 July 2022

0191-2615/© 2022 The Author(s). Published by Elsevier Ltd. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

subfield of artificial intelligence) to convincingly assess how discounts on train tickets for long-distance connections in Switzerland, the so-called ‘supersaver tickets’, affect demand, by exploiting a unique data set of the Swiss Federal Railways (SBB) that combines train utilization records with a survey of supersaver buyers.

More concisely, our study provides a promising use case of causal machine learning for business analytics in the railway industry: analyzing the causal effect of the discount on demand shifts among customers that would have booked the trip even without a discount. This is feasible because our unique survey contains information on how supersaver buyers would have decided in the absence of a discount, e.g. whether they are so-called ‘always buyers’ and would have booked the connection even at the regular fare. Using causal machine learning approaches, this novel survey allows us to answer the question of whether discounts causally affect individual buying decisions — an important and under-researched question. Moreover, and to the best of our knowledge, this is the first application of causal machine learning techniques in the public transportation literature.

For both our descriptive and causal analysis, we make use of appropriately tailored machine learning techniques, which learns the associations between the demand outcomes of interest, the discount rate, and further customer or trip-related characteristics in a data-driven way and helps avoiding model misspecification. Predictive machine learning permits optimizing forecasts about demand and customer behavior as a function of observed characteristics in order to descriptively analyze specific patterns in customer behavior. On the other hand, causal machine learning permits evaluating the causal effect of specific interventions like a discount regime for optimizing the offer of such discounts. When applying predictive machine learning, trip-related characteristics like seat capacity, utilization, departure time, and the discount rate, but also customer’s age turn out to be strong predictors.

Concerning the causal analysis, we impose (i) a selection-on-observables assumption implying that the discount rate is as good as randomly assigned when controlling for our rich set of trip- and demand-related characteristics and (ii) weak monotonicity of any individual’s decision to purchase an additional trip (otherwise not realized) in the discount rate, implying that a higher (rather than lower) discount does either positively or not affect any customer’s buying decision. We formally show how these assumptions permit tackling the selectivity of discount rates and survey response to identify the discount rate’s effect on demand shifts (rescheduling away from rush hours) for the subgroup ‘always buyers’, based on the survey information on how customers would have behaved in the absence of a discount. Furthermore, we discuss testable implications of monotonicity, namely that among all survey respondents, the share of additional trips must increase in the discount rate, and the selection on observables assumptions, requiring that conditional on trip- and demand-related characteristics, the discount must not be associated with personal characteristics (like age or gender) among always buyers. Hypothesis tests do not point to the violation of these implications.

Based on our causal identification strategy, we estimate the marginal effect of slightly increasing the (continuously distributed) discount rate based on the causal forest (CF), see [Wager and Athey \(2018\)](#) and [Athey et al. \(2019\)](#), and find that on average, increasing the discount rate by one percentage point increases the share of rescheduled trips by 0.16 percentage points among always buyers. In a second approach, we binarize the discount rates by splitting them into two discount categories of less than 30% (relative to the regular fare) and 30% or more. Applying double machine learning (DML), see [Chernozhukov et al. \(2018\)](#), we find that discount rates of 30% and more on average increase the share of rescheduled trips 3.6 percentage points, which is in line with the CF-based results. Our paper, therefore, provides the first empirical evidence (at least for Switzerland) that such discounts can help balancing out train utilization across time and reducing overload during peak hours, albeit the magnitude of the impact on always buyers appears limited.⁴ We benchmark our estimated effects based on linear regression and propensity score matching. The findings based on the latter approaches point in the same direction as our main results. Yet, the results become less significant and for propensity score matching, this goes together with an increase in the variance. This may be due to overfitting, given our relatively high number of potential control variables. Therefore, in similar scenarios as ours, causal machine learning may appear preferable, as the data-driven selection of important control variables (rather than using any of them) reduces the risk of overfitting. Finally, our study also highlights the advantages of flexibly including influential control variables among many observed covariates - a task for which causal machine learning is well suited.

When investigating the heterogeneity of effects across all of our observed characteristics using the CF, our results suggest that demand-related trip characteristics (like seat capacity, utilization, departure time, and distance) have some predictive power for the size of the discounts’ impact on shifting demand. Such information on heterogeneous effects appears interesting for optimizing the allocation of discounts for the purpose of shifting demand, as the SBB has (due to its monopoly in the Swiss long-distance passenger rail market) agreed with the Swiss price monitoring agency to provide a fixed amount of discounted tickets per year, but is free to choose the timing and connections for discounts. In a second heterogeneity analysis, we investigate whether effects differ systematically across a pre-selected set of characteristics, namely: age, gender, possession of a half fare travel card, travel distance, whether the purpose is business, commute, or leisure, and whether the departure time is during peak hours. Using the regression approach of [Semenova and Chernozhukov \(2021\)](#), we find that conditional on the other characteristics, the effects of increasing the discount by one percentage point on rescheduling are by more than 0.2 percentage points higher during peak hours and for leisure travelers, differences that are statistically significant at the 10% level when, however, not controlling for multiple hypothesis testing. These effects appear plausible as leisure travelers are likely more flexible and discounts during peak hours make trips at times of increased demand even more attractive. We do not find statistically significant effect differences for the other pre-selected

⁴ We note that the CF and DML estimators are but two approaches in the emerging and fast-developing literature on combining causal inference and machine learning, see for instance causal bayesian networks (e.g. [Heckerman, 2013](#)) or probabilistic graphical models (e.g. [Sucar, 2015](#)) for further examples. For surveys on recent developments in causal analysis based on machine learning and algorithms, we refer the reader to [Athey and Imbens \(2019\)](#) and [Peters et al. \(2017\)](#).

characteristics, which could, however, be due to the (for the purpose of investigating effect heterogeneity) limited sample of several thousand observations.

Our paper is related to a growing literature applying statistical and machine learning methods for analyzing transport systems, as well as to methodological studies on causal inference for so-called principal strata, see [Frangakis and Rubin \(2002\)](#), i.e. endogenous subgroups like the always buyers. Typically, it is hard to identify the causal effect of some treatment (or intervention) like a discount on such a non-randomly selected subgroup defined in terms how a post-treatment variable (e.g. buying decision) depends on the treatment (e.g. treatment). One approach is to give up on point identification and instead derive upper and lower bounds on a set of possible effects for groups alike the always buyers based on the aforementioned monotonicity assumption (and possibly further assumptions about the ordering of outcomes of always buyers and other individuals), see for instance ([Zhang and Rubin, 2003](#)), [Zhang et al. \(2008\)](#), [Imai \(2008\)](#), [Lee \(2009\)](#), and [Blanco et al. \(2011\)](#). Alternatively, the treatment effect on always buyers is point-identified when invoking a selection-on-observables or instrumental variable assumption for selection into the survey, see for instance ([Huber, 2014](#)), which requires sufficiently rich data on both survey participants and non-participants for modeling survey participation. In contrast to these previous studies, the approach in this paper point-identifies the treatment effect by exploiting the rather unique survey feature that customers were asked about their behavior in the absence of the discount, which under monotonicity permits identifying the principal stratum of always buyers directly in the data.

Furthermore, our work is related to conceptual studies on transport systems, considering for instance the previously mentioned positive externalities of an increased service for customers that are not accounted for by transportation providers. Such externalities typically arise from economies of scale due to fixed costs and a 'Mohring effect', implying that service frequency reduces waiting costs ([Mohring, 1972](#)). The study by [Parry and Small \(2009\)](#) suggests that lower fares can boost overall welfare by increasing economies of scale (off-peak) and decreasing pollution and accidents (at peaks). Similarly, [De Palma et al. \(2017\)](#) argue that time-dependent ticket prices may increase overall welfare as overcrowding during peak hours is suboptimal for both consumers and providers. As public transport is usually highly subsidized, governments may directly manage the trade-off mentioned above. As this involves taxpayer money, it is a question of general interest how the subsidies should be designed. Based on their results, [Parry and Small \(2009\)](#) conclude that even substantial subsidies are justified due to lower fares' positive welfare effect. In contrast, [Basso and Silva \(2014\)](#) find that the contribution of transit subsidies to welfare diminishes once congestion is taxed and alternatives are available, i.e., bus lanes. Irrespective of the specific policy instrument, the consumer's willingness to shift demand drives these policies' effectiveness. While many factors affect this willingness, most studies conclude that consumers are price sensitive ([Paulley et al., 2006](#)). In this context, we aim at contributing to a better understanding of how time-dependent pricing translates to consumer decisions. In particular, we shed light on the effects of a continuously distributed discount. The recent literature typically considers binary or categorical treatments (e.g., [Kholodov et al., 2021](#); [Shin, 2021](#); [Wallimann et al., 2021](#)), whereas studies on continuous treatments are rare in the public transportation literature.

More broadly, our paper relates to the literature on policies targeting demand shifts. Among these, the setting of car parking costs, fiscal regulations, or even free public transport has been analyzed (e.g. [Batty et al., 2015](#); [Rotaris and Danielis, 2014](#); [Zhang et al., 2018](#); [De Witte et al., 2006](#)). Another stream of literature applies machine learning algorithms in the context of public transport. Examples are short-term traffic flow forecasts for bus rapid transit ([Liu and Chen, 2017](#)) or metro ([Liu et al., 2019](#)) services. Further, [Hagenauer and Helbich \(2017\)](#) and [Omrani \(2015\)](#) implement machine learning algorithms to predict travel mode choices. [Yap and Cats \(2020\)](#) predict disruptions and their passenger delay impacts for public transport stops. In other research fields, also applications of causal (rather than predictive) machine learning are on the rise (see for instance [Yang et al., 2020](#); [Knaus, 2021](#)). This is, to the best of our knowledge, the first study using causal machine learning in the context of public transport. Finally, a growing literature discusses the opportunities of data-driven business decision-making ([Brynjolfsson and McElheran, 2016](#)) by assessing the relevance of predictive and causal machine learning. [Ascarza \(2018\)](#) and [Hünemann et al. \(2021\)](#) show that companies may gain by designing their policies based on causal machine learning. For instance, firms can target the relevant consumers much more effectively when accounting for their heterogeneity in terms of reaction to a treatment. Our study provides a use case of how the machine learning-based assessment of discounts could be implemented also in other businesses and industries facing capacity constraints.

