Transportation Research Part A 174 (2023) 103691

Contents lists available at ScienceDirect

Transportation Research Part A

journal homepage: www.elsevier.com/locate/tra

Price versus Commitment: Managing the demand for off-peak train tickets in a field experiment☆

Christoph Thommen [a,*], Beat Hintermann [b,*]

[a] Zurich University of Applied Sciences, Switzerland
[b] University of Basel, Switzerland

ARTICLE INFO

JEL Codes:
L92
R41
L11
C93

Keywords:
Field experiments
Public transport
Trains
Dynamic pricing
Peak pricing
Switzerland

ABSTRACT

Using data from a field experiment, we provide estimates for the own-price elasticity of train travel in Switzerland. Our estimates are based on exogenous changes to the level of discounts for long-distance trains and thus avoid the usual endogeneity problem between demand-dependent discounts. Besides the price, we also vary the gap between the early booking period and departure during the experiment, which allows us to recover the relative effectiveness of pricing and timing measures. We compute own-price elasticities of around −0.7. Ending the early booking period on midnight of the previous day rather than one hour before departure leads to a decrease in the sale of discount tickets by 18–30%, which is equivalent to a price increase by 26–43%. Last, we find that increasing the discount causes people to purchase their tickets at an earlier time, which allows us to quantify the value of commitment. Our results help design measures for peak-shifting in transport at least societal cost.

1. Introduction

Public transport is crucial to efficiently move workers, students and other users to their destination. In densely populated countries and cities, relying on road transport alone would not be cost-effective in terms of private and societal costs (Hintermann et al., 2021), nor would it be physically feasible as private transport requires significantly more space than public transport per person-km. In fact, one of the chief benefits of public transport is to reduce road congestion (Anderson, 2014). At the same time, public transport is itself limited by capacity constraints, especially during peak hours. One way of reducing the stress on the overall transport system and to alleviate crowding is to move trips from peak to off-peak hours for all modes. In this paper, we investigate to what extent this can be achieved by the level and timing of price discounts in public transport by collaborating with a transport company.

Using the setting of a field experiment, we exogenously vary the off-peak discounts and pre-sale deadlines on the purchase of train tickets in Switzerland in the autumn of 2019. We find that a one-percent decrease in the ticket price during off-peak hours causes a demand increase by 0.65–0.7 percent. In order to increase demand for off-peak tickets while protecting overall revenue, the transport company imposes pre-sale deadlines such that a discount, train-bound ticket has to be purchased at a specified minimum period before departure. Extending the pre-sale deadline from midnight on the previous day to one hour before departure leads to a demand increase

☆ We thank Pierre Chevalier, Reto Lüscher, Meredith Fowlie, Martin Huber, Hannes Wallimann, Ali Darudi, Renate Strobl, Matthias Niggli, Ingmar Schlecht and participants at the SSES Annual Congress 2022 in Fribourg and at the University of Basel SCCER Lunch Seminar 2021 for helpful comments and discussions. Magnus Vieten provided excellent research assistance.

* Corresponding authors.
E-mail addresses: christoph.thommen@zhaw.ch (C. Thommen), b.hintermann@unibas.ch (B. Hintermann).

https://doi.org/10.1016/j.tra.2023.103691
Received 15 July 2022; Received in revised form 9 April 2023; Accepted 12 April 2023
Available online 21 June 2023

by 18–30 percent, but also a reduction in the sale of regular tickets (i.e., a windfall to existing users). As would be expected, the effect of changes in price and pre-sale deadline is stronger for trains that depart immediately before or after peak hours (“shoulder” trains).

While fixed or flat rate fares are still the standard pricing scheme in many urban areas, peak spreading measures are increasingly being adopted to improve the functioning of transit systems. These policies usually target morning and evening hours, which are subject to high demand due to standardised working hours. According to Vickrey’s (1969) bottleneck model for private transport, congestion can be fully eliminated by using tolls that vary by the time of the day. De Palma et al. (2017) show that these conclusions also apply to transit systems, and that pricing and capacity expansions can work as complements rather than substitutes.[1]

Besides softening capacity constraints, peak spreading also serves to provide relief for crowding, which is a source of significant external costs in public transport (Hörcher and Graham, 2018; Tirachini et al., 2013; Wardman and Murphy, 2015).

There are different approaches to shifting people from peak hours to preceding or subsequent hours. In many cases, discounts are provided for shoulder periods (e.g., in Hong Kong and Singapore, see Halvorsen et al., 2016). This is sometimes combined with a surcharge for peak hours (e.g., in London and Washington D.C., see De Palma et al., 2016), which sometimes has proven to be more effective than discounts during off-peak hours (Paulley et al., 2006; Douglas et al., 2011).[2] Beheshtian et al. (2020) propose a market design inspired by electricity markets to price in the congestion costs associated with multi-modal transport. The success of peak-shifting policies is mixed as they are often undermined by the steadily increasing overall demand (Ma et al., 2019). In addition, more attractive conditions during shoulder periods lead to an overcrowding of shoulder periods, which is equivalent to an extension of the peak period (Currie, 2010). De Palma et al. (2017) further argue that a large share of workers lack the flexibility to change their trip times. Daniels and Mulley (2013) show that the potential to shift work trips to earlier morning hours is restricted by the presence of biological body rhythms.

Besides varying prices, some transport providers also rely on early commitment through train-specific tickets that have to be purchased in advance. Such “early bird” discounts are employed with the aim of achieving a greater price differentiation effect for a given discount volume as they allow for exploiting differences in demand elasticities.[3] The discounts offered by the transport company in our setting also rely on such pre-sale deadlines, which means that two people using the same (off-peak) train may have paid very different prices depending on the time of ticket purchase. This is similar to the dynamic pricing strategy that has been employed in airlines for many years.[4] Zhu and Zhao (2020) and Khwanpruk et al. (2021) study optimal pricing schemes based on several pre-sale periods using a theoretical framework. However, there is not much empirical work on the effect of the pre-sale deadline on demand for restricted tickets. Huber et al. (2022) study the propensity to reschedule a trip in the same setting and with the same type of tickets as we do. In doing so, they use survey-based data and construct a quasi-random discount rate dependent on a rich set of train and personal characteristics applying machine learning techniques. Van den Berg et al. (2009) conduct a stated preferences experiment and find own-price elasticities to be substantially higher for restricted (but cheaper) tickets than for unrestricted tickets. Ortega-Hortelano et al. (2016) analyse discounted train-specific tickets on Spanish high-speed trains but do not identify effects stemming from variation in the purchase time.

Our paper contributes to the literature in two ways. First, we identify demand elasticities based on an exogenous variation of the price. Many studies evaluate price elasticities by exploiting policy changes that suffer from endogeneity issues or with stated preferences experiments that are based on hypothetical scenarios (De Grange et al., 2013). This is one of the first studies that applies experimental price variation to a transit setting. Second, because we examine the effect of adjustments to price and pre-sale deadline within the same experiment, we can compare their relative effects as well as their interaction. To our knowledge, Huber et al. (2022) is the only other study that has estimated the joint effect of pricing and timing in the same empirical context.

Our results imply that the design of sales conditions can be tailored multi-dimensionally in order to smooth occupancy peaks and thus increase the share of public transport among overall travel. Operationally speaking, our estimates could be used to maximize profits (or minimize losses) when providing transport services. From a social point of view, dynamic pricing will lead to a better use of the available capacity of both road and rail transport.

Besides peak shifting, the off-peak train discounts also lead to an increase in the net demand for public transport, such that the net effect on peak ridership and crowding is not clear. To the extent that some of this demand increase constitutes a shift away from driving, however, off-peak discounts can contribute to meet the public transport demand in the future, reduce crowding and decarbonize the transport sector in the long run.

In the next section, we provide some background and describe the setting of the field experiment. In sections 3 and 4, we describe the methodology and the data, respectively. Section 5 contains the results, and section 6 concludes.

[1] Yang and Tang (2018) show that a fare-reward scheme that incentivizes a shift in departure time towards non-peak hours is welfare enhancing.

[2] Implementations also include free trips for a certain amount of trips in shoulder periods combined with information about occupancy rates (e.g. Melbourne and Sydney, see Henn et al., 2010). Other policies try to steer commuter currents away from local bottlenecks (e.g., Singapore, see (LTA, 2020)).

[3] The underlying reasoning is that customers are expected to be more price-elastic when planning a trip that takes place in the future compared to a situation in which they have already left their home/office and are about to board a train. Purchasing a ticket when boarding is still the standard in many places, including our setting in Switzerland.

[4] Xie and Shugan (2001) argue that dynamic pricing approaches satisfy different price sensitivities and thereby target different customer groups. Accordingly, late bookers tend to travel for business reasons, while early bookers can be characterised as leisure travellers.

Fig. 1. Screenshot from the transport company's online ticket shop. This screenshot shows a timetable query for the trip from Basel to Zurich and for a specific date and time. The availability of saver tickets is represented by the black triangle at the top left. The price shown in the red field represents the price for second-class and with a half-fare card. There is no information on how much longer a discount ticket will be available.

2. Context and design of the experiment

The following section describes the design and implementation of the field experiment on tickets in the Swiss long-distance train system. The experiment was conducted by a state-owned Swiss transport company and took place during four weeks between August 19 and September 15, 2019.

2.1. The "saver tickets" programme

The Swiss railway network includes 5,200 km of rails (SFSO, 2020a). Trains belonging to different transport providers cover a distance of 550,684 km (SFSO, 2020b) and generate 1.76 million person-trips per day, such that on average, each person residing in Switzerland travels 7.2 km per day on a train (8.8 km per day on overall public transport). Whereas peak trains can be severely crowded, the average occupancy of long-distance trains is only 28.9%.[5] The capacity constraint during peak hours combined with a low occupancy rate overall implies that shifting trips from peak to off-peak would provide societal benefits. To increase the occupancy rate of off-peak trains, one of the largest transport providers introduced discount or so-called "saver" tickets.[6] These tickets are only valid for a specific train connection with a predetermined departure time (which differs from the standard tickets, which are valid for any train on a given date) and they have to be purchased a certain time period before departure. Discount ticket sales account for about a quarter of the tickets sold for individual long-distance trips.[7]

Discount tickets are available only for long-distance trains.[8] Furthermore, in order for a trip to be eligible for a discount ticket, an occupancy forecast of 60% must not be exceeded.[9] Similar ticket schemes are also used in other European countries, such as the high-speed train systems in Germany or Spain (Ortega-Hortelano et al., 2016).

