

Hintermann, Beat; Thommen, Christoph

Working Paper

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

WWZ Working Paper, No. 2022/05

Provided in Cooperation with:

Center of Business and Economics (WWZ), University of Basel

Suggested Citation: Hintermann, Beat; Thommen, Christoph (2022) : Price versus commitment: Managing the demand for off-peak train tickets in a field experiment, WWZ Working Paper, No. 2022/05, University of Basel, Center of Business and Economics (WWZ), Basel, <https://doi.org/10.5451/unibas-ep89789>

This Version is available at:

<https://hdl.handle.net/10419/278182>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.


Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



University
of Basel

Faculty of
Business and Economics



July 2022

Price versus Commitment: Managing the Demand for Off-peak Train Tickets in a Field Experiment

WWZ Working Paper 2022/05

Beat Hintermann; Christoph Thommen

A publication of the Center of Business and Economics (WWZ), University of Basel.

© WWZ 2022 and the authors. Reproduction for other purposes than the personal use needs the permission of the authors.

Universität Basel
Peter Merian-Weg 6
4052 Basel, Switzerland
wwz.unibas.ch

Corresponding Author:
Prof. Dr. Beat Hintermann
Tel. +41 61 207 33 39
b.hintermann@unibas.ch

Price versus Commitment: Managing the Demand for Off-peak Train Tickets in a Field Experiment*

Christoph Thommen[†] Beat Hintermann[‡]

July 15, 2022

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 length of the pre-sale period 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. Extending the pre-sale deadline by one hour leads to an increase in the pre-sale of discount tickets by 2.1%, which is equivalent to a price decrease by 3.1%. Reducing the price by 10% causes customers to purchase the discount ticket 7 hours earlier. Our results help design measures for peak-shifting in transport at least societal cost.

Keywords: Field Experiments, Public Transport Systems, Train, Dynamic Pricing, Switzerland

JEL Codes: L92, R41, L11, C93

*We thank Pierre Chevalier, Reto Lüscher, Meredith Fowlie, Ali Darudi, Renate Strobl, Matthias Niggli, Ingmar Schlecht and participants at the University of Basel SCCER Lunch Seminar for helpful comments and discussions. Magnus Vieten provided excellent research assistance.

[†]Zurich University of Applied Sciences, Switzerland; E-mail: christoph.thommen@zhaw.ch

[‡]University of Basel, Switzerland; E-mail: b.hintermann@unibas.ch

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 is thus 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.

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. 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 by one hour leads to a demand increase by 2.1-2.5 percent, such that an extension by one hour is equivalent to a price discount of 3 percent. 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 (so-called *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. (2015) show that these conclusions also apply to transit

systems, despite the additional capacity constraints compared to road networks.¹ 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 at least in some cases has proven to be more effective than discounts during off-peak hours (Douglas et al., 2011; Paulley et al., 2006).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. (2015) further argue that a large share of workers lack the flexibility to change their trip times. Daniels and Mulley (2013) argue that the potential to shift work trips to earlier morning hours is restricted by the presence 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 Swiss transport provider 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

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

²Implementations 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).

³The 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.

⁴Xie and Shugan (2001) argue that dynamic pricing approaches satisfy different price sensitivities and

There are theoretical papers that study optimal pricing schemes based on several pre-sale periods (Zhu and Zhao, 2020 and Kankanit and Moryadee, 2021), but 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 and pre-sale deadlines. 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, there is no previous study besides Huber et al. (2022) 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, the estimates could be used to maximize profits (or minimize losses) when providing transport services. From a social point of view, peak shifting will lead to a better use of the available capacity of both road and rail transport. This will help to meet the transport demand in the future and to decarbonize the transport sector in the long run.

thereby target different customer groups. Accordingly, late bookers tend to travel for business reasons, while early bookers can be characterised as leisure travellers.

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

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 in August and September 2019.

2.1 The “saver tickets” programme

The Swiss railway network includes 5,200 kilometres of rails (SFSO, 2020a). Trains belonging to different transport providers cover a distance of 550,684 kilometres (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 binding 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.⁶

To increase the occupancy rate of off-peak trains, one of the largest state-owned rail service providers introduced discount or so-called “saver” tickets.⁷ 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

⁵This figure is based on data from the largest transport provider, which covers around 75% of all passengers in the Swiss transit system.