This paper proceeds as follows. Section 2 presents the institutional setting of passenger railway transport in Switzerland. Section 3 describes our data, coming from a unique combination of a customer survey and transport utilization data. Section 4 discusses the identifying assumptions underlying the causal machine learning approach as well as testable implications. Section 5 outlines the employed causal machine learning methods. Section 6 contains a descriptive analysis of the data, including predictive machine learning. Section 7 presents the causal results. Finally, Section 8 concludes.

2. Institutional background

The railway system in Switzerland is known for its high quality of service. Examples include the high level of system integration with frequent services, synchronized timetables, and comprehensive fare integration, see [Desmaris \(2014\)](#). In Switzerland, a country of railway tradition, the state owned incumbent Swiss Federal Railways (SBB) operates the long distance passenger rail market as a monopolist ([Thao et al., 2020](#)). Furthermore, nationally operating long-distance coaches may only be approved if they do not 'substantially' compete with existing services. Thus, the SBB competes exclusively with motorized private transport in Swiss long-distance traffic. The company also owns most of the rail infrastructure, which is funded by the Federal Government. However, since the end of 2020 the companies Berne-Lötschberg-Simplon Railways (BLS) and Southeast Railways (SOB) operate a few links on

behalf of the SBB. Different to regional public transport that Swiss taxpayers subsidize with approximately CHF 1.9 bn per year, the operation of the long distance public transport itself has to be self-sustaining (Wegelin, 2018).

Because of the monopoly position of the SBB in long distance passenger transport, the prices are screened by the Swiss ‘price watchdog’ (or price monitoring agency) to prevent abuse. Based on the price monitoring act, the watchdog keeps a permanent eye on how prices and profits develop. By the end of 2014, the watchdog concluded that the SBB charged too high prices. As a consequence and through a mutual agreement, the SBB and the Swiss price watchdog agreed on a significantly higher supply of supersaver tickets, which were first offered in 2009. Using a supersaver ticket, customers can travel on long distance public transport routes with a discount of up to 70%. Thereafter, additional agreements were regularly reached regarding the number and scope of the supersaver tickets. While only a few thousand supersaver tickets were sold in 2014, sales increased to about 8.8 million in 2019, see Lüscher (2020).

From the SBB’s perspective, these tickets can serve two purposes. First, the tickets might be used as means to balance out the utilization of transport services. For instance, supersaver tickets could reduce the high demand during peak hours which is a key challenge for public transport. Thus, balancing the demand may reduce delays and increase the number of free seats which is valued by the consumers. The average load of SBBs’ seats amounts to 30% in the long distance passenger transport.⁵ For this reason, there is in the literal sense room for improving the allocation. Second, price sensitive customers can be acquired during off-peak hours at rather negligible marginal costs.

Despite the increased interest in the supersaver tickets in recent years, many users of the Switzerland public transport network purchase a so-called ‘general abonnement’ travel ticket (GA). This (annually renewed) subscription provides free and unlimited access to the public transport network in Switzerland. In 2019, about 0.5 million individuals owned a GA in Switzerland, roughly 6% of the Swiss population. The GA’s cost amounts to 3860 and 6300 Swiss francs for the second and first class, respectively. In the same year, about 2.7 million individuals held a relatively cheap half fare travel ticket amounting to 185 Swiss francs. The latter implies a price reduction of 50% for public transport tickets in Switzerland. Overall, discounts provided through supersaver tickets are slightly lower for owners of half fare tickets, as the SBB aims to attract non-regular public transport users. In our causal analysis, we therefore also control for the possession of a half fare ticket.

3. Data

To investigate supersaver tickets’ effect, we use a unique cross-sectional data set provided by the SBB. Our sample consists of randomly surveyed buyers of supersaver tickets that purchased their tickets between January 2018 and December 2019. These survey data are matched with data on distances between any two railway stops as well as utilization-related information relevant for the supply and calculation of discounts. In Section 7, we provide descriptive statistics for these data.

3.1. Survey data

The customer survey is our primary data source. It for instance includes the outcome variable ‘demand shift’, a binary indicator of whether an interviewee rescheduled her or his trip due to buying a supersaver ticket. ‘Yes’ means that the departure time has been advanced or postponed because of the discount. A second variable characterizing customer behavior is an indicator for upselling, i.e. whether someone purchased a first rather than a second class ticket as a reaction to the discount. Another question asks whether an interviewee would have bought the train trip in the same or a higher class even without being offered a discount, which permits judging whether an additional trip has been sold through offering the discount and allows identifying the subgroup of always buyers under the assumptions outlined further below. Our continuously distributed treatment variable is the discount rate of a supersaver ticket relative to the standard fare, which may take positive values of up to 70%.

Furthermore, we observe two kinds of covariates, namely trip- or demand-related factors and personal characteristics of the interviewee. The former are important control variables for our causal identification strategy outlined below and include the difference between the days of purchase and travel, the weekday, month, and year, an indicator for buying a half fare ticket, departure time, peak hour,⁶ number of tickets purchased per person, class (first or second), indicators for leisure trips, commutes, or business trips, the number of companions (by children and adults if any) and a judgment of how complicated the ticket purchase was on a scale from 1 (complicated) to 10 (easy). Furthermore, it consists of indicators for the point of departure, destination, and public holidays. The personal characteristics include age, gender, migrant status, language (German, French, Italian), and indicators for owning a half fare travel ticket or other subscriptions like those of regional tariff associations, specific connections, and Gleis 7 (‘rail 7’). The latter is a travelcard for young adults not older than 25, providing free access to public transport after 7pm.

3.2. Factors driving the supply of supersaver tickets

In addition to the survey, we have access to factors determining the supply of supersaver tickets with various discounts. This is crucial for our causal analysis that hinges on controlling for all characteristics jointly affecting the discount rate and the demand shift outcome. While information on the distances between railway stops in Switzerland is publicly available,⁷ the SBB provides us

⁵ See <https://reporting.sbb.ch/verkehr>.

⁶ Peak hour is defined as a departure time between 6am and 8:59am or between 4pm and 18:59pm, from Monday to Friday. This time window is chosen on the base of the SBB’s train-path prices. For further details, see <https://company.sbb.ch/en/sbb-as-business-partner/services-rus/onestopshop/services-and-prices/the-train-path-price.html> (assessed on March 24 2021).

⁷ See the Open Data Platform of the SBB: <https://data.sbb.ch/explore/dataset/linie-mitbetriebspunkte> (accessed on March 24 2021), which provides the distances between any stops on a railway route.

for the various connections with information on utilization data, the number of offered seats, and contingency schemes, which define the quantity of offered discounts. This allows us to account for travel distance, offered seats, capacity utilization, and quantities of offered supersaver tickets for various discount levels as well as quantities of supersaver tickets already sold (both quantities at the time of purchase). Furthermore, we create binary indicators for the 27 different contingency schemes of the SBB present in our data, which change approximately every month.

The variables listed in the previous paragraph are important, as the SBB calculates the supply of supersaver tickets based on an algorithm considering four type of inputs: Demand forecasts, advance booking deadlines, number of supersaver tickets already sold, and contingency schemes defining the amount and the size of offered discounts based on the three previous inputs. The schemes are set as a function of the SBB's self-imposed goals such as customer satisfaction but also depend on the requirements imposed by the price watchdog. The algorithm calculates a journey's final discount as a weighted average of all discounts between any two adjacent railway stops along a journey. The weights depend on the distances of the respective subsections of the trip. To approximate the (not directly available) demand forecasts of the SBB, we consider the quarterly average of capacity utilization and the number of offered seats for any two stops, which are available by (exact) departure time, workday, class, and, weekend. In addition, we make use of indicators for place of departure, destination, month, year, weekday and public holidays. We use this information to reconstruct the amount and size of offered discounts by taking values from the contingency schemes that correspond to our demand forecast approximation combined with the difference between buying and travel days. Comparing this amount and size of offered discounts with a buyer's discount, we estimate the number of supersaver tickets already sold for the exact date of purchase.

3.3. Sample construction

Our initial sample contains 12,966 long-distance train trips that cover 61,469 sections between two adjacent stops. For 12.2% of these sections, there is no information on the capacity utilization available, which can be due to various reasons. First, for some cases, capacity utilization data is missing. Second, passengers traveling long-distance may switch to regional transport in exceptional cases causing problems for determining utilization. A further reason could be issues in data processing. Altogether, missing information occurs in 3967 trips of our initial sample. We tackle this problem by dropping all journeys with more than 50% of missing information, which is the case for 320 trips or 2.5% of our initial sample. After this step, our evaluation sample consists of 12,646 trips. For the remaining 3647 trips with missing information (which now account for a maximum of 50% of all sections of a journey), we impute capacity utilization as the average of the remaining sections of a trip. In our empirical analysis, we include an indicator for whether some trip information has been imputed as well as the share of imputed values for a specific trip as control variables. Finally, we note that our causal analysis makes (in contrast to the predictive analysis) only use of a subsample, namely observations identified as always buyers who would have traveled even without a discount, all in all 6112 observations.

4. Identification

We subsequently formally discuss the identification strategy and assumptions underlying our causal analysis of the discounts among always buyers.

4.1. Definition of causal effects

Let D denote the continuously distributed treatment 'discount rate' and Y the outcome 'demand shift', a binary indicator for rescheduling a trip due to being offered a discount. More generally, capital letters represent random variables in our framework, while lower case letters represent specific values of these variables. To define the treatment effects of interest, we make use of the potential outcome framework, see for instance (Rubin, 1974). To this end, $Y(d)$ denotes the potential outcome hypothetically realized when the treatment is set to a specific value d in the interval $[0, Q]$, with 0 indicating no discount and Q indicating the maximum possible discount. For instance, $Q = 0.7$ would imply the maximum discount of 70% of a regular ticket fare. The realized outcome corresponds to the potential outcome under the treatment actually received, i.e. $Y = Y(D)$, while the potential outcomes under discounts different from one received remain unknown without further statistical assumptions.

