KIEL
WORKING PAPER

Striking evidence: The impact of railway strikes on competition from intercity bus services in Germany

No. 2251 June 2023

Matthias Beestermöller, Levke Jessen-Thiesen, Alexander Sandkamp

Kiel Institute for the World Economy
ISSN 1862–1155

ABSTRACT

STRIKING EVIDENCE: THE IMPACT OF RAILWAY STRIKES ON COMPETITION FROM INTERCITY BUS SERVICES IN GERMANY*

Matthias Beestermöller, Levke Jessen-Thiesen, and Alexander Sandkamp

This paper investigates the impact of the largest rail strikes in German history on intercity buses – a then newly liberalised market. Using unique booking data of bus services, we exploit variation in rail service cancellations across routes to show that the disruption in rail transport increases bus ticket sales. Crucially, the effect persists beyond the strike, indicating that travellers do not return to their originally preferred mode of transport. It is particularly pronounced for passengers travelling on weekends. The findings suggest that customers were previously under-experimenting. Beyond transportation, our results highlight the importance of service reliability, as temporary disruptions can cause customers to permanently switch to competitors.

Keywords: Experimentation, inter-modal substitution, learning, optimisation, strike, switching costs, transport

JEL classification: C81, D83, L92, R41

Matthias Beestermöller

University of Munich (LMU)
Geschwister-Scholl-Platz 1
D-80539 München, Germany

www.lmu.de

Levke Jessen-Thiesen

University of Kiel (CAU) and
Kiel Institute for the World Economy
Kiellinie 66
D-24105 Kiel, Germany
Email:
levke.jessen-thiesen@ifw-kiel.de
www.ifw-kiel.de

Alexander Sandkamp
(Corresponding author)
University of Kiel (CAU),
CESifo, KCG and
Kiel Institute for the World Economy
Kiellinie 66
D-24105 Kiel, Germany
Email:
alexander.sandkamp@ifw-kiel.de
www.ifw-kiel.de

* An earlier version circulated as “Striking Evidence? Demand Persistence for Inter-City Buses from German Railway Strikes” (Beestermöller, 2017). This paper has been greatly improved by comments and suggestions from Daniel Baumgarten, Carsten Eckel, Florian Englmaier, Katharina Erhardt, Gabriel Felbermayr, Lisandra Flach, Lionel Fontagné, Inga Heiland, Julian Hinz, Anna Koukal, Anthonin Levelu, Chen Li, Lukas Mergele, Ferdinand Rauch, Till Requate, Jens Ruhose, Klaus Schmidt, Monika Schnitzer, Fabian Schrey, Vincent Stamer, Colin Vance, Lu Wei, Gerald Willmann, Yoto Yotov as well as numerous colleagues and seminar participants at the MGSE Colloquium and seminars in Munich, Kiel, and Paris, the RGS doctoral conference in Bochum, the BiGSEM Workshop, and the EPCS Meeting. Many thanks to MeinFernbus Flixbus GmbH (especially Tina Bosler and Anna Humpert) for providing data. Beestermöller gratefully acknowledges financial support from the German Research Foundation (DFG) through GRK 1928 and the Egon-Sohmen-Foundation.

1 Introduction

Once settled on a product, consumers do not usually revisit their choice for every purchase. After all, you do not evaluate whether your preferred e-mail provider, regular hairdresser, or favourite restaurant constitutes the optimal choice every time you make use of their services. Such behaviour might be a rational response to search costs and informational frictions, or it could be driven by inertia. Incumbent firms benefit from those habits, as they constitute a barrier to entry for new firms. Service interruptions, however, may force customers to experiment with substitutes, inducing some of them to permanently switch to competitors. For example, if a specific streaming service suddenly becomes unavailable, customers might switch to a rival and do not return even after the interruption has ended.

Persistent service reliability would consequently be an important factor for companies in maintaining their customer base. We test this hypothesis by exploiting a major strike in the German railway network as a natural experiment to investigate whether the resulting service disruptions caused customers to permanently switch from rail to bus travel. This study is, to the best of our knowledge, the first to present systematic evidence of this mode-of-transport-switch.

The railway strike is particularly well suited to investigate the impact of service disruptions on product switching because passenger transport is an experience good, which can only be properly evaluated during or after consumption. Barriers to switching transport mode may thus be especially high. In the fall of 2014, labour disputes shut down nearly all German long-distance trains for days, forcing train travellers to use alternative transport modes. For some travellers, this led to a first encounter with intercity buses. In introducing new customers to the railway's key rival, the strike potentially resulted in new, long-term customers for buses.[1] The German railway strike of 2014 provides several desirable features for a quasi-natural experimental setting: It was unprecedented in extent, announced in the short term, and exogenous to transport services in its geographic dispersion. In addition, it was the first German railway strike in which buses were available as a viable alternative.

We employ a unique data set: detailed booking data provided by one of the major intercity bus providers, *MeinFernbus* (MFB). The data set contains the universe of MFB ticket sales between any combination of 33 large German cities over a period of four months. We match this data with web-crawled rail itineraries and emergency timetables. For any of the 33 cities, we know how they are connected by rail at each hour of the day. This includes information on the length of the trip, the required changeovers, and the frequency of departures - both in normal times and during strikes.

[1] Intercity buses are defined as regularly scheduled services exceeding a distance of 50 km. In the literature, they are often interchangeably referred to as 'inter-urban' or 'long-distance' buses.

In a difference-in-differences setting, we estimate the effect of the railway strike on intercity bus use. In doing so, we exploit the fact that not all routes are affected equally by the strike. We show that bus ticket sales indeed increase during the strike (by 32% on average during the first wave). This is primarily driven by first-time customers who did not use the bus before the strike. Crucially, the effect persists beyond the strike, implying that several passengers permanently switched to bus travel after having been forced to experiment with new modes of transport. After the strike, ticket sales on affected routes are on average 8% higher. The post-strike effect is particularly pronounced for customers travelling over the weekend, indicating that buses are mostly an alternative for passengers travelling for leisure. Positive price effects suggest that the increase in ticket sales is the result of an increase in demand due to the strike and not driven by a positive supply shock (i.e. capacity expansion linked to a price offensive). Finally, we show that switching mainly occurs on relatively short routes, where buses are arguably a better substitute for trains.