Under regular operating conditions, there are different discount schemes with different levels associated with different occupancy rates. There may also be different availability limits for each discount level, as well as different pre-sale deadlines. The discount

[5] This figure is based on data from the largest transport provider, which covers around 75% of all passengers in the Swiss transit system.
[6] These are called "Sparbillette" in German, "Billets dégriffés" in French and "Biglietti risparmio" in Italian.
[7] 80% of long-distance train passengers hold a flat-rate subscription, such as a "general abonnement" that allows the holder to use the entire railway system at no marginal cost, or a local version that is valid within a particular region. As holders of such subscriptions do not need to buy a ticket, they cannot take advantage of the discount tickets. Discussions are currently underway to introduce flat-rate tickets that are valid for certain times only.
[8] Technically, this applies to trains that cross at least one tariff area border, as the tariff structure within the tariff areas is the responsibility of the respective organisations, which do not offer discount tickets themselves.
[9] This forecast is updated several times a day and is fed by different sources. This restriction prevents train overloading, which can already occur at a forecast much lower than 100%, as the estimate is relatively inaccurate and passengers may not be evenly distributed within the train.

Table 1
Experiment settings.

	Discount	
	70%	30%
One hour before departure	A	C
Previous day (midnight)	B	D

Table 2
Experiment design.

	Calendar week			
	34	35	36	37
Line 1	A	B	C	D
Line 2	B	A	D	C
Line 3	C	D	A	B
Line 4	D	C	B	A

schemes can further differ across the four ticket categories (first-class/half-fare, first-class/full-fare, second-class/half-fare and second-class/full-fare).[10] The schemes are adjusted depending on changes in the occupancy forecast. The discount is computed for each segment of a trip (a segment is the link between two stops); the total discount level for a trip corresponds to the weighted average of the discounts of all included sections. Note that most of these settings are not visible to the customers, who simply see a price (and an indication of a discount) when browsing the online schedule; see Fig. 1 for a trip from Basel to Zurich. They can then buy a discount ticket, a regular ticket (which is not train-bound), or no ticket at all.

Under these conditions, it would be impossible to identify the effect of a change in sales conditions on the number of discount tickets sold, as the latter affects the former (e.g., as more tickets are sold, the discount decreases). To overcome this issue, the transport company agreed to decouple the sales conditions from the actual demand during the experiment. This was done by disabling the availability restrictions and applying the same discount factor and pre-sale deadline for all occupancy rates. Furthermore, the occupancy rate was not updated based on the ticket sales.[11] These adjustments ensured that the price and the pre-sale deadline could be varied exogenously in the experiment.

There are other factors, e.g. weather-related, that can cause the occupancy forecast for a specific train to switch from below to above 60% during the pre-sale period, or vice versa. This would be problematic as in this case, a low number of sales would not necessarily be a sign of low demand, but because discount tickets were not always available during the pre-sale period. The procedure described in Appendix A allows for identifying and discarding these observations. Hence, it is possible to identify the unbiased and causal effect of the discount level and the pre-sale deadline on the number of discount tickets sold.

2.2. *Design of the experiment*

For the experiment, the transport company selected four train lines covering distances between approximately 50 and 90 km with travel times ranging from 35 to 75 minutes. Given that discount tickets had existed for years prior to the experiment and continued to be available for all other lines during the experimental phase, it is impossible to identify the effect of introducing discounts as there is no control group (which is why we refrain from calling our experiment a randomized controlled trial). However, by varying the level of the discount, we can identify the effect of a *change* in the discount level and in the pre-sale deadline on the purchase of tickets.

To create an exogenous variation for the price and the duration of the pre-sale period, the selected lines were randomly assigned to four different settings, as shown in Table 1. The settings were defined by interacting two discount levels (30% and 70%) with two different pre-sale deadlines (one hour before departure vs. midnight of the previous day).

Settings A and D have the most and least attractive conditions, respectively, from the perspective of a customer. Settings B and C are in between, but their relative attractiveness cannot be determined ex-ante as this depends on how customers value a higher level of the discount vs. a shorter pre-sale deadline. The settings were applied to the different lines for one week each. The experimental design is shown in Table 2. This design is known as the *Latin squares* approach in the Experiment Design literature (see Cochran and Cox, 1992). However, in our baseline analysis we use more standard econometric techniques that allow for controlling for unobserved heterogeneity across lines. The carried out Latin square approach works primarily as a robustness test (see Appendix B).

[10] Discount tickets are distributed as follows among the four ticket types (with the remainder being regular tickets for each): second-class/half-fare: 24.3%, second-class/full-fare: 23.6%, first-class/half-fare: 14.2%, first-class/full-fare: 15.6%.

[11] In other words, even if a large number of tickets were sold for a train that had a forecasted occupancy rate of 59%, the train remained eligible for discount tickets. Given that the vast majority of riders have a flat-rate subscription (see above), the sale of discount tickets will never lead to an actual crowding problem (of, say, more than 65% occupancy) on a train that was originally forecasted to be less than 60% of occupancy.

For Line 2, the discount level temporarily deviated from the experimental settings, and different segments were priced differently by mistake.[12] As a result, we removed Line 2 from the baseline analysis. However, we included it (by necessity) in the Latin square analysis and and also in a robustness test.

During the experiment phase, full-fare customers partly erroneously received a 50% discount when a 30% discount was supposed to be set. As we use the discount variable as a continuous variable, this irregularity can be taken into account in the analyses and even introduces further variance to the data. A similar situation arises for discount tickets for trains in the experiment period that had been sold under pre-experiment conditions. These observations account for 14.2% of discount ticket sales in the first week of the experiment. Because there were at least three days between the change of settings and the first train in the experiment, the availability restriction was never binding.

3. Methodology

This section presents the empirical strategy to identify the effect of the price and the pre-sale deadline on the volume and timing of discount ticket purchases. We also discuss possible challenges to the identification. Our baseline econometric approach is based on a fixed effects regression. We analyzed the data also according to the Latin squares approach, but we consider this a robustness test (see Appendix B).

3.1. Effect of price and pre-sale deadline on ticket sales

We analyze the effect of changes in the price and the pre-sale deadline on the number of discount tickets sold, first on the day-line-category and then on the train-category level. For the regressions in which the unit of analysis is the number of tickets sold per day, we estimate the following equation:

$$\log(\#\ tickets_{dlc}) = \beta_1 \cdot \log(price_{dlc}) + \beta_2 \cdot pre - sale\ deadline\ 1h_{dlc} + \Gamma_1 \cdot departure\ date_d + \Gamma_2 \cdot (weekday_d \times directional\ line_l \times category_c) + u_{dlc} \quad (1)$$

In words, the dependent variable (log number of discount tickets) is regressed on the log of the price (after the discount), a dummy that is equal to one if the pre-sale deadline dummy is one hour (and zero if it is the previous day), a series of departure date dummies and a number of fixed effects to allow for different demand by the day of the week, directional lines and ticket categories. The indices are defined as follows: d denotes the day of the experiment, running from 1 to 20; l refers to the directional line; and c measures the four main ticket categories defined by the interaction of full vs. half-fare and first vs. second-class.[13]

All observations are weighted according to the average number of tickets sold in each weekday/directional line/category-group. The identifying variation is thus derived from within-changes in ticket prices and the pre-sale deadline, which are varied exogenously according to the experiment settings described in Tables 1 and 2.

The coefficient β_1 identifies the own-price elasticity of discount ticket sales, provided that there is no unobserved heterogeneity in the sense that some other determinant of discount ticket sales co-varies with the price (see discussion below). Because the pre-sale deadline is included as a dummy, β_2 measures a semi-elasticity of changing from a day-ahead to an hour-ahead sale deadline.The parameter vectors Γ_1 and Γ_2 contain the coefficients on the dummy variables (fixed effects), which are not shown in the results.

We estimate eq. (1) using a Poisson pseudo-maximum likelihood (PPML) model. This approach corrects for the over-dispersion in the available data and addresses the problem of a potential bias that can arise when estimating a log-linearized model in the presence of heteroskedasticity (for a discussion, see Santos Silva and Tenreyro, 2006). Because the PPML model is estimated in exponentiated form, it allows for zeros (which would drop out if the equation were estimated in the log-linearized form in eq. (1)).

In additional regressions, we replace the dependent variable in (1) with the number of regular tickets (instead of discount tickets) sold. The resulting coefficient β_1 then captures the cross-price elasticity. Because the departure date is not recorded for regular tickets (see above), we approximate this with the date of purchase. In principle, it is possible to purchase a ticket many days or even weeks in advance. For this reason, we do not place the same confidence in the regressions involving the regular tickets as in the regressions involving the discount tickets.[14]

Because discount tickets are assigned to individual trains, we also carry out an analysis in which the unit of observation is the train-category level. Whereas we expect that the results of the day-level analyses are more robust as train-specific shocks are averaged out across the day, the train-specific analysis allows for examining the effect of the experimental settings on particular times of the day.

For the regression analysis on the train level, we use all trains identified as not having exceeded the 60% occupancy forecast threshold during the course of the purchase period (see Appendix A for details).

For these trains, the ticket types “second-class/half-fare”, “second-class/full-fare” and “first-class/half-fare” are considered. The

[12] More specifically, the discounts of the individual sub-sections were priced according to the overall policy of the transport company, which varied over time and also with the demand. The prices in the individual segments of line 2 were therefore sometimes above and sometimes below the experimental setting that should have been applied to this line, and the same applies to the pre-sale deadline.