⁶Additional room in the trains during peak hours will lead to a convex combination of the following: less crowding, generation of new trips (i.e., trips that would not have taken place in the absence of the peak shift) and a shift from private to public transport during peak hours. Since all three outcomes are desirable from a societal point of view, their combination will also be welfare-improving.

⁷These are called “Sparbillete” in German, “Billets dégriffés” in French and “Biglietti risparmio” in Italian.

purchased a certain time period before departure. Discount tickets sales account for about a quarter of the tickets sold for individual long-distance trips.⁸

Discount tickets are available only for long-distance trains.⁹ Furthermore, in order for a trip to be eligible for a discount ticket, an occupancy forecast of 60% must not be exceeded.¹⁰ 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 schemes can further differ across the four main ticket categories (first class/half-fare, first class/full-fare, second class/half-fare and second class/full-fare).¹¹ 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

⁸80% of long-distance passengers hold a flat-rate transport 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.

⁹Technically, 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.

¹⁰This 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.

¹¹Discount 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%.

2.2 Design of the experiment

For the experiment, the transport company selected four train lines covering distances between approximately 50 and 90 kilometres 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 exogenously varying the level of the discount, we can identify the effect of a *change* in the discount level and the pre-sale deadline.

To create an exogenous variation, 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 vs. midnight of the previous day).

Table 1: Experiment settings

		Discount	
		70%	30%
Pre-sale deadline	One hour	A	C
	Previous day	B	D

Settings A and D have the most and least attractive conditions, respectively. 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, we ended up analyzing it using more standard econometric techniques, such that the design symmetries were not used.

For Line 2, the discount level temporarily deviated from the experimental settings, and different segments were priced differently by mistake. Although the influence might not be

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

large, we removed Line 2 from the baseline analysis.

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 the “regular” conditions, i.e., before the settings were adjusted for the experiment. 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 change of settings and the first train in the experiment, the availability restriction was never binding.


3 Data

The following information is collected and processed for each discount ticket: Sale date, travel date, departure time, train number, origin and destination, half/full-fare, class, effective price and regular price. For regular tickets, no information on the travel date is available.¹³ But according to the train company, around 90% of tickets are bought on the day of travel.

Discount tickets data are available for the calendar weeks 30 to 37 (autumn) of 2019 and, in addition, for calendar weeks 10 to 23 (spring) of 2019 (see Figure 2 for the distribution of discounts for both periods). The analysis is based on the data from autumn.

¹³When 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 therefore not be used together.

Figure 2: Number of discount tickets sold for Lines 1, 3 and 4 between week 10 and 23 (spring) and 30 and 37 (autumn) 2019



We run our analyses separately on the day and on the train level. A line is defined by the route between a unique origin and a unique destination. 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,¹⁴ the number of stops as well as the occupancy forecast.¹⁵ Tables 3 and 4 provide descriptive statistics for the variables used in the regression analyses at the day and the train level.

The last data issue pertains to separating missings from zeros. If we only used data on discount ticket sales we would not be able to differentiate a missing from a zero, because in both instances, there would be no sales associated with a particular train. Therefore, as

¹⁴In 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. 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 per line in the experiment. 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.

¹⁵Forecasts 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 line, direction, departure time, weekday and train number.

Table 3: Number of discount tickets sold during the experiment (day-line-category level)

Line/Direction	Class	Half-fare	Obs.	Mean	SD
1a	1	0	20	3.5	2.0
1a	1	1	20	15.4	6.8
1a	2	0	20	110.0	32.0
1a	2	1	20	210.4	67.5
1b	1	0	20	2.6	1.6
1b	1	1	20	13.2	6.1
1b	2	0	20	92.7	34.0
1b	2	1	20	189.5	78.6
3a	1	0	20	4.4	4.2
3a	1	1	20	19.4	10.3
3a	2	0	20	77.6	27.4
3a	2	1	20	244.1	75.0
3b	1	0	20	4.5	4.2
3b	1	1	20	16.2	9.4
3b	2	0	20	70.2	16.2
3b	2	1	20	219.4	70.7
4a	1	0	20	1.8	1.8
4a	1	1	20	21.1	10.1
4a	2	0	20	19.1	6.7
4a	2	1	20	109.7	42.0
4b	1	0	20	1.4	1.7
4b	1	1	20	10.7	5.3
4b	2	0	20	14.6	4.7
4b	2	1	20	66.8	28.2

a pre-requisite for the analyses on the train level, all trains running on the involved lines during the period of the experiment need to be identified. In a subsequent step, 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.