The paper contributes to several strands of literature. First, we add to the classic literature relating to the way in which individuals decide between alternatives. There is a long-standing debate on rational decision-making (Simon, 1955; Weitzman, 1979; Morgan and Manning, 1985) and constraints such as search costs (Baumol and Quandt, 1964; Ben-Akiva and Morikawa, 1990), information imperfections or simply habits, specifically when it comes to the choice of transport mode (Moser et al., 2018; Donna, 2021). A permanent increase in demand for intercity bus services could also be explained by learning. Travellers may learn about the service and quality of buses by actually testing and experiencing them. Bus travel may thus be seen as an experience good, the quality of which is underestimated, by consumers ex-ante (Riordan, 1986; Bergemann and Välimäki, 2006).

Klemperer (1987c) observes that welfare may be reduced in the presence of switching costs. Consequently, Porter (1996) argues that exogenous shocks may help individuals to their optimal choice by triggering a period of experimentation. The underlying idea of experimentation due to exogenously-imposed constraints, such as the non-availability of rail services, applies to the setting in this paper. After all, bus services were available before the strikes and the availability of online bookings - the primary booking channel - remedied some of the search costs. Public transport strikes are typically considered to be highly economically damaging (Kennan, 1986). However, if the rail strike reveals information, it may actually be welfare-improving.

Shapiro (1983) and Villas-Boas (2006) argue that firms can address under-experimentation by customers through low introductory prices. We show that temporary service disruptions of rivals can have a similarly positive impact on sales. Our paper thus also adds to the industrial organisation literature surrounding competition in the presence of switching costs (Von Weizsäcker, 1984; Klemperer, 1987a,b,c, 1995). These are relevant in many industries, including the markets for pay-TV (Weeds, 2016), pharmaceuticals (Janssen, 2023) and even the deposit market (Zephirin, 1994). Switching costs are also relevant for effective policy design (Giulietti et al., 2005). We add to this literature by showing that service disruptions of incumbent firms can have pro-competitive effects, as rivals can permanently lure away some of the incumbent's customers. McDonald and Bloch (1999) study the impact of strikes on the profitability of competing firms. They show that firms can benefit from industrial action that disrupts a competitor's production, although effects decline with industry concentration. We show that such effects are present in an industry that is highly concentrated and dominated by a single company.

Our paper also relates to the literature investigating determinants of the choice of transport mode. This is a classical research subject in transport research going back to Train (1978).[2] Several studies have investigated how mode choice changes when one transport option is temporarily unavailable during a strike. A survey by van Exel and Rietveld (2001) reviews 13 such works and finds that a switch to the car is the most common, and that sometimes a permanent mode switch occurs that persists after the strike. More recently, Bauernschuster et al. (2017) exploit public transport strikes as a quasi-experimental shock to assess the negative externalities of increased car traffic.[3]

Anderson (2014) shows that a strike in Los Angeles increased road traffic congestion. Adler and van Ommeren (2016) and Adler et al. (2021) confirm this effect for Rotterdam and Rome, respectively. The authors conclude that public transport plays a major role in avoiding road overcrowding. Yang et al. (2022) provide evidence that the use of bike-sharing schemes in London increased during Tube strikes. These studies focus on urban transport and commuter mode choice and mostly investigate contemporary effects. One of the few studies investigating intercity travel is Yeung and Zhu (2022). The authors show that the number of booked seats with BlaBlaCar - an intercity ride-sharing app - increased by 33% during a railway strike in France. Our paper contributes to the literature by providing an estimate for the extent of the strike-driven modal switch from rail to bus in intercity transport during as well as after a strike.

For perfectly informed consumers, a strike should be just a temporary disruption, after which they return to their optimal mode of transport. However, Larcom et al. (2017) find that after a 2014 strike on the London Underground, up to 5% of travellers permanently changed their commuting route. These commuters seem to have been stuck with a suboptimal route before the strike forced them to experiment and discover a better option. Their results touch on Goodwin (1977), who states that "the traveller does not carefully and deliberately calculate

[2]See for example Small and Verhoef (2007) for an overview on urban transportation.

[3]In contrast, Chen and Whalley (2012) investigate a positive shock on the transport network, showing that the opening of a new rail transit system in Taipei reduced local air pollution. On the negative side, expanding transportation networks may contribute to the spread of viral diseases (Adda, 2016). Allen and Arkolakis (2022) analyse the welfare impacts of improvements in transportation infrastructure against the backdrop of traffic congestion.

each morning anew whether to go to work by car or by bus". While Larcom et al. (2017) provide evidence for permanent route switching within the London Transport Network, we show that customers even switch to competing transport networks (from rail, serviced by Deutsche Bahn to bus, serviced by MFB). Our findings thus also relate to Fung et al. (2021), who show that the number of bike-sharing trips in Glasgow persistently increased following temporary closure of the subway. We show that there is persistence of strike-induced modal switches also in intercity transport, suggesting that there had been under-experimentation between rail and bus on long-distance routes.

Finally, this paper is among the first of a small but growing literature studying the German market for long-distance buses. The primary concern of this literature has been to study the impact of the market liberalisation of German buses on rail ticket prices and services. Böckers et al. (2015) and Evangelinos et al. (2015) find that the effect on the rail network was larger at the periphery of the network. Bataille and Steinmetz (2013) provide theoretical models on the effect of liberalisation. These studies of inter-modal competition relate to a slightly older literature on the entry of low-cost airlines into Germany in the early 2000s (Friebel and Niffka, 2009). Durr et al. (2015) study competition within the intercity bus market, and estimate the price effect of a large merger of MFB and Flixbus (see also Gagnepain et al. (2011) for a more general review of bus market competition). Neither of these studies pays attention to the recent German railway strikes. The empirical literature relies on data from price comparison websites and usually offers few time-series observations. The descriptive statistics from MFB booking data presented here, therefore, contribute a much-improved insight into this young market and its dynamics.

The remainder of this paper is structured as follows: Section 2 sketches a model that motivates our empirical work. Section 3 describes the railway strikes in 2014 in more detail. Section 4 introduces the data sets and provides new descriptive statistics on the intercity bus market. Section 5 presents our empirical strategy and discusses estimation challenges. Section 6 provides the main results, while Section 7 reports robustness tests. Section 8 concludes.