A further complication for causal inference is that our survey data only consists of individuals that purchased a supersaver ticket, a decision that is itself an outcome of the treatment, i.e. the size of the discount. Denoting by S a binary indicator for purchasing a supersaver ticket and by $S(d)$ the potential buying decision under discount rate d , this implies that we only observe outcomes Y for individuals with $S = 1$. In general, making the survey conditional on buying introduces Heckman-type sample selection (or collider) bias, see Heckman (1976) and Heckman (1979), if unobserved characteristics affecting the buying decision S also likely affect the inclination of shifting the timing of the train journey Y . Furthermore, it is worth noting that $S = S(D)$ implies that buying a supersaver ticket is conditional on receiving a non-zero discount. For this reason, non-treated subjects paying regular fares (with $D = 0$) are not observed in our data. Yet, the outcome in our sample is defined relative to the behavior without treatment, as Y indicates whether a passenger has changed the timing of the trip because of a discount. This implies that $Y(0) = 0$ by definition, such that the causal effect of some positive discount d vs. no discount is $Y(d) - Y(0) = Y(d)$ is directly observable among observations that actually received d . However, it also appears interesting to investigate whether the demand shift effect varies across different (non-zero) discount rates $d \in (0, Q]$ to see whether the size matters. This is complicated by the fact that supersaver customers with different discount rates that are observed in our data might in general differ importantly in terms of background characteristics also affecting the outcome, exactly because they bought their trip and were selected into the survey under non-comparable discount

regimes. Our causal approach aims at tackling exactly this issue to establish customer groups that are comparable across discount rates in order to identify the effect of the latter.

Based on the potential notation, we can define different causal parameters of interest. For instance, the average treatment effect (ATE) of providing discount levels d vs. d' (for $d \neq d'$) on outcome Y , denoted by $\Delta(d, d')$, corresponds to

$$\Delta(d, d') = E[Y(d) - Y(d')]. \tag{1}$$

Furthermore, the average partial effect (APE) of marginally increasing the discount level at $D = d$, denoted by $\nabla(d)$, is defined as

$$\nabla(d) = \frac{\partial E[Y(d)]}{\partial d}. \tag{2}$$

Accordingly, $\theta = E[\nabla(D)]$ corresponds to the APE when marginally increasing the actually received discount of any individual (rather than imposing some hypothetical value d for everyone).

The identification of these causal parameters based on observable information requires rather strong assumptions. First, it implies that confounders jointly affecting D and Y can be controlled for by conditioning on observed characteristics. In our context, this appears plausible, as treatment assignment is based on variables related to demand (like weekdays or month), contingency schemes, capacity utilization, and supersaver tickets already sold — all of which is available in our data, as described in Section 3. Second, identification requires that selection S is as good as random (i.e., not associated with outcome Y) given the observed characteristics and the treatment, an assumption known as missing at random (MAR), see for instance (Rubin, 1976) and Little and Rubin (1987). However, the latter condition appears unrealistic in our framework, as our data lack important socio-economic characteristics likely affecting preferences and reservation prices for public transport, namely education, wealth, or income. For this reason, we argue that the ATE and APE among the individuals selected for the survey ($S = 1$), i.e. conditional on buying a supersaver ticket, which are defined as

$$\Delta_{S=1}(d, d') = E[Y(d) - Y(d')|S = 1], \theta_{S=1} = E[\nabla(D)|S = 1], \tag{3}$$

cannot be plausibly identified either. The reason is that if an increase in the discount rate induces some customers to buy a super saver ticket, then buyers with lower and higher discounts will generally differ in terms of their average reservation prices and related characteristics (as education or income), which likely also affect the demand-shift outcome Y .

To tackle this sample selection issue, we exploit the fact that our data provide information on whether the supersaver customers would have purchased a ticket for this specific train trip also in the absence of any discount. Provided that the interviewees give accurate responses, we thus have information on $S(0)$, the hypothetical buying decision without treatment. Under the assumption that each customer’s buying decision is weakly monotonic in the treatment in the sense that anyone purchasing a trip in a specific travel class (e.g., second class) without discount would also buy it for that class in the case of any positive discount, this permits identifying the group of always buyers. Importantly, we therefore define always buyers as those that would buy the trip not in a lower travel class (namely second rather than first class) without discount. For always buyers, $S(0) = S(d) = 1$ for any $d > 0$, such that their buying decision is always one and thus not affected by the treatment, implying the absence of the selection problem. In the denomination of Frangakis and Rubin (2002), the always buyers constitute a so-called principal stratum, i.e., a subpopulation defined in terms of how the selection reacts to different treatment intensities. Therefore, sample selection bias does not occur within such a stratum, in which selection behavior is by definition homogeneous. For this reason, we aim at identifying the ATE and APE on the always buyers:

$$\begin{aligned} \Delta_{S(0)=1}(d, d') &= E[Y(d) - Y(d')|S(0) = 1] = E[Y(d) - Y(d')|S(0) = S(d'') = 1] \text{ for any } d'' \in (0, Q], \\ \theta_{S(0)=1} &= E[\nabla(D)|S(0) = 1] = E[\nabla(D)|S(0) = S(d'') = 1], \end{aligned} \tag{4}$$

where the respective second equality follows from the monotonicity of S in D which is formalized further below.

Fig. 1 provides a graphical illustration of our causal framework based on a directed acyclic graph, with arrows representing causal effects. Observed covariates X that are related to demand are allowed to jointly affect the discount rate D and the demand-shift outcome Y . X may influence the potential purchasing decision under a hypothetical treatment $S(d)$, implying that buying a ticket given a specific discount depends on observed demand drivers like weekday, month, etc. Furthermore, unobserved socio-economic characteristics V (like the reservation price) likely affect both $S(d)$ and Y . This introduces sample selection when conditioning on S , e.g. by only considering survey respondents ($S = 1$). We also note that S is deterministic in D and $S(d)$ (as $S = S(D)$), even when controlling for X . This is the case because conditional on $S = 1$, D is associated with V , which also affects Y , thus entailing confounding of the treatment-outcome relation. A reason for this is for instance that buyers under higher and lower discounts are generally not comparable in terms of their reservation prices. In the terminology of Pearl (2000), S is a collider that opens up a backdoor path between D and Y through V . Theoretically, this could be tackled by jointly conditioning on the potential selection states under treatment values d vs. d' considered in the causal analysis, namely $S(d), S(d')$, as controls for the selection behavior. This is typically not feasible in empirical applications when only the potential selection corresponding to the actual treatment assignment is observed, $S = S(D)$. In our application, however, we do have information on $S(0)$ and can thus condition on being an always buyer under the mentioned monotonicity assumption.




Fig. 1. Causal framework.

4.2. Identifying assumptions

We now formally introduce the identification assumptions underlying our causal analysis.

Assumption 1 (Identifiability of Selection Under Non-treatment). $S(0)$, is known for all subjects with $S = 1$.

Assumption 1 is satisfied in our data in the absence of misreporting, as subjects have been asked whether they would have bought the train trip even in the absence of discount.

Assumption 2 (Conditional Independence of the Treatment). $Y(d), S(d) \perp D | X$ for all $d \in (0, Q]$.

By Assumption 2, there are no unobservables jointly affecting the treatment assignment on the one hand and the potential outcomes or selection states under any positive treatment value on the other hand conditional on covariates X . This assumption is satisfied if the treatment is quasi-random conditional on our demand-related factors X . Note that the assumption also implies that $Y(d) \perp D | X, S(0) = 1$ for all $d \in (0, Q]$.

Assumption 3 (Weak Monotonicity of Selection in the Treatment). $\Pr(S(d) \geq S(d') | X) = 1$ for all $d > d'$ and $d, d' \in (0, Q]$.

By Assumption 3, selection is weakly monotonic in the treatment, implying that a higher treatment state can never decrease selection for any individual. In our context, this means that a higher discount cannot induce a customer to not buy a ticket that would have been purchased under a lower discount. An analogous assumption has been made in the context of nonparametric instrumental variable models, see Imbens and Angrist (1994) and Angrist et al. (1996), where, however, it is the treatment that is assumed to be monotonic in its instrument. Note that monotonicity implies the testable implication that $E[1 - S(0) | X, S = 1, D = d] = E[1 - S(0) | X, S = 1, D = d]$ weakly increases in treatment value d . In words, the share of customers that bought the ticket because of the discount must increase in the discount rate in our survey population when controlling for X .

Assumption 4 (Common Support). $f(d | X, S(0) = 1) > 0$ for all $d \in (0, Q]$.

Assumption 4 is a common support restriction requiring that $f(d | X, S(0) = 1)$, the conditional density of receiving a specific treatment intensity d given X and $S(0) = 1$ (or conditional probability if the treatment takes discrete values), henceforth referred to as treatment propensity score, is larger than zero among always buyers for the treatment doses to be evaluated. This implies that the demand-related covariates X do not deterministically affect the discount rate received such that there exists variation in the rates conditional on X .