[13] There are additional categories such as tickets for dogs, bicycles or traveling groups, but these are excluded from the analysis.

[14] Our sample period does not contain public holidays, for which people may purchase a ticket in advance. Note also that seat reservations (which have to be purchased independently of the tickets themselves) are very rare in train travel within Switzerland and essentially only used for group and international travel.

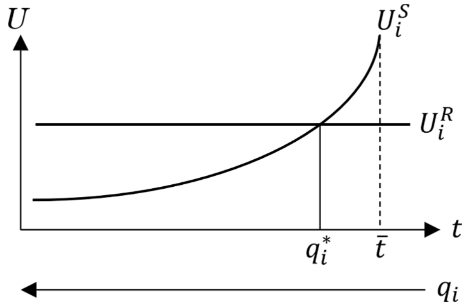

Fig. 2. Utility of discount and regular tickets.

fourth category "first-class/full-fare" is discarded due to too few observations. We use the following regression equation:

$$\log(\#\ tickets_{dtc}) = \beta_1 \cdot \log(price_{dtc}) + \beta_2 \cdot pre-sale\ deadline\ 1h\ _{dt} + \Gamma_1 \cdot dep.\ date_d + \Gamma_2 \cdot dep.\ time_t + \Gamma_3 \cdot (weekday_d \times directional\ line_t \times category_c) + \Gamma_4 \cdot forecast_{dtc} + \Gamma_5 \cdot segments_t + \Gamma_6 \cdot capacity_{dt} + u_{dtc} \tag{2}$$

The identification follows the same principles as on the day-line-category level, but the regression controls for additional, train-specific variables (denoted by the subscript t). They include information on the departure time, the occupancy forecast, the number of stops/segments and the capacity of each train.

Again, we estimate equation (2) in exponentiated form using a PPML model. The feature of this model to preserve zeroes is particularly important for the regressions on the train level, as many trains feature no discount ticket purchases at all. As a robustness test, we also employ a negative binomial distribution. The results are qualitatively similar to those from the PPML model and are included in Table C2 in the Appendix.

3.2. The cost of commitment

Buying a discount ticket constitutes a trade-off between committing to a particular train vs. a lower price. By studying the effect of the discount on the purchase time, we can study the cost of commitment.

Let $U_i(L)$ denote the utility of an individual i for a train trip on Line L form an origin to a destination. For simplicity, we assume that utility is linear in income. The regular ticket price is given by $p(L)$. The individual can buy a regular ticket, a discount ticket, or no ticket at all. We normalize the utility of not traveling by zero.

By buying a discount ticket, the individual commits to a particular departure time (which is not the case for a regular ticket, which can be used for any train on that day). We assume that the commitment cost increases with the commitment period, but at a decreasing rate. The discount ticket needs to be purchased at $t_i \leqslant \bar{t}$, where $\bar{t}$ defines the pre-sale deadline. The utility for a regular ticket (U_i^R) and a discount ticket (U_i^S) is thus given by

$$U_i^R = U_i(L) - p(L) \tag{3}$$

$$U_i^s = \begin{cases} U_i(L) - q_i^{\alpha_i} - (1-d)p(L) & if\ t_i \leqslant \bar{t} \\ 0 & if\ t_i > \bar{t} \end{cases} \tag{4}$$

It is assumed that $U_i(L)$ is unrelated to the ticket in use and remains constant for different purchase times. Hence, U_i^R is independent of time t. The utility of a discount ticket is reduced by commitment costs $q_i^{\alpha_i}$. As it is no longer possible to purchase a discount ticket after the pre-sale deadline, the corresponding utility is 0 for $t_i > \bar{t}$ (if the individual buys a regular ticket, he or she obtains U^R). The price corresponds to the price of a regular ticket p subtracted by the discount of dp, which results in a final price of $(1 - d)p(L)$. The specification $q_i^{\alpha_i}$ ensures the assumed concavity in the commitment costs (with $0 < \alpha_i < 1$). We allow the shape of the commitment cost curve (and thus the temporal flexibility) to vary across individuals.

We now assume that potential customers continuously compare the utility of buying a discount ticket with the constant utility of a regular ticket. Fig. 2 shows this as stylized form. As time progresses, U_i^S increases until dropping to zero at $\bar{t}$. A utility-maximizing customer will purchase the discount ticket at some point between q_i^* and $\bar{t}$. Buying earlier than q_i^* will lead to a too large commitment cost, whereas buying after $\bar{t}$ is not possible.

The discount ticket offer is designed in such a way that customers cannot anticipate from when and until when a specified discount offer is provided. Due to the reduced utility associated with the uncertainty of the future development of U^S (von Neumann and

Morgenstern, 1953), which includes the possibility of the discount to become unavailable,[15] we assume that the purchase of a discount ticket takes place as soon as U^S exceeds U^R. With these assumptions, and given equations (3) and (4), we can compute α_i based on the observed purchase time q_i and the level of discount:

$$\alpha_i = \frac{\log(dp(L))}{\log(q_i)} \tag{5}$$

This equation implies that, given a discount d, people with a lower α_i purchase the ticket earlier than those who perceive higher costs of commitment. The equation furthermore implies that all else equal, the purchase time is moved forward if the discount is increased.

To test this using our data, we compute the difference between the time of purchase and the departure time (in hours). We then use this time difference as the dependent variable in the following model:

$$\begin{aligned} \log(time\ diff._{dtc}) \quad &= \beta_1 \cdot discount_{dtc} + \beta_2 \cdot pre - sale\ deadline\ 1h_{dt} + \Gamma_1 \cdot dep.\ date_d + \Gamma_2 \cdot dep.\ time_t + \Gamma_3 \cdot (weekday_d \\ &\times directional\ line_t \times category_c) + \Gamma_4 \cdot forecast_{dtc} + \Gamma_5 \cdot segments_t + \beta_8 \cdot capacity_{dt} + u_{dtc} \end{aligned} \tag{6}$$

This equation is very similar to (2), the main difference being that instead of the price, the explanatory variable of interest is the discount (in percent).The coefficient β_1 thus measures the (proportional) effect of increasing the discount by 1 percentage point on the time difference between ticket purchase and departure time.

The pre-sale deadline is included as a control, but it is not a variable of interest in itself as increasing the pre-sale deadline will mechanically increase the difference between purchase and departure time. Despite this relationship, we stress that the pre-sale deadline and the (relative) time of purchase measure different things: whereas the former is exogenously determined by the transport company (and held fixed within an experimental treatment setting), the latter is a choice made by the customers. The pre-sale deadline provides a lower bound for the time of purchase, but most discount tickets are purchased well before the deadline. During our experiment, the average time difference between purchase and departure was about over 4 days, even though the longest possible pre-sale deadline is 24 hours (for a train departing shortly before midnight). The time of purchase thus depends also on other aspects, such as the level of discount, and the identification of this effect is the purpose of model (6). Note also that since the discount is fixed during the experiment, there is no reverse causality such that the coefficient β_1 can be measured without bias.

Although the discount was held constant within a treatment period, the regular discount tickets programme includes availability limits per discount level in terms of both quantity and timing. Once these have been reached, the discount is lowered or removed altogether. Regular customers of discount tickets will therefore try to lock in a certain discount level. This trade-off between obtaining a lower price (buying early) and a shorter commitment period (buying later) is a measure for consumers' commitment costs.

3.3. Identification

Ideally, we would have randomized the level of pricing and pre-sale deadline on the person level. This was not feasible for technical and also marketing-related reasons. Instead, we were allowed to exogenously set the level of the discount and the pre-sale deadline for individual lines. The settings were changed once per week (see Table 2).

Note that unobserved shocks that affect all lines are equally absorbed by the calendar day fixed effects, and time-invariant differences across lines are captured by the line fixed effects. Our identification thus relies on the assumption that there were no *line-specific* shocks that coincide with the change in the treatment. This assumption is essentially not testable, and we would also need to rely on such an assumption even if our experiment included a proper control group. For example, if an exogenous shock affecting tickets sales (say, the start of school vacation) coincided with the change in the experimental setting and were to affect line 1 differently to line 3, then this would bias our results. We chose the time period of the experiment with this in mind: Between mid August and mid September, there are no major public holidays or school vacations in Switzerland, such that these threats to identification can be ruled out.

Last, we note that while the experiment settings were applied to predefined origin-destination combinations, the discount tickets programme continued to run as usual on the remaining network. This means that a trip that runs through one of the pre-specified lines but has a different origin or destination does not receive the experiment settings and is therefore not part of the experiment. Since we do not consider the data of other lines for the analysis, we can ignore effects from the experimental changes on the remaining network. However, an effect of a shock in the remaining network on Lines 1–4 would be problematic, especially if this shock affects the four lines differently (otherwise it would be absorbed by the day fixed effects). To minimize this problem, the transport company agreed to keep the settings on the remainder of the network constant for two weeks before and during the experiment.

4. Data

The following information is collected and processed for each discount ticket: sale date, travel date, departure time, train number,

[15] In our experiment, the discount was always available until the pre-sale deadline. However, this pre-sale deadline is not communicated to customers, and in the usual programme, the level of the discount also varies over time. This means that by purchasing a discount ticket, consumers "lock in" a certain discount level, which is the reason for why they don't wait until the last moment.

Fig. 3. Number of discount tickets sold for Lines 1, 3 and 4 between week 10 and 23 (spring) and 30 and 37 (autumn) 2019.

origin and destination, half/full-fare, class, effective price and regular price. For regular tickets, no information on the travel date is available.[16] But according to the train company, around 90% of tickets are bought on the day of travel. Discount tickets data are available for the experiment period of calendar weeks 30 to 37 (autumn) of 2019 and, in addition, for calendar weeks 10 to 23 (spring) of 2019 (see Fig. 3 for the distribution of discounts for both periods). The analysis is based on the data from autumn.

A line is defined by the route between a unique origin and a unique destination. Because the incentives to purchase a discount ticket may differ between the first leg and the return trip, and the cities differ in size, we treat traveling from A to B as a different *directional* line than traveling from B to A. In contrast, a train is a composition of wagons that runs as a unit on a line from the origin to the destination at a departure time.