2 Conceptual framework

The conceptual framework guiding our investigation is based on Donna (2021), who models the role of switching costs in the choice of transport mode. Preferences of passengers can be represented with the following equation:[4]

[4]This is a simplified and slightly altered version of Equation (1) in Donna (2021). Note that Donna (2021) models the use of public transport and car use, whereas we focus on two different types of public transport, namely rail and bus services.

$$V_{irt}(m) = U_{irt}(m) - p_{rt}^{m} - \phi^{m}\mathbf{1}\{m_{irt-1} \neq m_{irt}\} \tag{1}$$

with $m \in \{rail, bus\}$, where $V_{irt}(m)$ is the net utility derived by passenger i from travelling with transport mode m on route r at time t. $U_{irt}(m)$ is the gross utility of travelling with mode m and captures travel time (and thus traveller specific time cost) of mode m and other individual preferences such as utility derived from the existence of dining cars or wifi. We allow $U_{irt}(m)$ to vary over time to capture the fact that characteristics such as travel time and frequency of service change during the strike, thus affecting a traveller's gross utility derived from using rail services.

p_{rt}^{m} is the price of travelling with mode m on route r at time t. $\phi^{m} > 0$ captures switching costs to mode m and $\mathbf{1}\{\cdot\}$ is an indicator function, ensuring that switching costs are only incurred if a passenger switches transport mode between periods $t-1$ and t. Switching costs include search costs such as having to look up information on bus schedules, routes or transfers.

In the model, passengers who travelled by rail in period $t-1$ will only switch to buses if $E[V_{irt}(bus)] > V_{irt}(rail)$, i.e. the expected value of switching (net of switching costs) is larger than the value of sticking to the current mode. We add an expectation term to capture the fact that transportation is an experience good. If a passenger has not used bus services before (i.e. $m_{irt-1} \neq m_{irt}$), he or she forms an expectation over $U_{irt}(bus)$. The true $U_{irt}(bus)$ is, however, only revealed once travelling by bus is experienced. Formally, passengers who previously used rail services will switch mode iff:

$$E[U_{irt}(bus)] - p_{rt}^{bus} - \phi^{bus} > U_{irt}(rail) - p_{rt}^{rail} \tag{2}$$

A switch can occur either through a change in relative prices (not the focus of our paper) or by a change in $U_{irt}(m)$.[5] During the strike, the gross utility derived from using rail, $U_{irt}(rail)$, diminishes. Ceteris paribus, this means that the expected value of using buses for some passengers now exceeds the value of rail transportation, inducing them to switch modes. After the strike, we assume that $U_{irt}(rail)$ returns to its original value. Passengers who switched to buses during the strike now face a new decision problem. They will continue using the bus iff:

$$U_{irt}(bus) - p_{rt}^{bus} > U_{irt}(rail) - p_{rt}^{rail} - \phi^{rail} \tag{3}$$

Inequality (3) differs from Inequality (2) in two important aspects. First, in Inequality (2), switching costs are incurred if a passenger chooses to travel by bus for the first time. In Inequality (3), switching costs occur if passengers choose to travel by train. The net utility of travelling by bus is thus higher in Inequality (3) than in Inequality (2). Note that this is true

[5] In Section 7, we show that bus ticket prices on routes affected by the strike increased during the strike. Our findings are thus not driven by a fall in prices for bus tickets.

even if ϕ^{rail} is close to zero, which probably is the case for those passengers who had used the train before. For some passengers, in the absence of switching costs, the value of using buses starts to exceed the value of using trains, leading them to permanently use bus services.

Second, in light of the literature on inertia and experience goods discussed above, it is also possible that $U_{irt}(bus) > E[U_{irt}(bus)]$. During the strike, passengers who are forced to experiment with buses, realise that they are a better substitute for rail transport than previously expected. The true utility from using buses (captured by Inequality 3) is thus higher than the expected utility from using buses (Inequality 2). There will thus be passengers for whom Inequality (2) does not hold, while Inequality (3) holds. These are passengers that used to travel by train, switched to buses during the strike (assuming a sufficient temporary drop in $U_{irt}(rail)$) and continue using buses also in its aftermath.

To sum up, two predictions can be derived from the model: First, the number of passengers using bus services increases during the strike following the temporary decline in $U_{irt}(rail)$. Second, some of this change persists even after the strike because ϕ^{bus} disappears once a switch occurred and because $U_{irt}(bus) > E[U_{irt}(bus)]$, at least for some passengers. In the remainder of the paper, we show that passenger numbers for bus services on affected routes indeed increase during the strike and that some of this change persists even after rail services resumed.

3 The German railway strikes of 2014-2015

High-speed railway services in Germany are predominantly provided by one publicly-owned firm: *Deutsche Bahn AG* (DB). The company controls the railway infrastructure and is legally shielded from competition. In the year 2013, it became legal to offer bus services in direct competition with existing railway connections.[6] Private bus operators entered the passenger transport market, offering intercity connections on routes more than 50 km apart.

The locomotive drivers' union (*Gewerkschaft Deutscher Lokomotivführer*; hereafter referred to as GDL) is relatively small but powerful and has a long history of disputes with DB. The 2014-2015 negotiations, however, constituted the most ferocious industrial action in the history of DB. Two factors contributed to the ferocity of the dispute: GDL was in a power struggle with a rival union, and new legislation was under review which threatened GDL's right to represent service personnel in future wage negotiations. Between September 2014 and May 2015, the dispute resulted in nine strike waves and 22 days affected by strikes – 354 hours of service

[6]The market was liberalised by law as of January 2013. Previously, the Passenger Transport Act only permitted intercity bus services if the state-owned railway company was unable to provide an acceptable service. Durr et al. (2015) provide more details on the liberalisation.

disruptions. Because of the importance of the rail network to the economy, the dispute was followed closely by the German media and the public.[7]

We study the effects of two major waves in October and November 2014.[8] We disregard strikes after January 2015 as this is when MFB merged with rival competitor Flixbus. In addition, we disregard minor warning strikes, as they only lasted a few hours and were announced with many days advance warning. Our data suggest that those strikes were too short to have any measurable impact on the bus market. Figure 1 shows the timeline of disruptions caused by the three strike waves relevant to our study. In the week between October 13th and October 19th of 2014, strikes disrupted rail services on Wednesday and Thursday, as well as on Saturday and Sunday. In the week between November 3rd and November 9th, rail services were cancelled due to the strike from Thursday, November 6th, onwards for three consecutive days.