Our assumptions permit identifying the conditional ATE given X (CATE), denoted by $\Delta_{X, S(0)=1}(d, d') = E[Y(d) - Y(d') | X, S(0) = 1]$ for $d \neq d'$ and $d, d' \in (1, Q]$. To see this, note that

$$\begin{aligned} \Delta_{X, S(0)=1}(d, d') &= E[Y | D = d, X, S(0) = 1] - E[Y | D = d', X, S(0) = 1], \\ &= E[Y | D = d, X, S(0) = 1, S = 1] - E[Y | D = d', X, S(0) = 1, S = 1], \end{aligned} \tag{5}$$

where the first equality follows from Assumption 2 and the second from Assumption 3, as monotonicity implies that asymptotically, $S = 1$ if $S(0) = 1$. Together with Assumption 1, which postulates the identifiability of $S(0)$, it follows that the causal effect on always buyers is nonparametrically identified, given that common support (Assumption 4) holds. It follows that the ATE among always buyers is identified by averaging over the distribution of X given $S(0) = 1, S = 1$:

$$\Delta_{S(0)=1}(d, d') = E[E[Y | D = d, X, S(0) = 1, S = 1] - E[Y | D = d', X, S(0) = 1, S = 1] | S(0) = 1, S = 1]. \tag{6}$$

Furthermore, considering (5) and letting $d - d' \rightarrow 0$ identifies the conditional average partial effect (CAPE) of marginally increasing the treatment given $X, S(0) = 1$, denoted by $\theta_{X,S(0)=1} = E[\nabla(D)|X, S(0) = 1]$:

$$\theta_{X,S(0)=1} = E \left[\frac{\partial E[Y|D = d, X, S(0) = 1, S = 1]}{\partial d} \Big| X, S(0) = 1, S = 1 \right]. \tag{7}$$

Accordingly, the APE among always buyers that averages over the distributions of X and D is identified by

$$\theta_{S(0)=1} = E \left[\frac{\partial E[Y|D = d, X, S(0) = 1, S = 1]}{\partial d} \Big| S(0) = 1, S = 1 \right]. \tag{8}$$

Our identifying assumptions yield a testable implication if some personal characteristics (like customer’s age) that affect $S(d)$ are observed, which we henceforth denote by W . In fact, D must be statistically independent of W conditional on $X, S(0) = 1, S = 1$ if X is sufficient for avoiding any confounding of the treatment-outcome relation. To see this, note that personal characteristics must by Assumption 2 not influence the treatment decision conditional on X . This statistical independence must also hold within subgroups (or principal strata) in which sample selection behavior (and thus sample selection/collider bias) is controlled for like the always buyers, i.e. conditional on $S(d), S = 1$.

5. Estimation based on causal machine learning

In this section, we outline the causal machine learning approach used in our empirical analysis. Let $i \in \{1, \dots, n\}$ be an index for the different interviewees in our evaluation sample of size n and $\{Y_i, D_i, X_i, W_i, S_i(0)\}$ denote the outcome, treatment, the covariates related to the treatment and the outcome, the observed personal characteristics, and the buying decision without discount of these interviewees that by the sampling design all satisfy $S_i = 1$ (because they are part of the survey). Therefore, Y_i represents customer i ’s demand shift (rescheduling behavior) under customer i ’s received discount rate D_i relative to no discount. Our analysis assesses the causal effect of increasing discount rates on demand shifts among always buyers satisfying $S_i(0) = 1$ while controlling for the selection into the survey and the non-random assignment of the treatment based on Assumptions 1 to 4 of Section 4. We apply the causal forest (CF) approach of Wager and Athey (2018), and Athey et al. (2019) to estimate the CAPE and APE of the continuous treatment, as well as the double machine learning (DML) approach of Chernozhukov et al. (2018) to estimate the ATE of a binary treatment of a discount $\geq 30\%$ vs. $< 30\%$ in the sample of always buyers.⁸

The CF adapts the random forest to the purpose of causal inference. It is based on first running separate random forests for predicting the outcome Y and the treatment D as a function of the covariates X using leave-one-out cross-fitting. The latter implies that the outcome or treatment of each observation is predicted based on all observations in the data but its own, in order to safeguard against overfitting bias. Second, the predictions are used for computing residuals of the outcomes and treatments, in which the influence of X has been partialled out. Finally, a further random forest is applied to average over so-called causal trees, see Athey and Imbens (2016), in order to estimate the CAPE. The causal tree approach contains two key modifications when compared to standard decision trees. First, instead of an outcome variable, it is the coefficient of regressing the residual of Y on the residual of D , i.e. the causal effect estimate of the treatment, that is to be predicted. Recursive splitting aims to find the largest effect heterogeneities across subsets defined in terms of X to estimate the CAPE accurately. Secondly, within each subset, different parts of the data are used for estimating (a) the tree’s splitting structure (i.e., the definition of covariate indicator functions) and (b) the causal effect of the treatment to prevent spuriously large effect heterogeneities due to overfitting.

The CAPE estimate obtained by CF can be thought of as a weighted regression of the outcome residual on the treatment residual. The random forest-determined weight reflects the importance of a sample observation for assessing the causal effect at specific values of the covariates. After estimating the CAPE given X , the APE is obtained by appropriately averaging over the distribution of X among the always buyers. For implementing CAPE and APE estimation, we use the *grf* package by Tibshirani et al. (2020) for the statistical software R. We set the number of trees to be used in a forest to 1000. We select any other tuning parameters like the number of randomly chosen covariates considered for splitting or the minimum number of observations per subset (or node) by the built-in cross-validation procedure.

We also estimate the ATE among always buyers in our sample based on DML for a binary treatment defined as $\tilde{D} = I\{D \geq 0.3\}$, with $I\{\cdot\}$ denoting the indicator function that is equal to one if its argument is satisfied and zero otherwise. Furthermore, let $\mu_d(X) = E[Y|\tilde{D} = d, X, S(0) = 1, S = 1]$ denote the conditional mean outcome and $p_d(X) = \Pr(\tilde{D} = d|X, S(0) = 1, S = 1)$ the propensity score of receiving treatment category d (with $d = 1$ for a discount $\geq 30\%$ and $d = 0$ otherwise) in that population. Estimation is based on the sample analog of the doubly robust identification expression for the ATE, see Robins et al. (1994) and Robins and Rotnitzky (1995):

$$A_{S(0)=1}(1, 0) = E \left[\mu_1(X) - \mu_0(X) + \frac{(Y - \mu_1(X)) \cdot \tilde{D}}{p_1(X)} - \frac{(Y - \mu_0(X)) \cdot (1 - \tilde{D})}{p_0(X)} \Big| S(0) = 1, S = 1 \right]. \tag{9}$$

We estimate (9) using the *causalweight* package for the statistical software R by Bodory and Huber (2018). As machine learners for the conditional mean outcomes $\mu_D(X)$ and the propensity scores $p_D(X)$ we use the random forest with the default options of the *SuperLearner* package of van der Laan et al. (2007), which itself imports the *ranger* package by Wright and Ziegler (2017) for random forests. To impose common support in the data used for ATE estimation, we apply a trimming threshold of 0.01, implying that we drop observations with estimated propensity scores smaller than 0.01 (or 1%) and larger than 0.99 (or 99%) from our sample.

⁸ An introduction to causal machine learning techniques as well as coding examples are provided in the textbook by Huber (2022).

Table 1
Mean and standard deviation by discount and type.

Discount	<30%		≥30%	
	No	Yes	No	Yes
always buyers				
Discount	0.21 (0.07)	0.19 (0.08)	0.57 (0.12)	0.53 (0.13)
Regular ticket fare	44.36 (29.38)	36.14 (25.47)	47.19 (30.14)	32.91 (23.78)
Age	47.22 (15.36)	47.68 (16.14)	45.59 (15.80)	48.77 (16.49)
Gender	0.51 (0.50)	0.55 (0.50)	0.53 (0.50)	0.59 (0.49)
Diff. purchase travel	3.42 (6.96)	3.23 (6.76)	7.72 (11.23)	7.19 (10.30)
Distance	136.49 (77.38)	127.86 (71.49)	126.15 (69.98)	116.76 (66.04)
Capacity utilization	35.51 (14.16)	39.19 (14.31)	26.46 (13.24)	33.15 (13.75)
Seat capacity	328.28 (196.19)	429.57 (196.10)	303.83 (185.42)	445.14 (188.54)
Offer total	33.95 (42.57)	44.10 (50.68)	70.97 (69.57)	98.34 (84.45)
Sold total	28.04 (41.92)	37.29 (50.31)	13.70 (36.37)	25.75 (53.67)
Half fare travel ticket	0.74 (0.44)	0.79 (0.40)	0.62 (0.49)	0.74 (0.44)
Leisure	0.77 (0.42)	0.69 (0.46)	0.82 (0.39)	0.76 (0.43)
Class	1.38 (0.48)	1.65 (0.48)	1.33 (0.47)	1.73 (0.44)
Swiss	0.89 (0.31)	0.92 (0.28)	0.88 (0.33)	0.88 (0.32)
Demand shift	0.31 (0.46)	0.19 (0.40)	0.31 (0.46)	0.23 (0.42)
Upselling	0.49 (0.50)	0.00 (0.00)	0.49 (0.50)	0.00 (0.00)
Obs.	1151	2221	5529	3745

Notes: Regular ticket fare is in Swiss francs. ‘diff. purchase travel’ denotes the difference between purchase and travel day. ‘Offer total’ and ‘sold total’ denote the total amount of supersaver tickets offered and the total amount of supersaver tickets sold respectively.

6. Descriptive analysis

Before discussing the results of our machine learning approaches, we first present some descriptive statistics for our data in [Table 1](#), namely the mean and the standard deviation of selected variables by always buyer status and binary discount category (≥30% and <30%). We see that discounts and regular ticket fares of always buyers are on average lower than those of other customers. Another interesting observation is that in either discount category, we observe fewer leisure travelers among the always buyers than among other customers, which can be rationalized by business travelers responding less to price incentives by discounts. This is also in line with the finding that always buyers tend to purchase more second class tickets. More generally, we see non-negligible variation in demand-related covariates across the four subsamples defined in terms of buying behavior and discount rates. For instance, among always buyers, the total amount of supersaver tickets offered is on average larger in the higher discount category, while it is lower among the remaining clients. This suggests that neither the treatment nor being an always buyer is quasi-random, a problem we aim to tackle based on our identification strategy outlined in [Section 4](#). Concerning the demand-shift outcome, we see that always buyers change the departure time less frequently than others. With regard to upselling, we recognize that the relative amount of individuals upgrading their 2nd class to a 1st class ticket is the same for both discount categories, i.e. ≥30% and <30%.

As a further descriptive analysis, we apply the random forest, a predictive machine learner described in [Appendix A](#), to investigate which variables among the covariates X, W as well as the size of the discount D importantly predict three outcomes of customer behavior. These outcomes indicate whether customers rescheduled their trip e.g. away from rush hours (demand shift), bought a first-class rather than a second-class ticket (upselling), or booked a trip otherwise not realized by train (additional trip), formally defined as $S_i - S_i(0)$. As $S_i = 1$ is equal to one for everyone in our sample, the later outcome corresponds to $1 - S_i(0)$ and indicates whether someone has been induced purchase the ticket because of the discount, i.e. is not an always buyer.