To obtain the complete set of trains, we developed a procedure that is described in detail in Section B in the Appendix. As a side benefit of this process, the trains that follow or precede non-eligible discount tickets trains can be identified, and we label them as *shoulder trains* (Liu & Charles, 2013). 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.

Table 4: Number of tickets sold during experiment period (train-category level)

Variable	Level	Obs.	Mean	SD
Pre-sale deadline	1 day	5,397	2.12	3.50
Pre-sale deadline	1 hour	6,093	2.62	4.41
Class	1	3,813	0.42	0.88
Class	2	7,677	3.41	4.56
Half-fare	0	3,822	1.85	2.47
Half-fare	1	7,668	2.70	4.57
Line/Direction	1a	2,811	2.28	3.38
Line/Direction	1b	2,376	2.30	3.65
Line/Direction	3a	1,725	3.58	5.86
Line/Direction	3b	1,896	2.90	4.53
Line/Direction	4a	1,392	1.94	3.19
Line/Direction	4b	1,290	1.19	1.92
Type	normal	9,981	2.01	3.07
Type	shoulder	1,509	5.04	7.19
Morning	0	8,814	2.20	3.60
Morning	1	2,676	3.12	5.10
Evening	0	9,498	2.26	3.69
Evening	1	1,992	3.13	5.20
Weighted mean purchase-time		11,490	4.15	6.12
Weighted mean discount		11,490	0.55	0.17
Occupancy forecast		11,490	0.36	0.12
Capacity		11,490	554.32	127.45
Segments		11,490	3.11	2.40

4 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.

Although the experiment was initially set up according to the Latin squares approach, we do not use this methodology in our main results. For the interested reader, the Latin square analysis can be found in Section A in the Appendix.

4.1 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 saver 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:

$$\begin{aligned} \log(\text{no tickets}_{dlc}) = & \beta_1 \log(\text{price}_{dlc}) + \beta_2 \text{pre-sale deadline}_{dlc} + \Gamma_1 \cdot \text{departure date}_d \\ & + \Gamma_2 \cdot \text{weekday}_d \times \text{line/direction}_l \times \text{category}_c + u_{dlc} \end{aligned} \quad (1)$$

The indices are defined as follows: d denotes the day of the experiment, running from 1 to 20; l refers to the line and direction (such that A to B is treated differently to B to A, leading to a total of 6 line-direction combinations);¹⁶ and c measures the four main ticket categories defined by the interaction of full vs. half fare and first vs. second class.¹⁷

To absorb systematic differences in terms of discount tickets sales across days, line/direction and ticket category, we include a series of fixed effects, some in interacted form. All observations are weighted according to the average number of tickets sold in each weekday/line/direction/category-group. The identifying variation is thus derived from within-changes in ticket prices and the pre-sale deadline, which are varied exogenously in the course of the experiment 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. In an alternative specification, this variable is included in continuous form, in which first the weighted average in hours is taken and is then aggregated on the day level.¹⁸ The parameter vectors Γ_1 and Γ_2 contain the coefficients on the dummy variables,

¹⁶We make this distinction as it may be easier to commit to a discount ticket in the morning than for the trip home, as the timing of the latter may be more difficult to foresee.

¹⁷There are additional categories such as tickets for dogs, bicycles or traveling groups, but these are excluded from the analysis.

¹⁸When the pre-sale deadline is set to the previous day, its continuous representation is defined as the num-

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 & Tenreyro, 2006). Because the PPML model is estimated in exponentiated form, it allows for zeros (which would drop out when log-linearizing the equation). This feature is particularly useful for the regressions on the train level (see below).