Figure 1: Timeline of rail strike in weeks October 13-20 and November 03-10, 2014

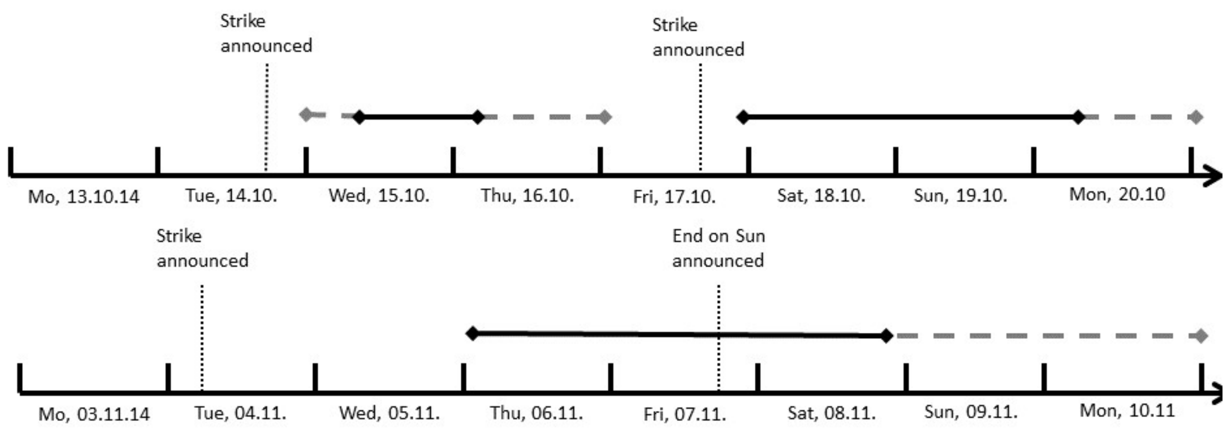

Note: Grey dashed lines indicate strike-related service disruptions. Disruptions started before the first strike wave because DB adopted its emergency timetables at the beginning of the departure day to minimise the overall impact of the strike. Disruptions lasted beyond the duration of each strike wave as it took time to return to normal timetable operations. The third rail strike wave ended prematurely on Saturday, although it had initially been announced to last until Sunday. Following public pressure, GDL announced it would return to work on Sunday, November 9^{th} to allow travellers to reach the anniversary festivities of the Fall of the Berlin Wall around the country.

The timing of the strikes was arguably exogenous. Strikes result from a breakdown of negotiations, the exact timing of which is unpredictable, as negotiations often collapse quickly and unexpectedly. Once negotiations have broken down, the exact timing of a strike remains unclear. It could be delayed by days, weeks, or months if the parties were hopeful of making progress or political pressure was exerted. The trade union centrally decides to go on strike

[7]This paper is concerned with passenger transport. Note, however, that the railway strikes affected both passenger and freight services by DB.

[8]Table A.3 provides a detailed account of the 2014/2015 public transport strikes.

after consulting its members. Importantly, there is no evidence to suggest that competition from buses played any role in the occurrence, timing, or length of the strikes. The strikes can be considered an exogenous positive demand shock to the German bus market. Having reached a decision, GDL usually announced strikes at short notice to maximise their impact. Each strike was announced no more than two days in advance.

GDL called for a strike nationwide. However, neither did GDL shut the network down entirely, nor were rail routes exposed to the same degree. GDL membership strength is weaker in West Germany because many West German train drivers have civil servant status – a relic of DB's historical status as a state company.[9] The emergency timetables operated during the rail strike reflect the varying power of GDL across Germany. The regional disparity in the change of service frequency specified in the emergency timetables was arguably exogenous to the bus market. DB did not strategically focus rail services on routes which were under particular threat of competition from buses. The emergency timetables were the same in all strike waves in 2014-2015, and they are almost identical to those employed by DB in the last railway strikes of 2007-2008; i.e. long before the liberalisation of the intercity bus market in 2013.

Switching between rail and bus is rather easy.[10] Bus terminals are located directly next to the rail station in most cities (Guihéry et al., 2016). Tickets can be bought online or on the bus. Travellers could arrive at the rail station and easily transfer to intercity buses when the implications of the rail strike became clear to them.

4 Data and descriptive statistics

This paper combines data from three sources: detailed booking data for intercity buses provided by MFB, DB emergency timetables, and a data set of all rail itineraries. The latter data are collected using a web crawler linked to the website of a leading price comparison website. We combine the emergency timetables and travel itineraries to create a data set of service cancellations and expected delays caused by the rail strike.

MFB booking data

MFB is Germany's largest bus provider in the sample period, with a market share of then roughly 50%. In addition to being the key player in the German intercity bus market, MFB's

[9] German civil servants have by law no right to strike or unionise.

[10] DB does not offer season passes on specific routes. It offers the *BahnCard* which grants fixed price reductions to cardholders. *BahnCard* subscriptions can be cancelled annually. This may have locked travellers in to DB services, in which case any lasting effect beyond the strike would not be visible until the medium or long term.

service quality as well as strategic use of local bus partners are representative of the entire intercity bus industry.[11]

Figure 2: Map and list of German cities in the sample

Cities:	
Augsburg	Heidelberg
Berlin	Karlsruhe
Bonn	Kassel
Braunschweig	Kiel
Bremen	Leipzig
Cologne	Mainz
Dortmund	Magdeburg
Dresden	Mannheim
Duesseldorf	Munich
Erfurt	Muenster
Essen	Nuremberg
Frankfurt (Main)	Rostock
Freiburg	Saarbruecken
Goettingen	Stuttgart
Hamburg	Ulm
Halle (Saale)	Wuerzburg
Hanover	

The data set provided by MFB contains the universe of MFB ticket sales between any combination of 33 large German cities for departure dates from August 27^{th} to December 16^{th} 2014. Figure 2 lists and maps all 33 cities in the sample. Any booking for a departure between these 33 cities is included, regardless of when the booking was made. The original data set contains about 2.2 million observations. Not all possible combinations of the 33 cities are actually routes served by bus services. Some routes are only served on weekdays or not at all. We restrict our sample to routes that were served by MFB during the strike.