For each of the outcomes, [Table 2](#) presents the 30 most predictive covariates in the training set ordered in decreasing order according to a variable importance measure. The latter is defined as the total decrease in the Gini index (as a measure of node impurity in terms of outcome values) in a tree when including the respective covariate for splitting, averaged over all trees in the forest. The results suggest that trip- and demand-related characteristics like seat capacity, utilization, departure time, and distance are important predictors. Concerning personal characteristics, also customer’s age appears to be relevant. Furthermore, also the

Table 2
Predictive outcome analysis.

Demand shift		Upselling		Additional trip	
Variable	Importance	Variable	Importance	Variable	Importance
Departure time	142.694	Capacity utilization	295.924	Seat capacity	147.037
Seat capacity	121.42	Offer level B	188.861	D	128.086
Age	119.846	Offer level C	149.911	Age	123.948
Capacity utilization	119.606	D	132.095	Departure time	123.516
D	112.474	Age	100.258	Capacity utilization	113.160
Distance	112.143	Departure time	98.909	Distance	101.730
Offer level B	84.142	Offer level A	93.303	Offer level B	84.989
Diff. purchase travel	81.167	Distance	87.319	Diff. purchase travel	80.236
Offer level C	76.238	Offer level D	85.408	Offer level A	78.507
Offer level A	75.971	Diff. purchase travel	62.841	Offer level C	77.097
Number of sub-journeys	73.096	Number of sub-journeys	55.978	Number of sub-journeys	69.443
Offer level D	61.763	Rel. sold level A	44.505	Ticket purchase complexity	64.498
Ticket purchase complexity	57.071	Ticket purchase complexity	41.819	Offer level D	56.888
Rel. sold level A	51.377	Offer level E	37.159	Class	51.456
Rel. amount imputed values	42.222	Rel. sold level B	34.462	Rel. sold level A	46.969
Rel. sold level B	38.144	Rel. amount imputed values	30.747	Rel. amount imputed values	38.869
Adult companions	34.176	Rel. sold level C	28.635	Rel. sold level B	36.484
Rel. sold level C	28.201	Adult companions	25.115	Half fare	35.785
Offer level E	25.714	2019	18.88	Adult companions	34.446
Gender	23.707	Gender	18.47	Halfe fare travel ticket	28.465
Amount purchased tickets	19.575	Rush hour	17.448	Gender	25.419
German	18.659	Saturday	16.173	Rel. sold level C	24.679
Travel alone	18.605	German	15.457	Offer level E	22.556
2019	18.082	Leisure	15.304	Leisure	20.438
French	17.906	Amount purchased tickets	15.112	No subscriptions	19.793
Saturday	17.487	Travel alone	14.792	Amount purchased tickets	19.283
Friday	17.272	Half fare	14.306	German	19.119
Peak hour	17.064	French	14.161	Travel alone	18.139
Class	16.973	Thursday	13.413	2019	17.192
Leisure	16.892	Scheme 20	13.411	French	17.026
Correct prediction rate	0.581		0.817		0.653
Balanced sample size	6962		6738		7000

Notes: 'Offer level A', 'offer level B', 'offer level C', 'offer level D' and 'offer level E' denote the amount of supersaver tickets with discount A, B, C, D and E respectively. 'Diff. purchase travel' denotes the difference between purchase and travel day. 'Rel. sold level A', 'rel. sold level B', 'rel. sold level C' and 'rel. sold level D' denote the relative amount of supersaver tickets offered with discount A, B, C and D respectively. The relative amounts are in relation to the seats offered. 'No subscriptions' indicates not possessing any subscription. For predicting upselling, the covariates 'class' and 'seat capacity' are dropped.

treatment intensity D has considerable predictive power. Interestingly, specific connections (defined by indicators for points of departure and destination) turn out to be less important characteristics conditional on the other covariates already mentioned. The resulting forecasts may for instance serve as a base for customer segmentation, e.g. into customer groups more and less inclined to book an additional trip or switch classes or departure times.

At the bottom of Table 2 we also report the correct classification rates for the three outcomes. While the accuracy in predicting a demand shift amounts to 58%, which is somewhat better than random guessing but not particularly impressive, the performance is more satisfactory for predicting decisions about additional trips with an accuracy of 65% and quite decent for upselling (82%).⁹ Tables C.1 and C.2 in the Appendix presents the predictive outcome analysis separately for subsamples with discounts $\geq 30\%$ and $< 30\%$, respectively. In terms of which classes of variables are most predictive (trip- and demand-related characteristics, age, discount rate) and also in terms of accuracy, the findings are rather similar to those in Table 2. In general, predictive machine learning appears useful for forecasting customer behavior in the context of demand for train trips, albeit not equally well for all outcomes of interest.

However (and in contrast to our causal machine learning approach for assessing the demand effects of discounts), the predictive analysis in this section does not yield the causal effects of the various predictors. In particular, our approach averages the predictions of the respective outcome over different levels of discounts D and thus different customer types in terms of reservation price (related to $S(0)$) and unobserved background characteristics that likely vary with the treatment level. Therefore, we also perform the prediction analysis within subgroups defined upon the treatment level D to see whether the set of important predictors is affected by the treatment intensity. To this end, we binarize the treatment such that it consists of two categories, namely (non-zero) discounts below 30%, i.e. covering the treatment range $d \in (0, 0.3)$, and more substantial discounts of 30% and more, $d \in [0.3, 0.7]$, as 70% is the highest discount observed in our data. The tables with the results of this approach are provided in Appendix C.

⁹ We note that when predicting upselling, we drop the variable 'class', which indicates whether someone travels in the first or second class, and 'seat capacity', which refers to the capacity in the chosen class, from the set of predictors. The reason is that upselling is defined as switching from second to first class, and therefore, the chosen class and the related seat capacities are actually part of the outcome to be predicted.

Table 3
Monotonicity tests.

	CF: average change	LR: coefficient	DML: $D \geq 0.3$ vs $D < 0.3$
Change in $(1 - S(0))$	0.564	0.544	0.184
Standard error	0.060	0.031	0.009
p-value	0.000	0.000	0.000
Trimmed observations			1760
Number of observations	12924		

Notes: ‘CF’, ‘LR’, and ‘DML’ stands for estimates based on causal forests, linear regression, and double machine learning, respectively. ‘trimmed observations’ is the number of trimmed observations in DML when setting the propensity score-based trimming threshold to 0.01. Control variables consist of X .




Fig. 2. Monotonicity given X .

7. Results

7.1. Testing the identification strategy

Before presenting the results for the causal analysis, we consider two different methods to partially test the assumptions underlying our identification strategy. First, we test Assumption 3 (weak monotonicity) by running the CF and DML procedures as well as a conventional linear regression in which we use buying an additional trip $(1 - S(0))$, i.e. not being an always buyer, as outcome variable and X as control variables in our sample of supersaver customers. The CF permits estimating the conditional change in the share of surveyed customers induced to buy an additional trip by modifying the discount rate D given X and $S = 1$, i.e. $E \left[\frac{\partial E[(1-S(0))|D=d, X, S=1]}{\partial d} \middle| X, S = 1 \right]$, as well as the average thereof across X conditional on sample selection, $E \left[\frac{\partial E[(1-S(0))|D=d, X, S=1]}{\partial d} \middle| S = 1 \right]$. DML, on the other hand, yields an estimate of the average difference in the share of additional trips across the high and low treatment categories conditional on sample selection, $E[E[(1 - S(0))|D < 0.3, X, S = 1] - E[(1 - S(0))|D \geq 0.3, X, S = 1]|S = 1]$. Finally, the linear regression of $(1 - S(0))$ on D and all X in our sample tests monotonicity when assuming a linear model.

Table 3 reports the results, which do not provide any evidence against the monotonicity assumption. When considering the continuous treatment D , the CF-based estimate of $E \left[\frac{\partial E[(1-S(0))|D=d, X, S=1]}{\partial d} \middle| S = 1 \right]$ is highly statistically significant and suggests that augmenting the discount by one percentage point increases the share of customers otherwise not buying the ticket by 0.56 percentage points on average. Furthermore, any estimates of the conditional change $E \left[\frac{\partial E[(1-S(0))|D=d, X, S=1]}{\partial d} \middle| X, S = 1 \right]$ are positive, as displayed in the histogram of Fig. 2, and 82.2% of them are statistically significant at the 10% level, 69.1% at the 5% level. Furthermore, the linear regression coefficient of 0.544 is highly significant. Likewise, the statistically significant DML estimate points to an increase in the share of additional trips by 18.4 percentage points when switching the binary treatment indicator from $D < 0.3$ to $D \geq 0.3$.

We also test the statistical independence of D and W conditional on X in our sample of always buyers, as implied by our identifying assumptions, see the discussion at the end of Section 4. To this end, we run a linear regression of D on X and include one variable in W at a time. Using a Wald test, we check whether the added variable improves the model. The results suggest that W and D are largely independent: the average p -value of the Wald test is 0.44, with only one regressor being marginally significant

Table 4
Effects on demand shift.

	CF	DML	LR	PS-matching
	APE	ATE $D \geq 0.3$ vs $D < 0.3$	Coefficient	ATE $D \geq 0.3$ vs $D < 0.3$
Effect	0.161	0.038	0.082	0.077
Standard error	0.063	0.010	0.042	0.082
p-value	0.011	0.000	0.053	0.352
Trimmed observations		151		
Number of observations	5903			

Notes: ‘LR’, ‘PS-Matching’, ‘CF’ and ‘DML’ stands for estimates based on ordinary least squares, propensity score matching, causal forests, and double machine learning, respectively. ‘trimmed observations’ is the number of trimmed observations in DML when setting the propensity score-based trimming threshold to 0.01. Control variables consist of both X and W .