In additional regressions, we replace the dependent variable in (1) with the number of regular tickets sold. The resulting coefficient β_1 then captures the cross-price elasticity. Because the date of travel is not recorded (see above), we approximate this with the date of purchase.¹⁹

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 this analysis, we use all trains identified as not having exceeded the 60% occupancy forecast threshold during the course of the purchase period (see Section B in the Appendix for details). For these trains, the ticket types “second class/half-fare”, “second class/full-fare” and “first class/half-fare” are considered. The fourth category “first class/full-fare” is

ber of hours between midnight and the departure time of the corresponding observation. For the continuous representation of the pre-sale deadline, this creates additional variance and allows a direct interpretation of the respective coefficient in terms of hours. At the same time, the implied linear functional form also entails a direct relationship with the departure time, since e.g. only observations with departure times between 7.30 and 8.30 am can have a pre-sale deadline of 8 hours. This, in turn, means that the effects do not differ for different departure times, in contrast to the case of a pre-sale deadline of one hour.

¹⁹In principle, it is possible to purchase a tickets 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.

discarded due to too few observations. We use the following regression equation:

$$\begin{aligned}
 \log(\text{no tickets}_{dtc}) = & \beta_1 \log(\text{price}_{dtc}) + \beta_2 \text{pre-sale deadline}_{dt} + \Gamma_1 \text{departure date}_d \\
 & + \Gamma_2 \text{departure time}_t + \Gamma_3 \text{weekday}_t \times \text{line/direction}_t \\
 & \times \text{category}_c + \Gamma_4 \text{forecast}_{dtc} + \Gamma_5 \text{segments}_t + \Gamma_6 \text{capacity}_{dt} + u_{dtc}
 \end{aligned} \tag{2}$$

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

Due to the non-negligible number of zeros in the dependent variable, we use a Poisson pseudo-maximum likelihood (PPML) model. This approach corrects for the over-dispersion in the available data. The PPML model is again estimated in exponentiated form.

As a robustness test, we also employ a negative binomial distribution. The results are qualitatively similar and included in Table C.3 in the Appendix).

4.2 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 thus learn something about the cost of commitment.

Let $U_i(L)$ denote the utility of an individual i for a train trip on Line L from A to B. 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. We assume that the commitment cost increases with the commitment period, but at a decreasing rate. The discount ticket needs to be purchased at $t_i \leq \bar{t}$, where \bar{t} defines the pre-sale

deadline. The utility for a regular trip (U^R) and a discount trip (U^S) is thus given by

$$U_i^R = U_i(L) - p(L) \quad (3)$$

$$U_i^S = \begin{cases} U_i(L) - q_i^{\alpha_i} - (1-d)p(L) & \text{if } t_i \leq \bar{t} \\ 0 & \text{if } t_i > \bar{t} \end{cases} \quad (4)$$

It is assumed that $U_i(L)$ is unrelated to the ticket in use and remains constant for different purchase times. Hence, U^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 a potential customer continuously compares her or his utility of buying a discount ticket with the constant utility of a regular ticket. Figure 3 shows this as stylized form. As time progresses, U^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.




Figure 3: Utility of discount and regular ticket

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 (see von Neumann and Morgenstern, 1953), 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 and the level of discount:

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

4.3 Effect of the discount on the time of purchase

Besides this more conceptual measure, we also analyze the effect of a change in the discount on the time of purchase using a regression analysis on the train level. To do this, we first compute the difference between the time of purchase and the departure time. We then use this time difference as the dependent variable in the following model:

$$\begin{aligned} \log(\text{time difference}_{dtc}) = & \beta_1 \text{discount}_{dtc} + \beta_2 \text{pre-sale deadline}_{dt} + \Gamma_1 \text{departure date}_d \\ & + \Gamma_2 \text{departure time}_t + \Gamma_3 \text{weekday}_d \times \text{line/direction}_t \\ & \times \text{category}_c + \Gamma_4 \text{forecast}_{dtc} + \Gamma_5 \text{segments}_t \\ & + \beta_8 \text{capacity}_{dt} + u_{dtc} \end{aligned} \quad (6)$$

This equation is very similar to (2), the main difference being that we include the discount rather than the log price as the main variable of interest.

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 a bit over 4 days, even though the longest possible pre-sale deadline 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.

The effect measured by β_1 not only incorporates the cost of commitment discussed above, but also the expectations of obtaining a larger discount by purchasing earlier. Although this feature was disabled during the experiment, 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, which provides a reason for purchasing earlier than what would be expected based on pure commitment costs.

4.4 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).