A booking observation includes detailed information on the bus service such as the route, price, departure date, and time as well as an anonymised e-mail address under which the booking was made. The e-mail address identifies first-time and repeat bookings by the same account, and thus allows following a customer over time. The key variable of interest is the

[11] For example, free wifi, luggage allowance, and legroom are almost identical across the industry. See Dürr et al. (2015) for a detailed introduction and comparison of players in the intercity bus market.

natural logarithm of the number of tickets sold at the route and departure day level.[12] We aggregate the individual bookings at the route and departure day level – the unit of analysis in this paper.[13] A route is the combination of an origin- and destination-city, meaning that different routes may be served by the same bus journey. For example, a bus ride from Munich to Berlin with a stop in Dresden serves three routes: Munich–Dresden, Munich–Berlin, and Dresden–Berlin.

While rail strikes continued beyond the sample period to May 2015, we restrict the sample period to 2014. This is because MFB unexpectedly merged with rival bus provider Flixbus in January 2015. Any changes after this date may be driven by the effects of the merger and not the rail strike. The final panel contains 312 routes and roughly 35,000 observations at the route and departure day level. The data set is balanced in the sense that all routes are observed over the entire sample period and through all strike waves.[14] MFB entered the market with an aggressive pricing strategy, where the cheapest tickets sell at only €4.39 for a one-way trip. The average ticket price is €14.5 (median: €12). The maximum price is €63, but only 1% of tickets sell at a price higher than €46. An average bus ticket costs €3.5 per scheduled hour of travel.[15]

The 2.29 million individual ticket sales we observe in the MFB data correspond to roughly 1.5 million ticket orders. In some cases, the order is made several months before the actual departure. On average, an order precedes a departure by around eight days. During the railway strike, we would expect an increase in tickets bought on short notice. Figure 3 compares cumulative bookings before departure for a day affected by a railway strike with a typical booking curve. The dashed vertical line indicates the moment of the strike announcement for the third strike wave on November 07, 2014.[16] As is apparent, ticket sales only diverge from their usual trend after the rail strike was announced. The small sales departure from the usual trend before the announcement suggests that a few travellers booked bus tickets after negotiations had broken down, but before the strike was announced; i.e. very few travellers anticipated the strike. If travellers book bus tickets for departure days before the strike in anticipation, our results would be downward biased. Figure 3 provides strong descriptive evidence that rail strikes drove the peak in ticket sales on striking days.

[12]The dependent variable is computed as $\ln(1 + tickets\ sold)$ at the route departure day level. This approach allows us to keep route-day observations with zero tickets sold. In the data set, zero observations only account for 0.3% of tickets sold and 7% of tickets sold to new customers.

[13]Note that there are two time dimensions to each individual booking: the date of booking and the date of departure. We aggregate ticket sales to the route and departure date dimension. 95% of bus travellers arrive on the same date as they depart.

[14]Note that, as a consequence, there are route-day combinations with zero ticket sales in our panel.

[15]A back of the envelope calculation: A DB *Sparpreis* (saver ticket) at €19 for the maximum travel distance of 250km at 200km/h travel speed would yield a per hour price of €15.2.

[16]See Figure 1.

Figure 3: Mean cumulative bookings for Friday departures

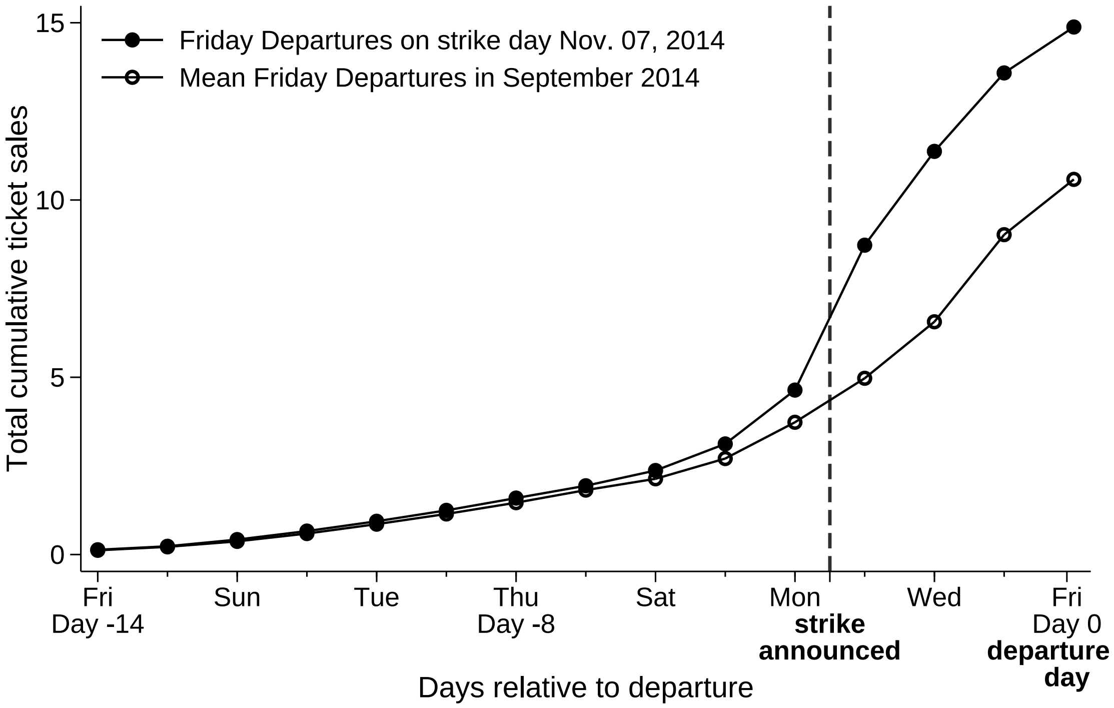

Note: Data are split into bookings for Friday departures in September, the month just preceding the rail strike, and bookings for departures for strike day November 07th, 2014. The strike was announced three days prior to the strike (as indicated by the dashed line). Note that ticket sales are not in log scale here.