The spring data is used to provide additional information for the train-level regressions. The train-specific variables include a train's capacity,[17] the number of stops as well as the occupancy forecast.[18] Tables 3 and 4 provide descriptive statistics for the variables used in the regression analyses at the day and the train level.

From the ticket sales data we can reconstruct which trains have run and how many tickets have been sold for a particular train. As can be seen in Table 5, there are 7,017 trains in the data. At the same time, we do not have information on trains with no discount tickets sold. There are two reasons for the absence of sold discount tickets on a train: 1. there was no demand, 2. there was no supply because the occupancy forecast of 60% was exceeded (see section 2.1). The first case must be considered when estimating the effects. However, trains that are affected in the second case have to be excluded from the analysis since no demand can be captured because there are no discount tickets for sale on these trains.

In a first step of the data processing, we identify all trains running on the involved lines during the period of the experiment. To obtain the complete set of trains, we developed a procedure that is described in detail in Appendix A.1. In addition to the 7,017 trains for which we have ticket sales in our data, there are 3,328 trains without sold discount tickets that are considered in order to have a complete set of train observations.

In a second step, based on all trains that ran, it must be determined whether trains that sold no discount tickets were subject to the 60% occupancy forecast restriction in the course of the purchase period (*crowded trains*). Among these resulting 10,345 trains, 1,718 trains are defined as *crowded trains* (i.e., trains that are not eligible for discount tickets because their occupancy forecast is too high) and are excluded from further analyses. Trains that are significantly slower due to an alternative route choice and trains that act as

[16] When a ticket is purchased, both the sale date and the travel date are recorded. For data protection reasons, however, these are stored separately and can not be used together. Note that this problem does not affect discount tickets, for which we know the date and time of purchase as well as the specific train for which they are valid.

[17] In about 10% of the cases, information about the capacity of autumn trains cannot be matched with corresponding spring trains. These gaps are filled with group median data based on departure date, train number and debit number from the spring data. The debit number defines different start and end stations or different rolling stock. For the lines in the experiment, there are between three to six different debit numbers. When there is no information available on that level, group median data based on train number and debit number is imputed. Any remaining record with missing capacity data is discarded.

[18] Forecasts are not available ex-post, and we do not have these data. To fix this problem, the train-specific forecast information was collected on March 4, 2020 for all trains until May 3, 2020 and was merged with the data from spring 2019. Finally, we connected the spring 2019 data with the experiment data on the level of directional line, departure time and weekday. We are aware that the data collection period for this variable coincides with the beginning of the Covid pandemic in Switzerland, which may have affected the predicted occupancy. However, since this variable only serves as a control and thus its coefficient is not of interest, we do not believe that this affects our results. To be certain, we estimated our regressions also without this variable and obtained negligible changes in the results.

Table 3
Summary statistics on the day-line-category level.

Directional line	Class	Half-fare	Obs	# Discount tickets			
				Mean	SD	Min	Max
1a	1	0	20	3.5	2.0	0	7
1a	1	1	20	15.4	6.8	2	28
1a	2	0	20	110.0	32.0	63	175
1a	2	1	20	210.4	67.5	105	342
1b	1	0	20	2.6	1.6	0	6
1b	1	1	20	13.2	6.1	2	23
1b	2	0	20	92.7	34.0	53	164
1b	2	1	20	189.5	78.6	88	364
3a	1	0	20	4.4	4.2	0	15
3a	1	1	20	19.4	10.3	5	38
3a	2	0	20	77.6	27.4	46	155
3a	2	1	20	244.1	75.0	141	412
3b	1	0	20	4.5	4.2	0	18
3b	1	1	20	16.2	9.4	3	34
3b	2	0	20	70.2	16.2	50	116
3b	2	1	20	219.4	70.7	122	373
4a	1	0	20	1.8	1.8	0	6
4a	1	1	20	21.1	10.1	6	46
4a	2	0	20	19.1	6.7	7	30
4a	2	1	20	109.7	42.0	56	209
4b	1	0	20	1.4	1.7	0	5
4b	1	1	20	10.7	5.3	4	24
4b	2	0	20	14.6	4.7	6	22
4b	2	1	20	66.8	28.2	27	127

Note: The data includes weekdays during the period of August 19 to September 15, 2019.

Table 4
Summary statistics on the train level.

Variable	Category	Obs.	Mean	SD	Min	Max
Discount ticket sales						
Pre-sale deadline	1 day	5,397	2.12	3.50	0	40
	1 hour	6,093	2.62	4.41	0	94
Class	1	3,813	0.42	0.88	0	8
	2	7,677	3.41	4.56	0	94
Fare type	Full	3,822	1.85	2.47	0	37
	Half	7,668	2.70	4.50	0	94
Directional line	1a	2,811	2.28	3.38	0	29
	1b	2,376	2.30	3.65	0	42
	3a	1,725	3.58	5.86	0	94
	3b	1,896	2.90	4.53	0	46
	4a	1,392	1.94	3.19	0	40
	4b	1,290	1.19	1.92	0	20
Time	Morning	2,676	3.12	5.10	0	94
	Evening	1,992	3.13	5.20	0	46
	Shoulder	1,509	5.04	7.19	0	94
Other variables						
Weighted mean purchase-time	All	11,490	4.15	6.12	0	60
Weighted mean discount	All	11,490	0.55	0.17	0.19	0.70
Occupancy forecast	All	11,490	0.36	0.12	0.01	0.86
Capacity	All	11,490	554.3	127.5	237	871
Segments	All	11,490	3.11	2.40	1	8

Note: This table shows the data from individual trains that ran along the lines 1, 3 and 4 during the period of the experiment (August 19 to September 15, 2019).

distributors between the origin and destination of a line are summarised as *detour trains*. These have a much lower demand due to different characteristics and cannot be compared with the remainder. Therefore, they are discarded. The same applies to trains that run on weekends (*weekend trains*). Eventually, 3,830 trains remain for the analysis. As a side benefit of this process, the trains that follow or precede non-eligible discount tickets trains (*crowded trains*) can be identified, and we label them as *shoulder trains*. It is reasonable to assume that the demand for discount tickets on these trains is higher than for trains that operate well before or after the peak periods. Details to the identification of train categories can be found in Appendix A.2.

Table 5
Data processing stages.

Total number of trains	**10,345**
Trains with discount tickets sold	7,017
Trains without discount tickets	3,328
Discarded trains	**6,515**
Weekend trains	3,334
Detour trains	1,463
Crowded trains	1,718
Trains used for analysis	**3,830**
Normal trains	3,327
Shoulder trains	503

Note: This table shows the data processing steps, starting from the total number of trains during the sample period.

Table 6
Effect of price and pre-sale deadline on discount ticket purchase.

	Dependent variable: log(# tickets)				
	Day-level	Train-level			
	Full sample	Full sample	Morning	Evening	Shoulder
log(price)	−0.703**	−0.667**	−0.571**	−0.951**	−0.815**
	(0.027)	(0.038)	(0.058)	(0.068)	(0.068)
Pre-sale deadline: 1 day	−0.299**	−0.183**	−0.001	−0.437**	−0.178**
	(0.024)	(0.035)	(0.046)	(0.043)	(0.053)
Observations	480	11,490	2,676	1,992	1,509
R^2	0.975	0.727	0.786	0.851	0.761

Note: Standard errors (in parentheses) corrected for heteroskedasticity and clustered at the line level. **: $p < 0.01$, *: $p < 0.05$, °: $p < 0.1$. The data covers the period of August 19 to September 15, 2019. The regressions additionally include a large nmber of fixed effects (see eqs. 1-2), which are excluded from the table.

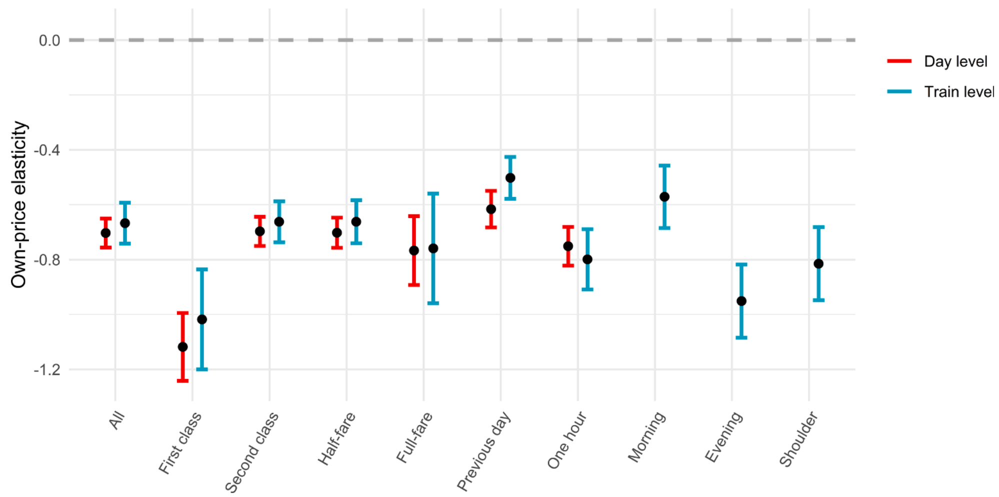

Fig. 4. Own-price elasticity for the full sample and different subsets.

5. Results

We first discuss the results of the price and pre-sale deadline on the number of sold tickets, followed by the effect on the time of purchase. This section focuses on the results of our main econometric approach. The results based on the Latin square methodology, which we consider a robustness test, can be found in Appendix B.