In order for our methodology to identify the causal effect of changing the level of discount and the length of the pre-sale deadline, these changes must not be correlated with other relevant determinants of train ticket sales. For example, if an exogenous shock affecting tickets sales (say, the start of school vacation) coincided with the change in the experimental setting, then this would bias our results. Although there is no untreated control group in our experiment, the fact that the experimental settings were assigned such that each line received a different treatment at any given time implies that common shocks are absorbed to the extent that they affect the various lines in the same way. As a further means of absorbing

shocks to overall demand, we include departure date 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 if our experiment included a proper control group.

We can ignore any heterogeneities induced by variations in pricing and comfort of trains that complicate the examination do not have to be considered: In the Swiss transit system, all trains are priced in the same way. Moreover, the trains being used for a line are generally the same in terms of travel duration and train comfort.

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 e.g. 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 both before and during the experiment.

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.

5.1 Effect on ticket sales

The first three columns in Table 5 show the results from estimating eq. (1) on the day level. The own-price elasticity is -0.70, and this result is robust to different specifications of the pre-sale deadline (columns 2 and 3). The point estimates based on eq. (2) on the train level

are 3-4 percentage points lower, but the confidence intervals overlap such that the results from the different levels of aggregation are largely consistent.

Figure 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 Tables C.2 to C.4 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. Furthermore, discount first-class tickets also attract customers that regularly 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.

Table 5: Baseline estimation with different pre-sale deadline specifications on the day and the train level

	log(No of tickets)					
	Day-level			Train-level		
	(1)	(2)	(3)	(4)	(5)	(6)
log(price)	-0.703** (0.027)	-0.709** (0.027)	-0.707** (0.027)	-0.667** (0.038)	-0.665** (0.037)	-0.663** (0.037)
Pre-sale deadline:1d	-0.299** (0.024)			-0.183** (0.035)		
Pre-sale deadline (cont.)		-0.025** (0.002)	-0.035* (0.014)		-0.021** (0.002)	-0.002 (0.010)
Pre-sale deadline (cont.) ²			0.001 (0.001)			-0.001 (0.001)
Observations	480	480	480	11,490	11,490	11,490
<i>psi</i>				4.048	3.977	3.966
R ²	0.975	0.975	0.975	0.727	0.730	0.729

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

We further compute the price elasticity for trains departing at different times during the course of a day. The own-price-elasticity is lower for customers of morning trains (departure until 10 am) than for 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). Note also that the possibilities to shift departure to an earlier time in the morning are also limited for biological and physical reasons (see Daniels and Mulley, 2013). Similarly, we find that the elasticity is higher for *shoulder trains*. This finding suggests that shoulder trains are a better substitute for the (non-discount) peak trains than a train running significantly before or after peak hours.




Figure 4: Own-price elasticity for the full sample and different subsets

Last, we examine whether the own-price elasticity varies for different pre-sale deadlines. Since the pre-sale deadline of the previous day 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, we expect the relative reaction of a price change to be smaller, as shown in Figure 4 and Table C.5 and C.6 in the Appendix.

Figure 5 and Table 5 shows the effect of varying the pre-sale deadline on ticket sales. On the day and the train level, the proportional effect per hour is between -0.025 and -0.021, implying that reducing the pre-sale deadline by one hour leads to an increase in ticket sales by between 2.1 and 2.5%.




Figure 5: Effect of pre-sale deadline for the full sample and different subsets

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 Figure 5).

The results of the analyses with regular tickets (see table C.1) indicate relatively strong windfall effects of discount tickets. I.e. the additional sales of discount tickets in the event of a price reduction are to some extent at the expense of regular tickets. Based on our estimates, if the discount increases from 30% to 70%, 72.3% of discount ticket purchasers would have also bought a regular ticket if 70%-discount tickets were not available. Conversely, a discount reduction from 70% to 30% leads to 40.5% of the former discount ticket purchasers buying

regular tickets.²⁰ Conversely, these estimates suggest that 28%-59% of the increase in discount tickets are due to *additional* trips as opposed to a shift from regular tickets. However, as we do not have complete data on the departure date for regular tickets (see Chapter 3), these results must be interpreted with caution.

5.2 Effects on the purchase time

Figure 6 displays different average values of α as computed by (5) and the corresponding commitment cost curves with the experimental data and the given functional form assumption. For regular ticket buyers, the two utility curves for purchasing a regular or a discount ticket depicted in Figure 3 do not intersect. Because we only consider actual discount ticket purchases for the illustrative example, the outside options are not realized.