During the strike, MFB experienced peaks in the number of bookings made by new customers. Panel A of Figure 4 shows the daily average tickets sold per route (bars). Strike days are indicated by the vertical dashed lines. Panel B shows the development of weekly averages of daily sales. We differentiate between three groups of tickets: all tickets, the subset of all tickets bought by first-time customers, and the subset of tickets bought by first-time customers that were bought no more than three days before departure (the spontaneous first-time customers). A ticket is marked as a first-time sale when the booking was made with an e-mail address not yet registered with MFB and a spontaneous sale when, in addition, the booking was made no more than three days prior to departure.[17] Figure 4 suggests that sales increases during the strike are driven by increases in ticket sales to first-time and spontaneous customers. For all three groups, there is a weekly pattern in ticket sales. Friday and Sunday departures sell the most tickets.

[17]On average 30% of bus passengers are first-time customers, two-thirds of whom undertake at least one more booking within our sample period.

Figure 4: Mean ticket sales per route

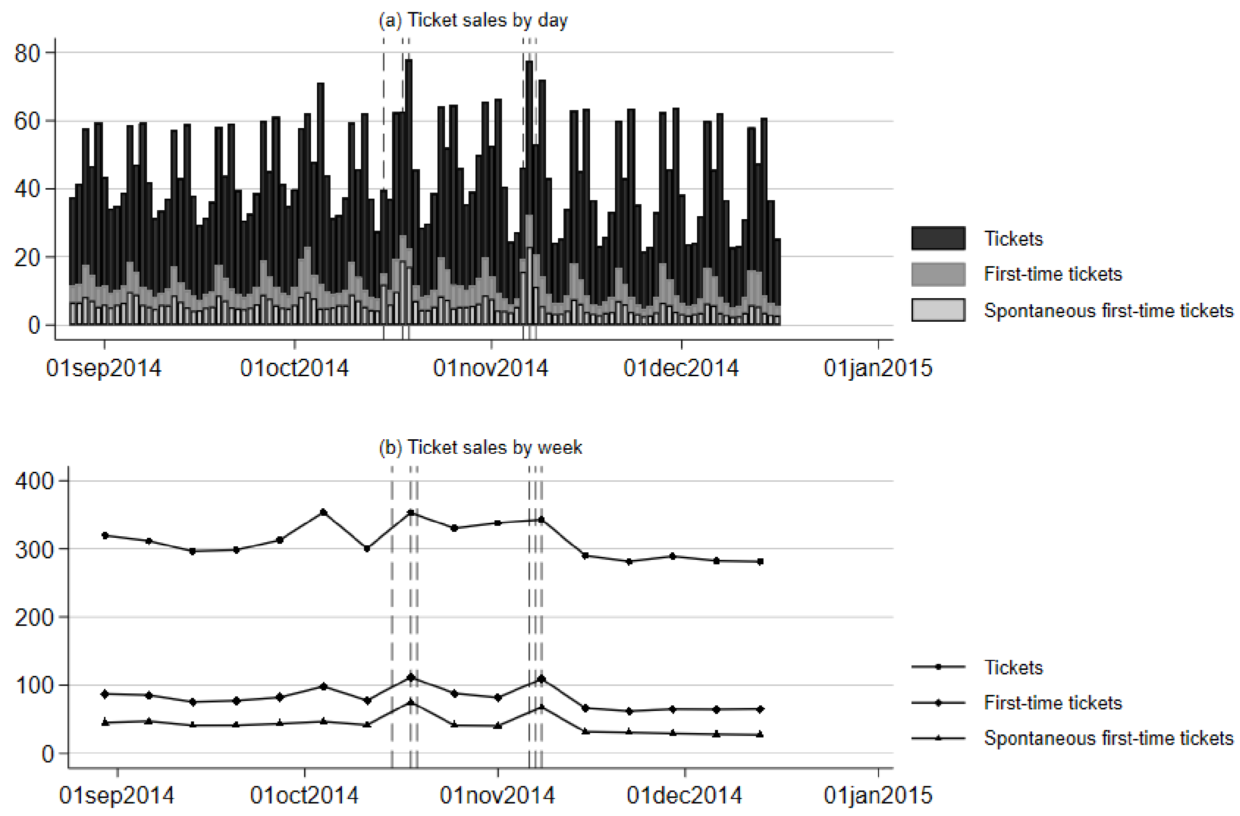

Note: The graph shows the average number of bus tickets sold per route. The bars (a) indicate daily averages, and the lines (b) show weekly averages. The vertical dashed lines indicate strike days.

DB Emergency timetables and itineraries

Rail service cancellations varied by route. DB maintained rail services on key routes, partly with reduced frequency. A route's exposure to cancellations is exogenous to the MFB service on this route. To capture the route-varying effect of the strike, we deduce to what extent a route was affected by the rail strike from DB's emergency timetables. The DB emergency timetables list DB services at the *line level*. For example, ICE line 25 from Hamburg to Munich halved its operations from once every hour to once every two hours. However, a typical itinerary involves stopovers and hence uses multiple rail lines. We combine the emergency timetables provided by DB with DB travel itineraries, which were collected using a web crawler linked to a leading price comparison website. Using actual itineraries takes into account that some DB routes are served through different paths in the rail network. We can then recapitulate all possible railway connections between any of the 33 cities in our sample, including the departure times, the number and length of stopovers, and the lines used. We deem a connection unfeasible if it uses more than five trains or requires a waiting period of more than 120 minutes. To the remaining connections in our sample, we assign a strike exposure, measuring the fraction of services cancelled during the strike along all connections that would regularly be available on this route.

One data limitation remains: the DB emergency timetables do not include information on regional trains. We disregard connections where more than 10% of itineraries include the use of regional trains. DB's Inter-City and Inter-City Express trains connect all 33 cities in our sample. 81% of all connections in our sample use at least one Inter-City Express train. Nonetheless, the lack of strike data for regional trains might seriously limit our findings if regional trains were strongly affected by service cancellations or if we drop primarily short routes on which the bus may be a closer substitute for the train (we address this in Section 6.3).