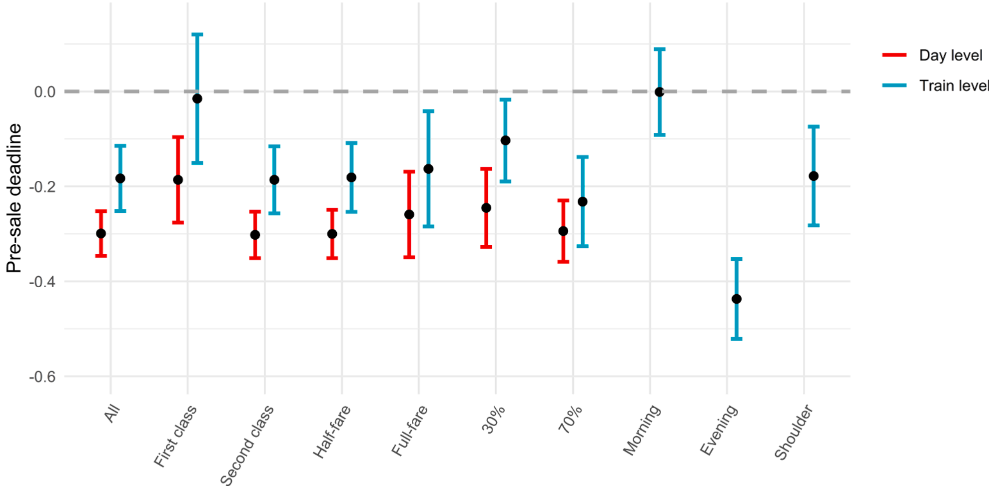

Fig. 5. Effect of pre-sale deadline for the full sample and different subsets.

Table 7
Effect of price and pre-sale deadline on regular tickets (day level).

	Dependent variable: log(# tickets)				
	Full sample	Only first-class	Only second-class	Only half-fare	Only full-fare
log(price)	0.120**	0.087**	0.123**	0.121**	0.141**
	(0.016)	(0.033)	(0.017)	(0.017)	(0.042)
Pre-sale deadline: 1 day	0.083**	0.077**	0.084**	0.087**	0.030
	(0.019)	(0.024)	(0.020)	(0.020)	(0.024)
Observations	480	240	240	240	240
R^2	0.987	0.966	0.982	0.985	0.990

Note: Standard errors (in parentheses) corrected for heteroskedasticity and clustered at the line level. **: $p < 0.01$, *: $p < 0.05$, °: $p < 0.1$.

5.1. Effect on ticket sales

The first column in Table 6 shows the results from estimating eq. (1) on the day level. The own-price elasticity is −0.70, implying that if the price for discounted tickets increases by 1%, demand will contract by 0.7%. The point estimates based on eq. (2) on the train level (column 2) is 3–4 percentage points lower, but the confidence intervals overlap such that the results from the different levels of aggregation are largely consistent. Columns 3–5 show the elasticities for different times during the course of a day. The own-price-elasticity is lower for customers on morning trains (departure until 10 am) than on evening trains (departure between 3 and 8 pm). This is in line with the conclusions made by Van den Berg et al. (2009), and also with studies that find a stronger propensity to postpone than to bring forward (see e.g. Douglas et al., 2011). Furthermore, the possibilities to shift departure to an earlier time in the morning are limited for biological and physical reasons (see Daniels and Mulley, 2013). Similarly, we find that the elasticity is higher for shoulder trains (for a definition, see Appendix A.2). This finding suggests that shoulder trains are a better substitute for the (non-discount) peak trains than trains running significantly before or after peak hours.

Fig. 4 illustrates the own-price elasticity results for the full sample and different subsamples, both for the day and train level. The regression results from the subsample analyses are shown in Table C1 and Table C2 in the Appendix. We estimate a higher price elasticity for first-class tickets for both types of regressions (day and train-level). Our explanation for this result is that first-class tickets are often used for business travel and by people with higher incomes, which suggests a lower elasticity. This is a known feature also in other markets, e.g., in the context of airline pricing (Williams, 2022). Furthermore, discount first-class tickets also attract customers that usually travel second-class. The higher elasticity for first-class tickets suggests that the latter effect dominates. We see no significant difference between full-fare vs. half-fare customers.

Last, we find that the own-price elasticity is greater for the hour-ahead pre-sale deadline relative to the deadline that ends on midnight of the previous day. Since the earlier pre-sale deadline is associated with higher relative commitment costs, the corresponding utility for the purchase of a discount ticket (compared to a regular ticket) is smaller than with a pre-sale deadline of one hour. As a consequence, the demand for discount tickets is less price-sensitive. The results are visualized in Fig. 4 and presented in detail in Table C3, Table C4 and Table C5 in the Appendix.

Fig. 5 and the second coefficient in Table 6 show the effect of varying the pre-sale deadline on ticket sales. The number of discount ticket sales decreases significantly when replacing the one-hour deadline with a pre-sale deadline that ends on midnight of the previous day.

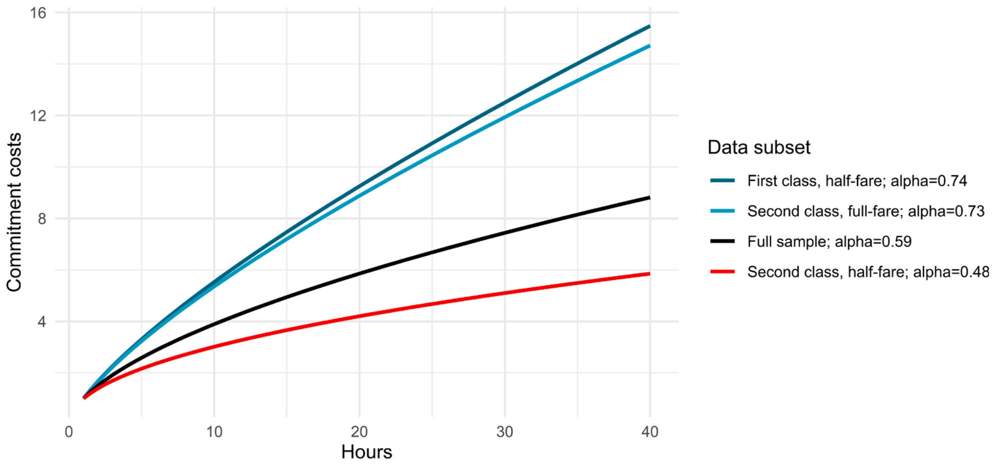

Fig. 6. Illustration of commitment costs with given functional form.

This effect is due to a combination of two things: As the pre-sale deadline is moved to the previous day, the commitment costs associated with purchasing a discount ticket will increase. At the same time, there is a reduction in the number of people that are "treated", as prospective passengers that seek to purchase a ticket on the same day do not see the discount if the pre-sale deadline is set to the end of the previous day. To identify commitment costs, we instead rely on the results from estimating eq. (6) (see section 5.2 below).

The pre-sale deadline effect as measured on the train level is substantially weaker for first-class ticket buyers than for the full sample. The number of tickets sold barely reacts to a change in the pre-sale deadline in this subset, which suggests an early purchase of these tickets. Similar to the price effect, the effect of the pre-sale deadline is much stronger in the evening than in the morning. This is intuitive, as the additional opportunities to purchase a discount ticket when the pre-sale deadline is reduced from the previous day to one hour are naturally greater for a train in the evening than in the morning. The effect is also stronger for *shoulder trains*, indicating that if customers have the opportunity to buy on short notice, they will take advantage of trains departing close to peak. Last, the effect of reducing the pre-sale deadline is greater if the discount is large (70%), relative to when it is small (see Fig. 5).

Table 7 shows the results of the analyses with regular ticket sales as the dependent variable. These cross-price elasticities indicate quite high "windfall effects" for customers in the sense that on average, 52% of the increase in discount ticket sales comes at a reduction in the sale of regular tickets.[19] Conversely, about half of the increase in discount tickets constitutes a net increase in ticket sales. Since we do not have information about transport choices on the individual level, we cannot identify the part that is due to new trip generation (i.e., trips that would not have taken place in the absence of the discount) vs. a shift away from driving. All we can say in the presence of unpriced externalities (which are greater for driving than for the train; see Hintermann et al. (2021)) is that the welfare effects from the discount tickets programme will increase with the share of the mode shift and decrease with the share of additionally generated trips. Moreover, the welfare gain will be greater if the mode shift takes place during peak hours vs. during off-peak due to the presence of road congestion.[20]

5.2. Commitment costs

Fig. 6 displays different average values of α as computed by expression (5) and the corresponding commitment cost curves with the experimental data and the given functional form assumption. Among half-fare customers, those that buy first-class tickets have a higher α and thus a steeper commitment cost curve than those who buy second-class tickets. This can be explained by a higher willingness to pay for first-class customers and the associated lower sensitivity to price changes. This fact translates into a higher relative cost of committing to a specific train. The full-fare customers of the second-class exhibit an α almost as high as for first-class ticket buyers. Since full-fare customers are mostly occasional train users, it can be assumed that they also have other options and are

[19] The values reported in Tables 6-7 are elasticities. To convert them into absolute ticket sales, we need to multiply by the number of regular and discount tickets, which vary across treatment periods. During weeks with a 30%-discount, 17,328 discount tickets and 73,345 regular tickets were sold. During weeks with a 70%-discount, the corresponding numbers are 28,437 discount tickets and 67,406 regular tickets. Taking these values as the extremes leads to a "windfall" range of 40%-72%. The value of 52% reported in the main text is computed using the average number of ticket sales.

[20] Assuming that all of the additional train trips are due to a mode shift from driving, a net increase in discount ticket sales by, e.g., 10,000 could be due to a "direct" shift of that many drivers that would have driven during off-peak hours (but now are taking the train during these hours). However, it could also be due to 10,000 drivers taking the train during peak hours and an equal number of previous train peak customers shifting to off-peak hours, and any convex combination of these two extremes. Because we do not know the time at which regular train tickets are used, we cannot quantify the "timing" of the mode shift.

Table 8
Baseline estimations for the effect on the purchase time for the full sample and subsets on the train level.