Figure 6: Illustration of commitment costs with given functional form

As can be seen in Figure 6, 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

²⁰Since we do not have information about the trains that were used with the regular tickets, we cannot measure to what extent these shifts from regular to discount tickets represent a shift from peak to offpeak trains.

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 therefore less willing to commit themselves to a discounted ticket.

The regression results from estimating eq. (6) are shown in Table 6. 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 4.2.²¹

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).

Table 6: Baseline estimations for the effect on the purchase time for the full sample and subsets on the train level

	log(Time difference)							
	Full sample	Only first class	Only second class	Only half-fare	Only full-fare	Only morning	Only evening	Only <i>shoulder</i>
Discount	0.746** (0.046)	0.928** (0.072)	0.718** (0.049)	0.679** (0.048)	1.302** (0.135)	0.627** (0.083)	0.788** (0.109)	0.671** (0.104)
Pre-sale deadline	0.124** (0.019)	0.139** (0.039)	0.124** (0.020)	0.123** (0.021)	0.151** (0.037)	0.009 (0.033)	0.162** (0.040)	0.030 (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 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

²¹As explained above, the coefficient may also capture customers trying to “lock in” a discount level. Although the discount remained fixed during the experiment, the regular discount ticket programme usually contains such a feature, such that regular users of discount tickets presumably expect the discount to decrease over time.

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 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.

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 does increase the number of tickets by 18.3%. If we evaluate the changes in the pre-sale deadline in monetary terms, an increase of the pre-sale deadline by one hour has the same effect on discount ticket sales as a price decrease of about 3 percent.

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 a bit 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 ticket 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 time.

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 subscription 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.

References

- Anderson, M. L. (2014). Subways, strikes, and slowdowns: The impacts of public transit on traffic congestion. *American Economic Review*, *104*(9), 2763–96.
- 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/https://doi.org/10.1016/j.tra.2019.09.021>
- Cochran, W. G., & Cox, G. M. (1992). *Experimental designs*. New York: Wiley press.
- 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. *Transport 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, Toronto, May 1-4, 2016*.
- De Palma, A., Lindsey, R., & Monchambert, G. (2015). *The economics of crowding in public transport* (Working paper). University of British Columbia, Vancouver, Canada.
- 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. *Transportation Research Record*, *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 elasticities: Some comparative warnings. *Transportation Research Part A: Policy and Practice*, *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. *Transportation Research Part A: Policy and Practice*, *41*(10), 1021–1035.
- Huber, M., Meier, J., & Wallimann, H. (2022). Business analytics meets artificial intelligence: Assessing the demand effects of discounts on swiss train tickets. *Transportation Re-*

- search Part B: Methodological*, 163, 22–39. <https://www.sciencedirect.com/science/article/pii/S0191261522001059>
- Kankanit, S. K. C. U.-t., & Moryadee, K. L. L. S. (2021). A study on high-speed rail pricing strategy for thailand based on dynamic optimal pricing model. *International Journal of Intelligent Engineering and Systems*, 14(3), 97–108.
- Litman, T. (2004). Transit price elasticities and cross-elasticities. *Journal of Public Transportation*, 7(2), 3.
- Liu, Y., & Charles, P. (2013). *Spreading peak demand for urban rail transit through differential fare policy: A review of empirical evidence* (Conference paper). 36th Australasian Transport Research Forum (ATRF), Brisbane.
- LTA. (2020). Launch of travel smart journeys to distribute peak hour travel demand. Retrieved January 20, 2021, from. 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.
- McCullom, 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. *Transportation Research Record*, 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. *Transport policy*, 13(4), 295–306.

- Santos Silva, J. M. C., & Tenreyro, S. (2006). The Log of Gravity. *The Review of Economics and Statistics*, 88(4), 641–658. %7Bhttps://ideas.repec.org/a/tpr/restat/v88y2006i4p641-658.html%7D
- 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>
- 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. *American Economic Review*, 59(2), 251–260.
- von Neumann, L. J., & Morgenstern, O. (1953). *Theory of Games and Economic Behavior*. Princeton: Princeton university press.
- Xie, J., & Shugan, S. M. (2001). Electronic tickets, smart cards, and online prepayments: When and how to advance sell. *Marketing Science*, 20(3), 219–243.
- Yang, H., & Tang, Y. (2018). Managing rail transit peak-hour congestion with a fare-reward scheme. *Transportation Research Part B: Methodological*, 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: Materials Science and Engineering*, 768(5), 052017.