The average train trip in our sample takes three hours. The median length of a trip is 182 minutes in pure travel time. Trips longer than six hours are uncommon in our data. The high-speed-railway trains can technically reach a speed of up to 300km per hour, yet the railway infrastructure does not permit a speed above 200km per hour on most routes. Ticket prices depend on the distance travelled. They regularly would not exceed €139 (business class: €225) for a one-way trip in 2013. DB also offers tickets at a dynamic price (Sparpreis) that varies with the expected demand, where a single economy class trip begins at €19 for an economy class trip below 250km and at €29 above. Typically, a connection between two cities in our 33-city sample uses no more than two trains and requires no more than 14 minutes of waiting time at changeovers. Of those connections that do require a changeover, 95% do not exceed a total waiting time of 53 minutes on all stops.

We measure each route's exposure to the rail strike by the fraction of train connections cancelled on a given route during the strikes (*trains cancelled (%)*) for each day of the week for each route. Where there are several possible connections on one route, the strike exposure is a weighted average of the connections' exposure. *Trains cancelled (%)* captures how many of the possible departures on all available connections on a route were inoperative during the strike.

Figure 5 plots the variable *trains cancelled (%)* against the rail travel time under the regular schedule. There is no visible systematic relationship between normal rail travel time and the fraction of services cancelled during the strike. Only one route-day combination maintains full service on all connections under the emergency timetable: Berlin to Hannover on Wednesday. In this case, *trains cancelled (%)* equals zero. Since the fraction is aggregated over all connections of a route, made up of up to five trains per connection, most routes are affected by the strike at least partially. Note that we can only capture the cancellations according to the emergency timetable here. Additional delays and cancellations - as are common to occur with DB services also outside of the strike for all kinds of reasons - are just noise to our analysis. Our treatment variable measures strike exposure solely as the discrepancy between the regular and the strike schedule.

Figure 5: Train travel time without strike and fraction of trains cancelled during strike

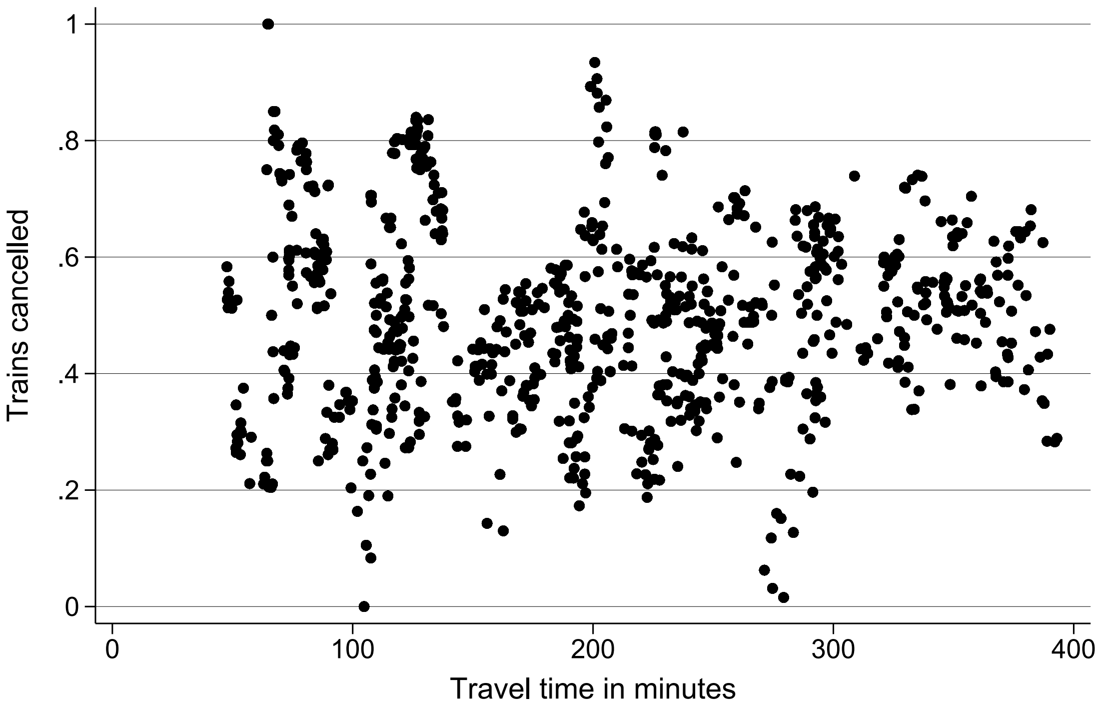

Note: Data from DB itinerary and emergency timetables.

5 Estimation strategy

We test for the effect of a route's exposure to the rail strikes on MFB ticket sales in an estimation based on the following baseline specification:[18]

$$\ln y_{ijt} = \sum_{s}^{Ns} \beta_s(exposure_{ij} \times strike^s_t) + \mu_{ijdow} + \mu_{iw} + \mu_{jw} + \boldsymbol{X}_{it}\boldsymbol{\gamma}^{'} + \boldsymbol{X}_{jt}\boldsymbol{\delta}^{'} + \epsilon_{ijt} \tag{4}$$

with $s \in \{1, 2, 3, post\}$, where $\ln y_{ijt}$ is the natural logarithm of the number of bus tickets sold on a route connecting origin-city i to destination-city j at time t (measured in days).[19] Treatment is defined by the interaction term $(exposure_{ij} \times strike^s_t)$. $Exposure_{ij}$ captures the extent to which a route was affected by the strike, measured by the variable *trains cancelled (%)*. $Strike^s_t$ is a dummy that equals one on a strike day and zero otherwise. As we want to know whether the different strike waves affected bus travel to different extents, we include three

[18] We implement all OLS regressions using *reghdfe* (Correia, 2016). PPML regressions are performed using the Stata command *ppmlhdfe* (Correia et al., 2020).

[19] ln(ticket sales + 1).

separate dummies capturing the first, second and third wave respectively.[20] A fourth dummy, $strike_t^{post}$ identifies the post-strike period (dates after November 11th, 2014).[21] The coefficient β_s measures the extent to which higher strike exposure of affected routes impacts the number of bus bookings relative to less affected routes.