(a dummy for French-speaking with a p -value of 0.049). This result does not provide compelling statistical evidence that W is associated with D conditional on X .

7.2. Assessing the causal effect of discounts

Table 4 presents the main results of our causal analysis, namely the estimates of the discount rate’s effect on the demand shift outcome, which is equal to one if the discount induced rescheduling the departure time and zero otherwise. We note that all covariates, i.e. both the trip- or demand-related factors X and the personal characteristics W , are used as control variables, even though we have previously claimed that X is sufficient for identification. There are, however, good reasons for including W as well in the estimations. First, conditioning on the personal characteristics available in the data may reduce estimation bias if X is – contrarily to our assumptions and to what our tests suggest – not fully sufficient to account for confounding. Second, it can also reduce the variance of the estimator, e.g. if some factors like age are strong predictors of the outcome. Third, having W in the CF allows for a more fine-grained analysis of effect heterogeneity based on computing more ‘individualized’ partial effects that (also) vary across personal characteristics.

Considering the estimates of the CF, we obtain an average partial effect (APE) of 0.161, suggesting that increasing the current discount rate among always buyers by one percentage point increases the share of rescheduled trips by 0.16 percentage points. This effect is statistically significant at the 5% level. As a word of caution, however, we point out that the standard error is non-negligible such that the magnitude of the impact is not very precisely estimated. When applying DML, we obtain an average treatment effect (ATE) of 0.038 that is significant at the 1% level, suggesting that discounts of 30% and more on average increase the number of demand shifts by 3.8 percentage points compared to lower discounts, which is qualitatively in line with the CF. Furthermore, we find a decent overlap or common support in most of our sample in terms of the estimated propensity scores across lower and higher discount categories considered in DML, see the propensity score histograms in Appendix B. This is important as ATE evaluation hinges on the availability of observations with comparable propensity scores across treatment groups. Only 151 out of our 5903 observations are dropped due to too extreme propensity scores below 0.01 or above 0.99 (pointing to a violation of common support).¹⁰ In summary, our results clearly point to a positive average effect of the discount rate on trip rescheduling among always buyers, which is, however, not overwhelmingly large.

For benchmarking purposes, we also consider the results of more ‘traditional’ estimators, namely linear regression (using the continuous discount as for the CF) and nearest neighbor propensity score matching (using the binary discount definition as for the DML). To implement the latter, we estimate the propensity scores of having a larger discount by a probit specification, run nearest neighbor matching with the default values of the *Matching* package by Sekhon (2011), and bootstrap the previous steps 2000 times to obtain a standard error for propensity score matching. Both linear regression and matching indicate that a higher discount leads to more consumers shifting their demand, yet, neither estimate is significant at the 5% level, while the linear regression effect is still significant at the 10% level. For propensity score matching, the lower significance goes together with a lower precision (or higher variance) of the effect estimate relative to DML. This might point to an overfitting issue related to a limited number of observations but a relatively large set of covariates to be included into the propensity score. Therefore, the data-driven selection of important covariates (rather than using any of them) by means of causal machine learning algorithms like DML may appear preferable to reduce the risk of overfitting and increase precision.

7.3. Effect heterogeneity

In this section, we assess the heterogeneity of the effects of D on Y across interviewees and observed characteristics. Fig. 3 shows the distribution the CF-based conditional average effects (CAPE) of marginally increasing the discount rate given the covariates values of the always buyers in our sample (which are also the base for the estimation of the APE). While the CAPEs are predominantly positive, they are quite imprecisely estimated. Only 2.9% and 0.8% of the positive ones are statistically significant at the 10% and

¹⁰ Our findings of a positive ATE remain robust when setting the propensity score-based trimming threshold to 0.02 (ATE: 0.042) or 0.05 (ATE: 0.045).




Fig. 3. CAPEs.

Table 5
Most important covariates for predicting CAPEs.

Covariate	Importance
Seat capacity	11.844
Offer level C	11.164
Capacity utilization	5.144
Departure time	5.122
Distance	4.287
Offer level D	4.015
Class	3.434
Saturday	2.933
Age	2.429
Number of sections	2.373
Diff. purchase travel	2.110
Offer level A	1.634
Offer level B	1.610
Half fare	1.524
Scheme 17	1.496
Half fare travel ticket	1.373
Rel. sold level B	0.901
Ticket purchase complexity	0.847
Leisure	0.773
Rel. sold level A	0.770

Notes: ‘Offer level A’, ‘offer level B’, ‘offer level C’ and ‘offer level D’ denote the amount of supersaver tickets with discount A, B, C and D respectively. ‘Rel. offer level A’, ‘rel. offer level B’ and ‘rel. offer level C’ denote the relative amount of supersaver tickets offered with discount A, B and C. The relative amounts are in relation to the seats offered.

5% levels, respectively. Further, only 0.1% of the negative ones are statistically significant at the 10% level. Yet, the distribution points to a positive marginal effect for most always buyers and also suggests that the magnitude of the effects varies non-negligibly across individuals.

Next, we assess the effect heterogeneity across observed characteristics based on the CF results. First, we run a conventional random forest with the estimated CAPEs as the outcome and the covariates as predictors to assess the covariates’ relative importance for predicting the CAPE, using the decrease in the Gini index as importance measure as also considered in Section 6. Table 5 reports the 20 most predictive covariates ordered in decreasing order according to the importance measure. Demand-related characteristics (like seat capacity, utilization, departure time, and distance) turn out to be the most important predictors for the size of the effects, also customer’s age has some predictive power. Similarly as for outcome prediction in Section 6, specific connections (characterized by points of departure or destination) are less important predictors of the CAPEs given the other information available in the data.

Table 6
Effect heterogeneity analysis.

	Effect	Standard error	p-value
<i>Discounts categories ($D \geq 0.3$ vs $D < 0.3$)</i>			
APE for $D < 0.3$ (constant)	0.209	0.089	0.019
Difference APE $D \geq 0.3$ vs $D < 0.3$ (slope coefficient)	-0.104	0.122	0.395
<i>Customer and travel characteristics</i>			
Constant	-0.154	0.295	0.602
Age	-0.002	0.004	0.556
Gender	-0.022	0.129	0.866
Distance	-0.000	0.001	0.697
Leisure trip	0.297	0.165	0.072
Commute	0.241	0.241	0.316
Half fare travel ticket	0.228	0.142	0.109
Peak hours	0.222	0.133	0.094

Notes: Business trip is the reference category for the indicators 'leisure trip' and 'commute'.

While Table 5 provides information on the best predictors of effect heterogeneity, it does not give insights on whether effects differ importantly and statistically significantly across specific observed characteristics of interest. For instance, one question relevant for designing discount schemes is whether (marginally) increasing the discounts is more effective among always buyers so far exposed to rather small or rather large discounts. Therefore, we investigate whether the CAPEs are different across our binary treatment categories defined by \tilde{D} (30% or more and less than 30%). To this end, we apply the approach of Semenova and Chernozhukov (2021) based on (i) plugging the CF-based predictions into a modified version of the doubly robust functions provided within the expectation operator of (9) that is suitable for a continuous D and (ii) linearly regressing the doubly robust functions on the treatment indicator \tilde{D} . The results are reported in the upper panel of Table 6. While the point estimate of -0.104 suggests that the demand shifting effect of increasing the discount is on average smaller when discounts are already quite substantial (above 30%), the difference is far from being statistically significant at any conventional level.

Using again the method of Semenova and Chernozhukov (2021), we also investigate the heterogeneity among a limited and pre-selected set of covariates that appears interesting for characterizing customers and their travel purpose, namely age, gender, and travel distance, as well as indicators for leisure trip and commute (with business trip being the reference category), traveling during peak hours, and possession of a half fare travel tickets. As displayed in the lower panel of Table 6, we find no important effect heterogeneities across the age or gender of always buyers or as a function of travel distance conditional on the other information included in the regression, as the coefficients on these variables are close to zero. In contrast, the effect of demand shift is (given the other characteristics) substantially larger among always buyers with a half fare travel tickets and among commuters, however, neither coefficient is statistically significant at the 10% level (even though the half fare coefficient is close).

For leisure trips, the coefficient is even larger (0.297), suggesting that all other included variables equal, a one percentage point increase in the discount rate increases the share of rescheduled trips by 0.29 percentage points more among leisure travelers than among always buyers traveling for business. The coefficient is statistically significant at the 10% level, even though we point out that the p -value does not account for multiple hypothesis testing of several covariates. This finding can be rationalized by leisure travelers being likely more flexible in terms of timing than business travelers. Also the coefficient on peak hours is substantially positive (0.222) and statistically significant at the 10% level (again, without controlling for multiple hypothesis testing). This could be due to peak hours being the most attractive travel time, implying that costumers are more willing to reschedule their trips when being offered a discount within peak hours. We conclude that even though several coefficients appear non-negligible, statistical significance in our heterogeneity analysis is overall limited, which could be due to the (for the purpose of investigating effect heterogeneity) limited sample of several thousand observations.

8. Conclusion

In this study, we applied causal machine learning to assess the demand effects of discounts on train tickets issued by the Swiss Federal Railways (SBB), the so-called 'supersaver tickets'. To this end, we analyzed a unique data set that combines a survey of supersaver customers with rail trip- and demand-related information provided by the SBB. Our main result suggested that increasing the discount rate by one percentage point entails an average increase of 0.16 percentage points in the share of rescheduled trips among so-called always buyers (who would have traveled even without discount). We also found limited evidence for effect heterogeneity, e.g., that leisure travelers might be more responsive to discounts than other customer groups.

Our study appears to be the first (at least for Switzerland) to provide empirical evidence on how discounts on train tickets affect customers' willingness to reschedule trips — an important information for designing discount schemes aiming at balancing out train utilization across time and reducing overload during peak hours. Even though the overall impact on the demand shifts on always buyers might not be as large as one could hope for, the causal forest pointed to the existence of customer segments that are likely more responsive and could be scrutinized when collecting a larger amount of data than available for our analysis. Furthermore, our empirical approach may also be applied to other countries or transport industries facing capacity constraints. For instance, we would expect that in a setting with higher competition from alternative public transport modes like long distance bus services (not

present in Switzerland), the impact of train discounts may well be different. More generally, our study can be regarded as a use case for how predictive and in particular causal machine learning can be fruitfully applied for business analytics and as decision support for optimizing specific interventions like discount schemes based on impact evaluation.