	Dependent variable: log(Time difference)							
	Full sample	Only first-class	Only second-class	Only half-fare	Only full-fare	Only morning	Only evening	Only *shoulder*
Discount	0.746**	0.928**	0.718**	0.679**	1.302**	0.627**	0.788**	0.671**
	(0.046)	(0.072)	(0.049)	(0.048)	(0.135)	(0.083)	(0.109)	(0.104)
Pre-sale deadline: 1 day	0.124**	0.139**	0.124**	0.123**	0.151**	0.009	0.162**	0.030
	(0.019)	(0.039)	(0.020)	(0.021)	(0.037)	(0.033)	(0.040)	(0.034)
Observations	11,490	3,813	7,677	7,668	3,822	2,676	1,992	1,509
R^2	0.332	0.314	0.325	0.400	0.233	0.406	0.383	0.420

Note: Standard errors (in parentheses) corrected for heteroskedasticity and clustered at the line level. **: $p < 0.01$, *: $p < 0.05$, °: $p < 0.1$. The data covers the period of August 19 to September,15, 2019.

therefore less willing to commit themselves to a discounted ticket.

The regression results from estimating eq. (6) are shown in Table 8. The effect of the price on the purchase time is significant and has the expected sign: the higher the discount, the earlier the purchase time, all else equal. This is consistent with the notion of commitment costs discussed in Section 3.2. Note that since we control for the pre-sale deadline in this regression (and thus for the number of "treated" customers), this effect is entirely driven by commitment costs.

At a discount level of 30%, an increase in the discount by 10 percentage points leads to an increase in the purchase time of 7.5%, or around 7 hours (the average purchase time is 4.15 days before departure).

The response of a discount change is on average stronger for evening train tickets. This may be explained by customers trying to lock in a high discount on the previous day (or even earlier), as purchasing it on the morning of travel risks the chance of not obtaining a discount. Since these trains depart later, the time difference between purchase and departure is automatically greater.

The effect of a discount increase on the purchase time is stronger for first-class tickets (compared to second-class). Finally, half-fare ticket buyers are less responsive to a discount change with regard to the purchase time than their full-fare counterpart. This can be explained by the fact that the discount, in Swiss Francs, is twice as high for full-fare than half-fare tickets. This larger discount is able to "buy" a larger commitment cost, all else equal (note that the commitment cost is driven by both d and p in eq. (5)).

5.3. Discussion

In our experiment, changes in the level of discount cause slightly larger effects than changes in the pre-sale deadline. An increase in the discount by 40 percentage points leads to 26.7% more sold tickets, whereas a reduction in the pre-sale deadline from the previous day to one hour increases the number of tickets by 18–30% which is equivalent to a price increase by 26–43%.

Huber et al. (2022) find a propensity to reschedule a trip of 0.16% at the individual level when the discount rate increases by 1 percentage point. In contrast to our study, they only have data on individuals who participated in the survey and who would have also bought a ticket at the regular price (so-called always takers). Since we have data at the train level and based on a field experiment, we can also consider the more price-sensitive individuals who would not have bought a ticket at the regular price. In the meta-analyses of Paulley et al. (2006), Holmgren (2007) and Hensher (2008), own-price elasticities of between −0.2 and −0.75 are reported. The estimates without considering peak hours are slightly higher and range from −0.3 to −0.8. Our own estimates are therefore at the upper range of what has been reported in the literature.

As shown by McCollom and Pratt (2004), own-price elasticities depend to a large extent on the operating environment, the type of services, the characteristics of customers and the overall market. The fact that our estimates are at the upper end of the literature can be explained by a higher price-sensitivity of the customers in the experiment, which excludes two important groups. First, holders of flat-rate subscriptions cannot take advantage of the discount tickets (although it is clear that purchasing, or not, a flat-rate subscription is also a choice that may be affected by the presence of the discount tickets programme). This group includes mainly regular commuters, many of whom may be constrained by their working hours (see also Litman, 2004). Second, some customers may be able to switch within peak hours (e.g., from a train with an occupancy rate of 95% to a train with an occupancy rate of 65%), but not all the way into shoulder or off-peak trains. This customer group is also excluded as discounts are only available for trains with an occupancy of less than 60%. This group may also be less price-elastic than those that are able to travel outside of peak hours.

6. Conclusions

We examine how changes in the characteristics of train-specific tickets for long-distance trains affect the demand for them. The underlying data come from a field experiment, which makes the results more plausible compared to studies based on policy changes or stated preferences approaches. We induce an exogenous variation about both the price (or discount) and the pre-sale deadline. The combined analysis of these characteristics allows for a direct comparison of their effect on ticket sales in public transport.

The different measures likely target different audiences. Whereas a decrease in the price potentially affects all customers, reducing the pre-sale deadline from a day to an hour leads to an increase in ticket purchases on the day of travel only. Whether offering "same-day" discounts is desirable depends on the goals (maximization of welfare vs. profits) and the characteristics of the transport system (e.

g., the level of road congestion and the likelihood of modal substitution).

The standardised conditions of the evaluated transit and ticket sales system as well as the field experiment setting ensure a high internal validity of the results. On the other hand, the study was rather short and the experimental treatments varied by week, such that customers may not have had sufficient time to learn about the conditions. Furthermore, it took place only on four lines (three of which were used for the analysis), and the discount offer does not apply to holders of flat-rate subscriptions which provide the majority of the ridership. In this sense, the study only provides a snapshot of a small part of the Swiss public transport system. A longer observation period, the consideration of additional lines and the inclusion of a broader base of customers would help to increase the external validity of the findings.

Although we do not observe the ridership in peak trains in our data, the high own-price elasticities for *shoulder train* tickets suggests that peak shifting is an important contributor for our results. Moreover, the shifting potential is particularly high for tickets for first-class and for evening trains. In addition to peak smoothing objectives, variable prices and sales conditions are becoming progressively more important to consolidate sales in view of the growing competition from other means of transport. Our results offer a starting point for the design of new selling terms as well as for further research in the field of multidimensional sales conditions.

CRediT authorship contribution statement

Christoph Thommen: Investigation, Methodology, Formal analysis, Data curation, Writing – original draft, Visualization, Project administration. **Beat Hintermann:** Conceptualization, Resources, Writing – original draft, Supervision, Funding acquisition.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

The data will be made available if the article is accepted, subject to the transport company's approval. This may or may not be possible, but if it is we will be happy to share the data.

Appendix

A. Identification of trains eligible for discount tickets

A.1. Identification of trains

Ideally, we would have a train timetable or supplementary data from the train company, with information on all trains that run during the experiment period. As we only have those trains in the data with at least one discount ticket sold based on the available ticket data, we first have to identify all trains that ran during the experiment period. In order to do this, we first check whether and how the timetables differ on weekdays for all lines. It can be shown that the timetable for weekdays from Monday to Thursday is always the same for the lines examined and that there may be deviations on Fridays for some lines. On Saturdays and Sundays, trains run at different times than on the remaining days of the week. On this basis, we create groups for each line and all days of the experiment where we assume that trains left at the same times. For Mondays to Thursdays, the groups consist of the weekdays of the same week (without Fridays). For the remaining days (i.e., Fridays, Saturdays and Sundays), weekdays of different weeks are assigned to different groups. In the following, we assume that, based on the available departure times, trains run equally on each day in a specific group.

This strategy bears the risk that there are also trains included that only run on a single day or for a specific occasion. In order to identify those trains, a further loop is necessary. Here, we first focus on trains that only appear once, i.e. on a single day within a group. Then we check whether these trains also appear in so-called control groups. Control groups in this context are groups of days that are very similar to the group under investigation (e.g., groups that cover Monday to Thursday of a given line of two or more consecutive weeks). If a train is observed only once in one group but does not appear at all in the corresponding control groups, we assume that this is a so-called extra train. Accordingly, this train most probably only ran on a specific day, but not on the days that are in the same group. The 206 extra trains we found, are then discarded for the process of the train timetable completion.

A.2. Categorization of trains.

After having completed the hypothetical train timetable, we categorize the included trains. First, we exclude "slow" trains and trains that detour. The former may run from the origin to the destination of a predefined line, but their main purpose is to bring customers from the origin of a line to less frequented stops or to bring them from less frequented stops to the destination of a line. The second category describes trains that take an alternative route between the origin and the destination.[21] In summary, both types, which

[21] In all analysed cases, one of the two routes is faster.

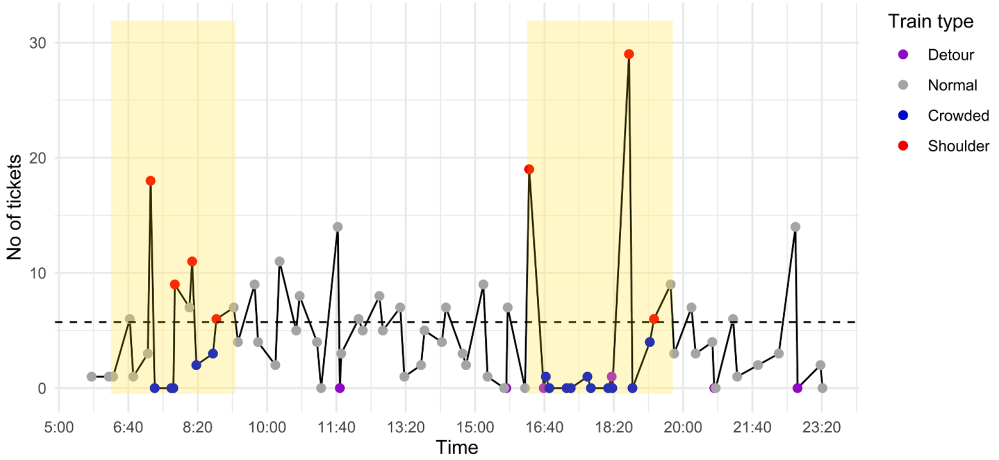

Fig. A1. Number of trains and categorisation line 1, July 26, 2019.