Appendix

A 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. 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 and C.

Table A.1: Summary statistics of number of sold saver tickets and share of saver 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

Settings	Sum		Mean	
	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

In terms of numbers of tickets sold, as shown in Table A.1, A is ahead of B ahead of C ahead of D. Table A.2 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 saver tickets among all tickets is considered, all differences are statistically significant and correspond to the expected order, as can be seen in the right-hand column in Table A.2. This means that changes in the remaining sale of regular tickets are also indirectly taken into account and thus represent a more comprehensive estimation.

Table A.2: Application of *Latin squares* approach

Settings	Differences	
	Total	Share
A/B	1065*	0.087**
A/C	1732**	0.137**
A/D	2415**	0.200**
B/C	668	0.051*
B/D	1351**	0.113**
C/D	683	0.063**
<hr/>		
Standard error	292	0.013
Standard error for difference between two means	414	0.019
Critical value stat. sign. 1% (t-value: 2.947)	1219	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 A.3: ANOVA Total saver tickets

	d.f.	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 A.4: ANOVA Share saver tickets

	d.f.	SS	MS	F
Rows (Lines)	3	0.029	0.010	13.79**
Columns (Weeks)	3	0.001	0.0001	0.266
Settings	3	0.004	0.001	40.56**
Error	6	0.004	0.001	
Total	15	0.120		

Under this consideration, 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 result is equivalent to 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. As Tables A.3 and A.4, in both analyses, the influence of the settings as well as the lines (due to their different sizes) are significant. The week, however, as desired by the design, has no significant effect.

This result gives 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.

B Identification of trains eligible for saver tickets

In a first step, we 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, 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. For example, 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 probably only ran on a specific day, but not on the days that are in the same group. Extra trains are then discarded for the process of the train timetable completion.

After having completed the hypothetical train timetable, we characterise 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.²² In summary, both types, which we call *detour trains*, have a much lower demand for the predefined lines due to the characteristics

²²In all analysed cases, one of the two routes is faster.

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 saver tickets per train and per debit number is below 15% of the average number of saver tickets per train and per line on that day, then these trains are defined as *detour trains*. These trains are discarded from further analyses.²³

After completing the train timetable and discarding *detour trains*, we identify trains that do not exceed the 60% occupancy forecast at some point in time, which we call *crowded trains*. 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.²⁴ In addition, one of the following four conditions have to be fulfilled:

- No sale of saver tickets
- Below daily average sale of saver tickets per train and at least one *neighbouring train*²⁵ without a saver ticket sold.
- Below daily average sale of saver tickets per train, at least one second-order *neighbouring train* without a saver ticket sold and a lower sale of saver tickets per train on the *neighbouring train* in between.
- Below daily average sale of saver tickets and last saver ticket sale two or more days before the travel date.

As can be seen in Table B.1, there are 7,017 trains in the data. In addition, there are 3,328 trains for which no saver tickets are sold that were added in order to have a complete set of train observations. Among these resulting 10,345 trains, 1,718 trains are defined as *crowded trains* (i.e., trains that are not eligible for saver tickets because their occupancy forecast is

²³For the calculation of the daily average sale of tickets only those trains are considered for which at least one saver ticket is sold.

²⁴Alternative definitions of peak periods are used as robustness checks.

²⁵A *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.

Table B.1: Data processing stages

	In the data	7,017 trains (thereof 127 <i>extra trains</i>)
Complementary trains (without saver tickets sold)		3,328 trains
		10,345 trains
	Weekend trains	- 3,334 trains
	<i>Detour trains</i>	- 1,463 trains
	<i>Crowded trains</i>	- 1,718 trains
	For the analysis	3,830 trains
	<i>Normal trains</i>	3,327 trains
	<i>Shoulder trains</i>	503 trains

too high) and are excluded from further analyses. Trains that are significantly slower due to an alternative route choice and trains that act as 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. Of the remaining 3,830 trains, 503 are defined as *shoulder trains* and the rest are *normal trains*. Figures B.1, B.2 and B.3 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.