Acknowledgments

We are grateful to the SBB Research Fund, Switzerland for financial support. Furthermore, we are indebted to Pierre Chevalier, Reto Lüscher, Philipp Wegelin and Kevin Blättler for their helpful discussions.

Appendix A. Predicting customer behavior based on the random forest

The random forest is a nonparametric machine learner suggested by Breiman (2001) for predicting outcomes as a function of observed variables, so-called predictors. Random forests rely on repeatedly drawing subsamples from the original data and averaging over the predictions in each subsample obtained by a decision tree, see Breiman et al. (1984). The idea of decision trees is to recursively split the predictor space, i.e. the set of possible values of predictors, into a number of non-overlapping subsets (or nodes). Recursive splitting is performed such that after each split, a statistical goodness-of-fit criterion like the sum of squared residuals, i.e. the difference between the outcome and the subset-specific average outcome, is minimized across the newly created subsets. Intuitively, this can be thought of as a regression of the outcome on a data-driven choice of indicator functions for specific (brackets of) predictor values. At each split of a specific tree, only a random subset of predictors is chosen as potential variables for splitting in order to reduce the correlation of tree structures across subsamples, which together with averaging predictions over all subsamples reduces the estimation variance of the random forest when compared to running a single tree in the original data. Even when using an excessive number of splits (or indicator functions for predictor values) such that some of them do not importantly predict the outcome, averaging over many samples will cancel out those non-predictive splits that are only due to sampling noise. Forest-based predictions can be represented by smooth weighting functions that bear some resemblance with kernel regression, with the important difference that random forests detect good predictors in a data-driven way.

To predict the outcomes of customer behavior as discussed in Section 6, we use the *randomforest* package by Breiman (2018) for the statistical software R to implement the random forest based on growing 1000 decision trees. To this end, we create three distinct datasets in which the values of the respective binary outcome are balanced, i.e. 1 (for instance, upselling) for 50% and 0 (no upselling) for 50% of the observations. We balance our data because we aim to train a model that predicts both outcome values equally well. Taking the demand shift outcome as an example, our data with non-missing covariate or outcome information contain 3481 observations with $Y = 1$ and 9576 observations with $Y = 0$. We retain all observations with $Y = 1$ and randomly draw 3481 observations with $Y = 0$ to obtain such a balanced data set. In the next step, we randomly split these 6962 observations into a training set consisting of 75% of the data and a test set (25%). In the training set, we train the random forest using the treatment D and all covariates X, W as predictors. In the test set, we predict the outcomes based on the trained forest, classifying e.g. observations with a demand shift probability ≥ 0.5 as 1. We then compare the predictions to the actually observed outcomes to assess model performance based on the correct classification rate (also known as accuracy), i.e. the share of observations in the test data for which the predicted outcome corresponds to the actual one.

Appendix B. Propensity score plots

See Figs. B.1 and B.2.




Fig. B.1. Propensity score estimates in the higher discount category ($D \geq 0.3$).




Fig. B.2. Propensity score estimates in the lower discount category ($D < 0.3$).

Appendix C. Further tables

See Tables C.1 and C.2.

Table C.1
Predictive outcome analysis, $D < 0.3$.

Demand shift		Upselling		Additional trip	
Variable	Importance	Variable	Importance	Variable	Importance
Departure time	37.33	Capacity utilization	41.387	Seat capacity	25.342
Seat capacity	27.871	Offer level D	27.669	Age	21.639
Capacity utilization	26.508	Age	22.145	Capacity utilization	20.168
Distance	26.31	D	19.077	Distance	18.970
Age	26.223	Offer level C	17.324	Departure time	18.527
D	25.08	Departure time	16.538	D	18.076
Number of sub-journeys	17.403	Distance	15.897	Ticket purchase complexity	16.637
Offer level C	15.643	Offer level B	15.696	Offer level C	12.085
Diff. purchase travel	15.299	Rel. sold level B	10.992	Rel. sold level B	11.709
Ticket purchase complexity	15.116	Number of sub-journeys	10.019	Offer level D	11.641
Rel. sold level B	15.03	Diff. purchase travel	9.573	Number of sub-journeys	11.347
Offer level D	15.012	Rel. sold level C	9.367	Diff. purchase travel	10.328
Rel. sold level C	14.625	Offer level A	7.857	Offer level B	10.185
Offer level B	14.413	Rel. sold level D	7.33	Rel. sold level C	8.993
Rel. sold level A	11.856	Ticket purchase complexity	7.319	Rel. sold level A	8.162
Offer level A	11.329	Offer level E	7.183	Offer level A	7.643
Rel. amount imputed values	10.079	Rel. sold level A	6.769	Class	7.381
Rel. sold level D	9.625	Rel. amount imputed values	5.422	Rel. sold level D	6.964
Adult companions	8.503	Adult companions	4.881	Rel. amount imputed values	6.189
Offer level E	6.511	Rush hour	4.143	Adult companions	5.785
Gender	5.214	Leisure	3.602	Leisure	5.398
Leisure	5.154	Gender	3.599	Offer level E	4.866
Destination Geneva Airport	4.83	2019	3.597	Gender	4.692
Departure Zuerich	4.736	Travel alone	3.047	German	3.801
Class	4.598	Friday	2.92	Halfe fare travel ticket	3.686
Travel alone	4.59	German	2.825	Travel alone	3.540
Peak hour	4.545	French	2.479	French	3.493
Friday	4.524	Departure Zuerich	2.429	Friday	3.419
German	4.522	Destination Zuerich Airport	2.427	Half fare	3.163
Amount purchased tickets	4.349	Scheme 20	2.427	2019	3.136
Correct prediction rate	0.555		0.772		0.605
Balanced sample size	1642		1140		1202

Notes: ‘Diff. purchase travel’ denotes the difference between purchase and travel day. ‘Rel. offer level A’, ‘rel. offer level B’, ‘rel. offer level C’ and ‘rel. offer level D’ denote the relative amount of supersaver tickets offered with discount A, B, C and D respectively. The relative amounts are in relation to the seats offered. ‘Offer level A’, ‘Offer level B’, ‘Offer level C’, ‘Offer level D’ and ‘Offer level E’ denotes the amount of supersaver tickets with discount A, B, C, D and E respectively. ‘No subscription’ indicates not possessing any subscription. For predicting upselling, the covariates ‘class’ and ‘seat capacity’ are dropped.

Table C.2

Predictive outcome analysis, $D \geq 0.3$.

Demand shift		Upselling		Additional trip	
Variable	Importance	Variable	Importance	Variable	Importance
Departure time	114	Seat capacity	246.396	Capacity utilization	133.936
Seat capacity	95.799	Offer level B	178.212	Age	105.107
Age	95.209	Offer level C	127.327	Departure time	100.091
Capacity utilization	89.422	D	100.618	Capacity utilization	97.889
Distance	85.503	Offer level A	88.947	Distance	85.647
D	80.447	Tageszeitinmin	82.658	D	83.671
Diff. purchase travel	69.276	Age	78.886	Offer level B	73.399
Offer level B	68.75	Distance	73.885	Offer level A	69.936
Offer level A	65.766	Offer level D	72.452	Diff. purchase travel	67.823
Offer level C	60.513	Diff. purchase travel	55.321	Offer level C	64.689
Number of sub-journeys	57.767	Number of sub-journeys	48.622	Number of Sub-journeys	58.100
Offer level D	44.626	Rel. sold level A	38.997	Class	54.857
Ticket purchase complexity	44.434	Ticket purchase complexity	31.991	Ticket Purchase complexity	49.348
Rel. sold level A	39.796	Offer level E	27.041	Offer level D	46.685
Rel. amount imputed values	32.144	Rel. amount imputed values	25.586	Rel. sold Level A	42.589
Adult companions	25.925	Adult companions	24.73	Rel. amount imputed values	35.411
Rel. sold level B	20.536	Rel. sold level B	17.099	Half fare	31.308
Gender	18.629	Gender	15.835	Adult companions	28.732
Offer level E	17.521	2019	15.416	Half fare travel ticket	23.900
Travel alone	15.15	Saturday	14.256	Gender	20.562
Amount purchased tickets	14.939	Amount purchased tickets	13.396	Rel. sold level B	20.036
German	14.855	Rush hour	12.947	Offer level E	18.595
French	14.415	German	12.878	German	16.257
2019	14.387	Leisure	12.344	Amount purchased tickets	15.847
Sunday	13.821	Travel alone	12.28	Leisure	15.500
Destination Zuerich Airport	13.387	Half fare	11.882	No subscription	15.434
Saturday	13.378	Friday	11.646	Travel alone	15.132
Class	13.27	Scheme 20	11.559	Swiss	14.951
Half fare	13.258	French	11.477	Saturday	14.613
Rel. amount imputed values	13.048	Sunday	11.39	2019	14.279
Correct prediction rate	0.589		0.809		0.629
Balanced sample size	5320		5598		5798

Notes: 'Diff. purchase travel' denotes the difference between purchase and travel day. 'Rel. offer level A', 'rel. offer level B', 'rel. offer level C' and 'rel. offer level D' denote the relative amount of supersaver tickets offered with discount A, B, C and D respectively. The relative amounts are in relation to the seats offered. 'Offer level A', 'Offer level B', 'Offer level C', 'Offer level D' and 'Offer level E' denotes the amount of supersaver tickets with discount A, B, C, D and E respectively. 'No subscription' indicates not possessing any subscription. For predicting upselling, the covariates 'class' and 'seat capacity' are dropped.