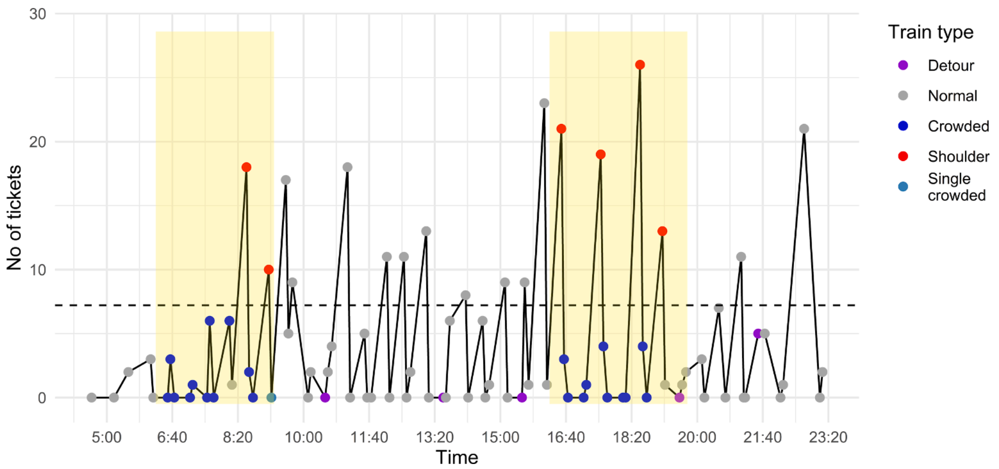

Fig. A2. Number of trains and categorisation line 3, August 7, 2019.

we call *detour trains*, have a much lower demand for the predefined lines due to the characteristics described and cannot be compared with the rest of the trains. For this reason, these trains are discarded for the analyses. To identify them, we resort to the debit number. If for a specific day and line, the average number of discount tickets per train and per debit number is below 15% of the average number of discount tickets per train and per line on that day[22], then these trains are defined as *detour trains*.[23]

After discarding *detour trains* and trains that run on weekends, we identify trains that do not exceed the 60% occupancy forecast at some point in time, which we call *crowded trains*.[24] At a first stage, *crowded trains* have to fulfil the mandatory requirement of a departure time within a peak period. In a default setting, the peak periods include trains that depart between 6.15 am and 9.15 am and between 4.15 pm and 8.45 pm.[25] In addition, at least one of the following four conditions have to be fulfilled:

- No sale of discount tickets.
- Below daily average sale of discount tickets per train and at least one *neighbouring train*[26]

[22] Alternative threshold values do not alter the estimates relevantly.
[23] For the calculation of the daily average sale of tickets only those trains are considered for which at least one discount ticket is sold.
[24] Unfortunately, there is no information available on which trains were eligible for discount tickets during the experiment period from the train company.
[25] Alternative definitions of peak periods are used as robustness checks.
[26] A *neighbouring train* is a train directly following or preceding in the timetable. Consequently, a second-order *neighbouring train* is a train before the preceding and after the following train.

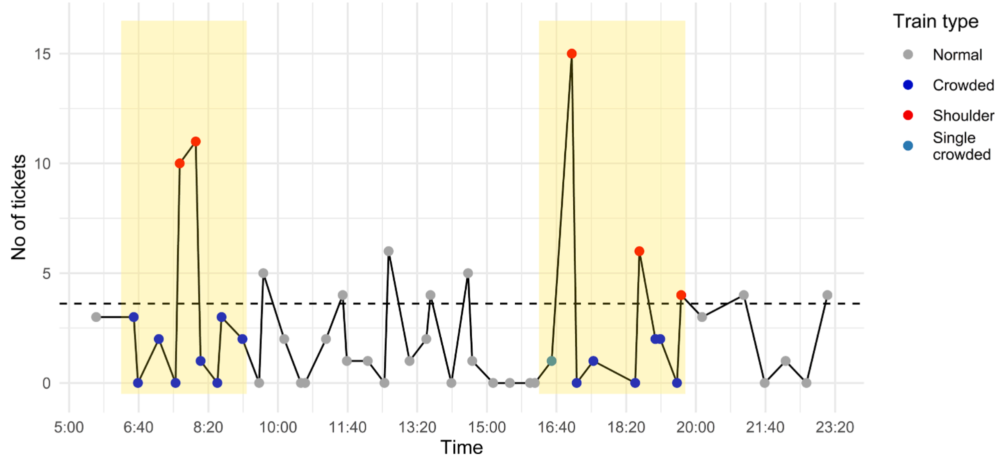

Fig. A3. Number of trains and categorisation line 4, September 5, 2019.

without a discount tickets sold.

- Below daily average sale of discount tickets per train, at least one second-order *neighbouring train* without a discount ticket sold and a lower sale of discount tickets per train on the *neighbouring train* in between.
- Below daily average sale of discount tickets and last discount ticket sale two or more days before the travel date.

Crowded trains are discarded because we assume that they exceeded the 60% occupancy forecast threshold during the sales period. Figure A1, A2 and A3 depict some examples of classifications for specific lines and days during the course of a day. They serve as a visual confirmation that the steps presented here lead to a plausible classification of trains.

B. Latin squares approach

As outlined in Section 2.2, the experimental set-up was chosen, among other things, such that the outcomes of the experiment can be evaluated with the non-parametric *Latin squares* method (see Cochran and Cox, 1992). Accordingly, the number of lines must correspond to the number of periods and the number of treatments. This approach requires a special sequencing of the different settings, which in turn are randomly assigned. It also requires the use of four lines for the four treatments, such that we have to include line 2 (for which the experimental conditions were not always implemented correctly, as discussed in the main text).

The following analysis allows a first assessment of the effects of the different sales settings and thus of the trade-off between a price decrease and a pre-sale deadline reduction. With this approach, a statement can be made about the direction and the significance of a difference between the outcomes of different settings. However, no conclusion can be made on the absolute size of an effect. In the present case, the outcomes of the four settings are compared on an aggregated weekly level. The four experiment settings can a priori be clearly ranked in terms of their attractiveness for buyers, with the exception of B vs. C.

Table B1
Summary statistics of the number of sold discount ticket and the share of discount tickets among all tickets.

Lines	Weeks								Sum	
	34		35		36		37			
	Total	Share	Total	Share	Total	Share	Total	Share	Total	Share
1	5,855	0.275	4,743	0.202	3,796	0.177	2,843	0.121	17,237	0.775
2	1,062	0.327	1,423	0.394	491	0.157	775	0.240	3,751	1.118
3	4,653	0.270	3,543	0.200	7,504	0.383	5,590	0.294	21,290	1.148
4	1,055	0.211	1,440	0.251	1,939	0.319	2,811	0.436	7,245	1.216
Sum	12,625	1.084	11,149	1.046	13,730	1.036	12,019	1.091	49,523	4.257

	Sum		Weekly mean	
Settings	Total	Share	Total	Share
A	17,593	1.488	4,398	0.372
B	13,334	1.142	3,333	0.286
C	10,664	0.939	2,666	0.235
D	7,932	0.688	1,983	0.172

Table B2
Application of *Latin squares* approach.

Settings	Differences	
	Total	Share
A/B	1,065*	0.087**
A/C	1,732**	0.137**
A/D	2,415**	0.200**
B/C	668	0.051*
B/D	1,351**	0.113**
C/D	683	0.063**
Standard error	292	0.013
Standard error for difference between two means	414	0.019
Critical value stat. sign. 1% (t-value: 2.947)	1,219	0.055
Critical value stat. sign. 5% (t-value: 2.131)	882	0.040
Critical value stat. sign. 10% (t-value: 1.753)	725	0.033

Table B3
ANOVA Total discount tickets.

	deg. of freedom	SS	MS	F
Rows (Lines)	3	50,951,603	16,983,868	49.64**
Columns (Weeks)	3	882,051	294,017	0.86
Settings	3	12,703,710	4,234,570	12.38**
Error	6	2,052,923	342,153	
Total	15	66,590,288		

Table B4
ANOVA Share discount tickets.

	deg. of freedom	SS	MS	F
Rows (Lines)	3	0.029	0.010	13.79**
Columns (Weeks)	3	0.001	0.0001	0.27
Settings	3	0.004	0.001	40.56**
Error	6	0.004	0.001	
Total	15	0.120		

Table C1
Baseline estimation of full sample and subsets to ticket types on the day level.

	Dependent variable: log(# tickets)				
	Full sample	Only first-class	Only second-class	Only half-fare	Only full-fare
log(price)	−0.703**	−1.118**	−0.697**	−0.702**	−0.767**
	(0.027)	(0.063)	(0.027)	(0.028)	(0.064)
Pre-sale deadline: 1 day	−0.299**	−0.186**	−0.302**	−0.300**	−0.259**
	(0.024)	(0.046)	(0.025)	(0.026)	(0.046)
Observations	480	240	240	240	240
R^2	0.975	0.905	0.966	0.972	0.951

Note: Standard errors (in parentheses) corrected for heteroskedasticity and clustered at the line level. **: $p < 0.01$, *: $p < 0.05$, °: $p < 0.1$.

Table C2
Baseline estimation of full sample and subsets to ticket types on the train level.

	Dependent variable: log(# tickets)					
			Quasi-poisson			*Negative binomial*
	Full sample	Only first-class	Only second-class	Only half-fare	Only full-fare	Full sample
log(price)	−0.667**	−1.018**	−0.662**	−0.662**	−0.759**	−0.654**
	(0.038)	(0.093)	(0.038)	(0.040)	(0.102)	(0.031)
Pre-sale deadline: 1 day	−0.183**	−0.015	−0.186**	−0.181**	−0.163**	−0.188**
	(0.035)	(0.069)	(0.036)	(0.037)	(0.062)	(0.028)
Observations	11,490	3,813	7,668	7,677	3,822	11,490
θ						10.98
R^2	0.727	0.432	0.714	0.698	0.504	0.544

Note: Standard errors (in parentheses) corrected for heteroskedasticity and clustered at the line level. θ describes the applied dispersion parameter to correct for overdispersion. **: $p<0.01$, *<0.05, °:$p<0.1$.

Table C3
Estimations with fixed settings on the day level.