Figure B.1: Number of trains and categorisation line 1, July 26, 2019




Figure B.2: Number of trains and categorisation line 3, August 7, 2019




Figure B.3: Number of trains and categorisation line 4, September 5, 2019

C Own-price elasticities and pre-sale deadline effects

Table C.1: Baseline estimation with effect on regular tickets of full sample and subsets to ticket types on the day level

	log(No of tickets)				
	Full sample	Only first class	Only second class	Only half fare	Only full fare
log(Price)	0.120** (0.016)	0.087** (0.033)	0.123** (0.017)	0.121** (0.017)	0.141** (0.042)
Pre-sale deadline	0.083** (0.019)	0.077** (0.024)	0.084** (0.020)	0.087** (0.020)	0.030 (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. **: $p < 0.01$, *: $p < 0.05$, °: $p < 0.1$.

Table C.2: Baseline estimation of full sample and subsets to ticket types on the day level

	log(No of tickets)				
	Full sample	Only first class	Only second class	Only half-fare	Only full-fare
log(Price)	-0.703** (0.027)	-1.118** (0.063)	-0.697** (0.027)	-0.702** (0.028)	-0.767** (0.064)
Pre-sale deadline	-0.299** (0.024)	-0.186** (0.046)	-0.302** (0.025)	-0.300** (0.026)	-0.259** (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 C.3: Baseline estimation of full sample and subsets to ticket types on the train level

	log(No of tickets)					
	<i>Quasi-poisson</i>					<i>Negative binomial</i>
	Full sample	Only first class	Only second class	Only half-fare	Only full-fare	Full sample
log(Price)	-0.667** (0.038)	-1.018** (0.093)	-0.662** (0.038)	-0.662** (0.040)	-0.759** (0.102)	-0.654** (0.031)
Pre-sale deadline	-0.183** (0.035)	-0.015 (0.069)	-0.186** (0.036)	-0.181** (0.037)	-0.163** (0.062)	-0.188** (0.028)
Observations	11,490	3,813	7,668	7,677	3,822	11,490
ϕ	4.048	0.434	5.882	4.206	3.161	
θ						10.98(0.259)
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. **: $p < 0.01$, *: $p < 0.05$, °: $p < 0.1$.

Table C.4: Baseline estimation of full sample and subsets (by departure time) on the train level

	log(No of tickets)			
	Full sample	Only morning	Only evening	Only <i>shoulder</i>
log(Price)	-0.667** (0.038)	-0.571** (0.058)	-0.951** (0.068)	-0.815** (0.068)
Pre-sale deadline	-0.183** (0.035)	-0.001 (0.046)	-0.437** (0.043)	-0.178** (0.053)
Observations	11,490	2,676	1,992	1,509
ϕ	4.048	4.902	4.109	8.640
R ²	0.727	0.786	0.851	0.761

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

Table C.5: Estimations with fixed settings on the day level

	log(No of tickets)				
	Full sample	Discount 30%	Discount 70%	Pre-sale deadline:1 day	Pre-sale deadline:1 hour
log(Price)	-0.703** (0.026)			-0.616** (0.034)	-0.751** (0.036)
Pre-sale deadline	-0.299** (0.024)	-0.245** (0.042)	-0.294** (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 C.6: Estimations with fixed settings on the train level

	log(No of tickets)				
	Full sample	Discount 30%	Discount 70%	Pre-sale deadline:1 day	Pre-sale deadline:1 hour
log(Price)	-0.667** (0.038)			-0.502** (0.039)	-0.799** (0.056)
Pre-sale deadline	-0.183** (0.035)	-0.103* (0.044)	-0.232** (0.048)		
Observations	11,490	5,061	6,429	5,397	6,093
ϕ	4.048	3.159	4.140	3.217	4.004
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 C.7: Robustness check estimations to full sample on the day level

	log(No of tickets)			
	Baseline	Incl. Sat & Sun	Incl. Line 2	Incl. Sat & Sun and Line 2
log(Price)	-0.703** (0.026)	-0.695** (0.029)	-0.704** (0.027)	-0.697** (0.028)
Pre-sale deadline	-0.299** (0.024)	-0.333** (0.024)	-0.296** (0.023)	-0.329** (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. **: $p < 0.01$, *: $p < 0.05$, °: $p < 0.1$.