References

- Angrist, J., Imbens, G., Rubin, D., 1996. Identification of causal effects using instrumental variables. *J. Amer. Statist. Assoc.* 91, 444–472.
- Ascarza, E., 2018. Retention futility: Targeting high-risk customers might be ineffective. *J. Mar. Res.* 55 (1), 80–98.
- Athey, S., Imbens, G., 2016. Recursive partitioning for heterogeneous causal effects. *Proc. Natl. Acad. Sci.* 113, 7353–7360.
- Athey, S., Imbens, G.W., 2019. Machine learning methods that economists should know about. *Annu. Rev. Econ.* 11.
- Athey, S., Tibshirani, J., Wager, S., 2019. Generalized random forests. *Ann. Statist.* 47, 1148–1178.
- Basso, L.J., Silva, H.E., 2014. Efficiency and substitutability of transit subsidies and other urban transport policies. *Am. Econ. J. Econ. Policy* 6 (4), 1–33.
- Batty, P., Palacin, R., González-Gil, A., 2015. Challenges and opportunities in developing urban modal shift. *Travel Behav. Soc.* 2 (2), 109–123.
- Blanco, G., Flores, C.A., Flores-Lagunes, A., 2011. Bounds on Quantile Treatment Effects of Job Corps on Participants' Wages. Discussion Paper.
- Bodory, H., Huber, M., 2018. The Causalweight Package for Causal Inference in R. SES Working Paper 493, University of Fribourg.
- Breiman, L., 2001. Random forests. *Mach. Learn.* 45, 5–32.
- Breiman, L., 2018. randomForest: Breiman and cutler's random forests for classification and regression. R package version 4.6-12. Software Available At URL: <https://cran.r-project.org/package=randomforest>.
- Breiman, L., Friedman, J., Olshen, R., Stone, C., 1984. Classification and Regression Trees. Wadsworth, Belmont, California.
- Brynjolfsson, E., McElheran, K., 2016. The rapid adoption of data-driven decision-making. *Amer. Econ. Rev.* 106 (5), 133–139.
- Chernozhukov, V., Chetverikov, D., Demirer, M., Duflo, E., Hansen, C., Newey, W., Robins, J., 2018. Double/debiased machine learning for treatment and structural parameters. *Econom. J.* 21, C1–C68.
- De Palma, A., Lindsey, R., Monchambert, G., 2017. The economics of crowding in rail transit. *J. Urban Econ.* 101, 106–122.
- De Witte, A., Macharis, C., Lannoy, P., Polain, C., Steenberghen, T., Van de Walle, S., 2006. The impact of "free" public transport: The case of Brussels. *Transp. Res. A Policy Pract.* 40 (8), 671–689.
- Desmaris, C., 2014. The reform of passenger rail in Switzerland: More performance without competition. *Res. Transp. Econ.* 48, 290–297.
- Frangakis, C., Rubin, D., 2002. Principal stratification in causal inference. *Biometrics* 58, 21–29.
- Hagenauer, J., Helbich, M., 2017. A comparative study of machine learning classifiers for modeling travel mode choice. *Expert Syst. Appl.* 78, 273–282.
- Heckerman, D., 2013. A Bayesian approach to learning causal networks. arXiv preprint.
- Heckman, J.J., 1976. The common structure of statistical models of truncation, sample selection, and limited dependent variables, and a simple estimator for such models. *Ann. Econ. Soc. Meas.* 5, 475–492.
- Heckman, J.J., 1979. Sample selection bias as a specification error. *Econometrica* 47, 153–161.
- Huber, M., 2014. Treatment evaluation in the presence of sample selection. *Econometric Rev.* 33, 869–905.

- Huber, M., 2022. Causal Analysis - Impact Evaluation and Causal Machine Learning with Applications in R. MIT Press.
- Hünernmund, P., Kaminski, J., Schmitt, C., 2021. Causal machine learning and business decision making.
- Imai, K., 2008. Sharp bounds on the causal effects in randomized experiments with 'truncation-by-death'. *Statist. Probab. Lett.* 78, 144–149.
- Imbens, G.W., Angrist, J., 1994. Identification and estimation of local average treatment effects. *Econometrica* 62, 467–475.
- Kholodov, Y., Jenelius, E., Cats, O., van Oort, N., Mouter, N., Cebecauer, M., Vermeulen, A., 2021. Public transport fare elasticities from smartcard data: Evidence from a natural experiment. *Transp. Policy* 105, 35–43.
- Knaus, M.C., 2021. A double machine learning approach to estimate the effects of musical practice on student's skills. *J. R. Stat. Soc. Ser. A (Stat. Soc.)* 184 (1), 282–300.
- van der Laan, M.J., Polley, E.C., Hubbard, A.E., 2007. Super learner. In: *Statistical Applications in Genetics and Molecular Biology*. Vol. 6. De Gruyter, Berlin, Boston.
- Lee, D.S., 2009. Training, wages, and sample selection: estimating sharp bounds on treatment effects. *Rev. Econom. Stud.* 76, 1071–1102.
- Little, R., Rubin, D., 1987. *Statistical Analysis with Missing Data*. Wiley, New York.
- Liu, L., Chen, R.-C., 2017. A novel passenger flow prediction model using deep learning methods. *Transp. Res. C* 84, 74–91.
- Liu, Y., Liu, Z., Jia, R., 2019. DeepPF: A deep learning based architecture for metro passenger flow prediction. *Transp. Res. C* 101, 18–34.
- Lüscher, R., 2020. 10 Jahre Sparbillette – Attraktive Preise Und Höhere Nachfrage Für Den Öv. Discussion Paper, Blog Litra.
- Mohring, H., 1972. Optimization and scale economies in urban bus transportation. *Am. Econ. Rev.* 62 (4), 591–604.
- Omrani, H., 2015. Predicting travel mode of individuals by machine learning. *Transp. Res. Procedia* 10, 840–849.
- Parry, I.W., Small, K.A., 2009. Should urban transit subsidies be reduced? *Amer. Econ. Rev.* 99 (3), 700–724.
- Paulley, N., Balcombe, R., Mackett, R., Titheridge, H., Preston, J., Wardman, M., Shires, J., White, P., 2006. The demand for public transport: The effects of fares, quality of service, income and car ownership. *Transp. Policy* 13 (4), 295–306.
- Pearl, J., 2000. *Causality: Models, Reasoning, and Inference*. Cambridge University Press, Cambridge.
- Peters, J., Janzing, D., Schölkopf, B., 2017. *Elements of Causal Inference: Foundations and Learning Algorithms*. The MIT Press.
- Redman, L., Friman, M., Gärling, T., Hartig, T., 2013. Quality attributes of public transport that attract car users: A research review. *Transp. Policy* 25, 119–127.
- Robins, J.M., Rotnitzky, A., 1995. Semiparametric efficiency in multivariate regression models with missing data. *J. Amer. Statist. Assoc.* 90, 122–129.
- Robins, J.M., Rotnitzky, A., Zhao, L., 1994. Estimation of regression coefficients when some regressors are not always observed. *J. Amer. Statist. Assoc.* 90, 846–866.
- Rotaris, L., Danielis, R., 2014. The impact of transportation demand management policies on commuting to college facilities: A case study at the University of Trieste, Italy. *Transp. Res. A Policy Pract.* 67, 127–140.
- Rubin, D.B., 1974. Estimating causal effects of treatments in randomized and nonrandomized studies. *J. Educ. Psychol.* 66, 688–701.
- Rubin, D.B., 1976. Inference and missing data. *Biometrika* 63, 581–592.
- Sekhon, J.S., 2011. Multivariate and propensity score matching software with automated balance optimization: The matching package for R. *J. Stat. Softw.* 42, 1–52.
- Semenova, V., Chernozhukov, V., 2021. Debiased machine learning of conditional average treatment effects and other causal functions. *Econometr. J.* 24 (2), 264–289.
- Shin, E.J., 2021. Exploring the causal impact of transit fare exemptions on older adults' travel behavior: Evidence from the Seoul metropolitan area. *Transp. Res. A Policy Pract.* 149, 319–338.
- Sucar, L.E., 2015. *Probabilistic Graphical Models*. Springer.
- Thao, V.T., von Arx, W., Frölicher, J., 2020. Swiss cooperation in the travel and tourism sector: long-term relationships and superior performance. *J. Travel Res.* 59 (6), 1044–1060.
- Tibshirani, J., Athey, S., Friedberg, R., Hadad, V., Hirshberg, D., Miner, L., Sverdrup, E., Wager, S., Wright, M., 2020. grf: Generalized random forests. R Package Version 1.2.0, URL <https://github.com/grf-labs/grf>.
- Wager, S., Athey, S., 2018. Estimation and inference of heterogeneous treatment effects using random forests. *J. Amer. Statist. Assoc.* 113, 1228–1242.
- Wallimann, H., Blättler, K., von Arx, W., 2021. Do price reductions attract customers in urban public transport? A synthetic control approach. arXiv preprint.
- Wegelin, P., 2018. Is the mere threat enough? An empirical analysis about competitive tendering as a threat and cost efficiency in public bus transportation. *Res. Transp. Econ.* 69, 245–253.
- Wright, M.N., Ziegler, A., 2017. ranger: A fast implementation of random forests for high dimensional data in C++ and R. *J. Stat. Softw.* 77, 1–17.
- Yang, J.-C., Chuang, H.-C., Kuan, C.-M., 2020. Double machine learning with gradient boosting and its application to the Big N audit quality effect. *J. Econometrics* 216 (1), 268–283.
- Yap, M., Cats, O., 2020. Predicting disruptions and their passenger delay impacts for public transport stops. *Transportation* 1–29.
- Zhang, J., Lindsey, R., Yang, H., 2018. Public transit service frequency and fares with heterogeneous users under monopoly and alternative regulatory policies. *Transp. Res. B* 117, 190–208.
- Zhang, J., Rubin, D.B., 2003. Estimation of causal effects via principal stratification when some outcomes are truncated by 'death'. *J. Educ. Behav. Stat.* 4, 353–368.
- Zhang, J., Rubin, D.B., Mealli, F., 2008. Evaluating the effects of job training programs on wages through principal stratification. In: Millimet, D., Smith, J., Vytlačil, E. (Eds.), *Advances in Econometrics: Modelling and Evaluating Treatment Effects in Econometrics*. Vol. 21. Elsevier Science Ltd. pp. 117–145.