	Dependent variable: log(# tickets)				
	Full sample	Discount 30%	Discount 70%	Pre-sale deadline:1 day	Pre-sale deadline:1 hour
log(price)	−0.703**			−0.616**	−0.751**
	(0.026)			(0.034)	(0.036)
Pre-sale deadline: 1 day	−0.299**	−0.245**	−0.294*		
	(0.024)	(0.042)	(0.033)		
Observations	480	240	240	240	240
R^2	0.975	0.976	0.983	0.976	0.973

Note: Standard errors (in parentheses) corrected for heteroskedasticity and clustered at the line level. **: $p < 0.01$, *: $p < 0.05$, °: $p < 0.1$.

Table C4
Estimations with fixed settings on the train level.

	Dependent variable: log(# tickets)				
	Full sample	Discount 30%	Discount 70%	Pre-sale deadline:1 day	Pre-sale deadline:1 hour
log(price)	−0.667**			−0.502**	−0.799**
	(0.038)			(0.039)	(0.056)
Pre-sale deadline: 1 day	−0.183**	−0.103*	−0.232**		
	(0.035)	(0.044)	(0.048)		
Observations	11,490	5,061	6,429	5,397	6,093
R^2	0.727	0.756	0.766	0.790	0.760

Note: Standard errors (in parentheses) corrected for heteroskedasticity and clustered at the line level. **: $p < 0.01$, *: $p < 0.05$, °: $p < 0.1$.

Table C5
Robustness check estimations to full sample on the day level.

	Dependent variable: log(# tickets)			
	Baseline	Incl. Sat & Sun	Incl. Line 2	Incl. Sat & Sun and Line 2
log(price)	−0.703**	−0.695**	−0.704**	−0.697**
	(0.026)	(0.029)	(0.027)	(0.028)
Pre-sale deadline: 1 day	−0.299**	−0.333**	−0.296**	−0.329**
	(0.024)	(0.024)	(0.023)	(0.023)
Observations	480	672	636	888
R^2	0.975	0.964	0.976	0.966

Note: Standard errors (in parentheses) corrected for heteroskedasticity and clustered at the line level. **: $p < 0.01$, *: $p < 0.05$, °: $p < 0.1$.

In terms of numbers of tickets sold, as shown in Table B1, A is ahead of B ahead of C ahead of D. Table B2 indicates whether the differences between the settings are statistically significant. There is no significant difference between B and C (nor between C and D). In contrast, if the share of discount tickets among all tickets is considered, all differences are statistically significant and correspond to the expected order. This means that changes in the remaining sale of regular tickets are also indirectly taken into account and thus represent a more comprehensive estimation.

It can be concluded that setting B with a discount of 70% and a pre-sale deadline of the previous day induces a higher demand than setting C with a discount of 30% and a one-hour pre-sale deadline. This implies a higher average evaluation of the increase in discount by 40%-points, starting at a discount of 30% compared to a shortening of the pre-sale deadline from the previous day to one hour. The analysis of variance (ANOVA) in Table B3 and B4 indicate that for both the total tickets and the shares, settings as well as the lines (due to their different sizes) are significant determinants of the overall variation. The week, however, as desired by the design, has no significant effect.

The results from this Latin squares approach provide a first indication of the valuation of the commitment period and also serves as a robustness check for subsequent analyses. In the main results section, this relationship is examined in more detail.

References

Anderson, M.L., 2014. Subways, strikes, and slowdowns: the impacts of public transit on traffic congestion. Am. Econ. Rev. 104 (9), 2763–2796.

Beheshtian, A., Geddes, R.R., Rouhani, O.M., Kockelman, K.M., Ockenfels, A., Cramton, P., Do, W., 2020. Bringing the efficiency of electricity market mechanisms to multimodal mobility across congested transportation systems [Developments in Mobility as a Service (MaaS) and Intelligent Mobility]. Transportation Research Part A: Policy and Practice 131, 58–69. https://doi.org/10.1016/j.tra.2019.09.021.

Cochran, W.G., Cox, G.M., 1992. Experimental designs. Wiley press, New York.

Currie, G. (2010). Quick and effective solution to rail overcrowding: Free early bird ticket experience in Melbourne, Australia. *Transportation research record*, *2146* (1), 35–42.

Daniels, R., Mulley, C., 2013. The paradox of public transport peak spreading: Universities and travel demand management. International Journal of Sustainable Transportation 7 (2), 143–165.

de Grange, L., González, F., Muñoz, J.C., Troncoso, R., 2013. Aggregate estimation of the price elasticity of demand for public transport in integrated fare systems: the case of Transantiago. Transp. Policy 29, 178–185.

De Palma, A., Lindsey, R., Monchambert, G., 2016. Optimal transit pricing with crowding and traffic congestion: A dynamic equilibrium analysis. Canadian Transportation Research Forum 51st Annual conference-North American transport challenges in an era of change.

De Palma, A., Lindsey, R., Monchambert, G., 2017. The economics of crowding in rail transit. J. Urban Econ. 101, 106–122.

Douglas, N. J., Henn, L., & Sloan, K. 2011. *Modelling the ability of fare to spread am peak passenger loads using rooftops* (Conference paper). 34th Australasian Transport Research Forum (ATRF), Adelaide.

Halvorsen, A., Koutsopoulos, H.N., Lau, S., Au, T., Zhao, J., 2016. Reducing subway crowding: analysis of an off-peak discount experiment in Hong Kong. Transp. Res. Rec. 2544 (1), 38–46.

Henn, L., Karpouzis, G., Sloan, K., 2010. *A review of policy and economic instruments for peak demand management in commuter rail* (Conference paper). 33th Australasian Transport Research Forum (ATRF), Canberra.

Hensher, D.A., 2008. Assessing systematic sources of variation in public transport elas- ticities: Some comparative warnings. Transp. Res. A Policy Pract. 42 (7), 1031–1042.

Hintermann, B., Schoeman, B., Molloy, J., Götschi, T., Castro, A., Tchervenkov, C., Tomic, U., & Axhausen, K. W. 2021. Pigovian transport pricing in practice.

Holmgren, J., 2007. Meta-analysis of public transport demand. Transp. Res. A Policy Pract. 41 (10), 1021–1035.

Hörcher, D., Graham, D.J., 2018. Demand imbalances and multi-period public transport supply. Transp. Res. B Methodol. 108, 106–126.

Huber, M., Meier, J., Wallimann, H., 2022. Business analytics meets artificial intelligence: assessing the demand effects of discounts on Swiss train tickets. Transp. Res. B Methodol. 163, 22–39.

Khwanpruk, S., U-tapao, C., Khwanpruk, K., Laokhongthavorn, L., Moryadee, S., 2021. A study on high-speed rail pricing strategy for Thailand based on dynamic optimal pricing model. Int. J. Intelligent Eng. Systems 14 (3), 97–108.

Litman, T., 2004. Transit price elasticities and cross-elasticities. J. Public Transportation 7 (2), 3.

LTA, 2020. Launch of travel smart journeys to distribute peak hour travel demand. Retrieved January 10, 2021, from https://www.lta.gov.sg/content/ltagov/en/newsroom/2020/2/news- releases/launch- of- travel- smart- journeys- to- distribute- peak-hour-travel-d.html.

Ma, Z., Basu, A. A., Liu, T., & Koutsopoulos, H. N. 2019. *Behavioral response to transit demand management promotions: Sustainability and implications for optimal promotion design* (Conference paper). Transportation Research Board 98th Annual Meeting, Washington D.C.

McCollom, B. E., & Pratt, R. H. 2004 *Traveler response to transportation system changes. Chapter 12-Transit Pricing and Fares* (Working Paper). Transportation Research Board of the National Academies, Washington, D.C.

Ortega-Hortelano, A., Guzman, A.F., Preston, J., Vassallo, J.M., 2016. Price elasticity of demand on the high-speed rail lines of Spain: impact of the new pricing scheme. Transp. Res. Rec. 2597 (1), 90–98.

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.

Santos Silva, J.M.C., Tenreyro, S., 2006. The log of gravity. Rev. Econ. Stat. 88 (4), 641–658.

SFSO, 2020a. Infrastruktur und Streckenlänge. Retrieved January 10, 2021, from https:// www.bfs.admin.ch/bfs/de/home/statistiken/mobilitaet-verkehr/verkehrsinfrastruktur- fahrzeuge/streckenlaenge.html.

SFSO, 2020b. Passenger transport performance. Retrieved January 10, 2021, from https://www.bfs.admin.ch/bfs/en/home/statistics/mobility-transport/passenger-transport/performance.html.

Tirachini, A., Hensher, D.A., Rose, J.M., 2013. Crowding in public transport systems: effects on users, operation and implications for the estimation of demand. Transportation research part A: policy and practice 53, 36–52.

Van den Berg, V., Kroes, E., Verhoef, E.T., 2009. *Choice of train ticket: A Study on Dutch travellers* (Working paper). VU University, Amsterdam, Netherlands.

Vickrey, W.S., 1969. Congestion theory and transport investment. Am. Econ. Rev. 59 (2), 251–260.

von Neumann, L.J., Morgenstern, O., 1953. Theory of Games and Economic Behavior. Princeton University Press, Princeton.

Wardman, M., Murphy, P., 2015. Passengers' valuations of train seating layout, position and occupancy. Transp. Res. A Policy Pract. 74, 222–238.

Williams, K.R., 2022. The welfare effects of dynamic pricing: evidence from airline markets. Econometrica 90 (2), 831–858.

Xie, J., Shugan, S.M., 2001. Electronic tickets, smart cards, and online prepayments: When and how to advance sell. Mark. Sci. 20 (3), 219–243.

Yang, H., Tang, Y., 2018. Managing rail transit peak-hour congestion with a fare-reward scheme. Transp. Res. B Methodol. 110, 122–136.

Zhu, G., Zhao, Y., 2020. Research on the joint decision model of high-speed railway dynamic pricing and ticket allocation. IOP Conference Series: Mater. Sci. Eng. 768 (5), 052017.