μ_{ijdow} is an interaction of route and day-of-the-week fixed effects. They capture unobserved time-invariant route-specific characteristics (such as the fact that some routes are more popular than others) that may or may not correlate with the degree of exposure to the strike. They also capture route-specific differences that vary across different days of the week. For example, some routes might be more popular during weekdays, while others may be frequented more often during weekends. As shown in Figure 4, Fridays and Sundays are particularly popular travel days. μ_{ijdow} also captures such variation in ticket sales that is common over all routes.

μ_{iw} and μ_{jw} are origin-week and destination-week fixed effects, respectively. They capture variation in temporal factors common to all departures from or arrivals in each city, such as national holidays, MFB marketing campaigns, origin and destination-specific changes in ticket prices or seasonal fluctuations. Finally, the specification includes vectors of control variables for city-specific daily events, such as holidays or soccer games $\boldsymbol{X}_{it}$ and $\boldsymbol{X}_{jt}$.[22] ϵ_{ijt} is an error term. To address potential serial correlation within routes and over time, we cluster standard errors by route throughout the paper.

Identification relies on the assumption that the trend in log-linearised bus ticket sales on different routes does not vary systematically with the extent to which these routes are disrupted by strikes. Selection into strike exposure is not a threat, but must not be specific to a particular dose. We argue that the level of strike exposure is exogenous to the trend in bus ticket sales. DB did not strategically focus rail services on routes that were under particular threat of competition from buses. The emergency timetables were the same in all strike waves in 2014-2015, and they are almost identical to those employed by DB in the last railway strikes of 2007-2008; i.e. long before the liberalisation of the intercity bus market in 2013.

While we assume that routes follow a common trend, the levels of our dependent variable vary across routes and time. Some routes sell up to several hundred tickets daily, whereas others sell no more than two on some days. Furthermore, we know that ticket sales are responsive to major holidays. For example, Figure 4 shows that uncommonly many tickets were sold in the week of October 3rd. October 3rd is Germany's national holiday, which fell on a Friday in the

[20] Note that the treatment is not staggered, as all treatment groups are treated at the same time, albeit with different intensity.

[21] As illustrated in Figure 1, rail services remained disrupted on Saturday November 9th and Sunday November 10th. These two dates neither fall in the 3rd wave treatment group, nor are they captured by the post-strike dummy. They are thus part of the control group, so that $\hat{\beta}_{post}$ may underestimate the true treatment effect.

[22] These are school holidays, public holidays, the day before a long weekend, Football World Cup games, Bundesliga games, Oktoberfest, Stuttgarter Wasn, and Gamescon. Note that German holidays vary at the state level.

year 2014, creating a long weekend off for students and many employees. On the Thursday of this week, in particular, MFB sold more tickets than usual. Further, seasonality might affect our data: the vacation period ends and travel is less frequent in the winter months. Overall, a decreasing trend is present in the data. Before the strike, MFB sold on average 14,000 tickets per day. After the strike, average daily sales were at 12,500 tickets. Sales then drastically increased again for Christmas (outside of our observational period). We control for these specific events and aggregate trends using our battery of controls and fixed effects. Consequently, our key identifying assumption cannot be spoiled by such seasonality, as long as there is no difference in underlying trends across routes that coincides with our exposure measure. We put the common trend assumption to the test in an event study estimation, illustrated by Figure 6 in Section 6.2.

We identify treatment effects by relying on the variation in exposure to the strike, measured by the fraction of rail services cancelled. This strategy captures the impact of the strike itself. However, the post-strike effect cannot be exactly identified because not all individuals regularly travel on the same routes. If a route is affected by the strike and causes travellers to switch to the bus, this is captured by the respective beta coefficient. However, if the strike induces a person to permanently switch to bus travel after the strike, this is only captured if the individual uses the same route after the strike. To give an example, say that the route Hamburg-Munich was strongly affected during the strike, causing people to switch to the bus. If their experience convinces travellers to permanently switch to intercity buses, they will not only use them to travel from Hamburg to Munich (which is captured by β_{post}) but also to travel to other destinations such as Berlin. If Hamburg-Berlin was unaffected during the strike, this observation would fall in the control group, leading to an underestimation of the post-strike treatment effect. Our estimates for the post-strike effects should thus be seen as a lower bound of the true effect. In Section 6.3, we propose an alternative measure to better capture post-strike effects.

6 Results

6.1 Baseline results

Column 1 of Table 1 provides the results of our baseline estimation as specified in Equation 4. The railway strike significantly increases bus ticket sales on affected routes during all three strike waves. Specifically, the coefficient of 0.507 indicates that a one percentage point increase in the fraction of trains cancelled on a particular route during the first strike wave increases bus ticket sales on that route by 0.66%.[23] On average, about half of the possible train connections

[23] $e^{0.507} - 1 = 0.66$.

are cancelled on strike-affected routes. Our estimates suggest that moving from no cancellations to the average cancellation rate of 49% yields an increase of 32% in ticket sales - which would be around 14 additional tickets with respect to the mean of 45 tickets sold per route per day.

Table 1: The effect of rail strikes on bus ticket sales

	(1)	(2)	(3)	(4)
Dependent variable: ln(*tickets*)	all	first	spont.	all
Trains cancelled (%) × strike 1	0.507***	0.924***	1.026***	0.474***
	(0.0565)	(0.0856)	(0.0899)	(0.0564)
Trains cancelled (%) × strike 2	0.366***	0.906***	1.357***	0.464***
	(0.0480)	(0.0740)	(0.0783)	(0.0480)
Trains cancelled (%) × strike 3	0.367***	0.725***	1.357***	0.396***
	(0.0416)	(0.0587)	(0.0664)	(0.0418)
Trains cancelled (%) × post-strike	0.155*	0.110	0.143**	0.0360
	(0.0904)	(0.0898)	(0.0648)	(0.0884)
Trains cancelled (%) × post × weekend				0.382***
				(0.0276)
R^2	0.918	0.824	0.789	0.920

Note: OLS regressions with origin-week, destination-week and route-dow fixed effects as well as controls. Standard errors clustered by route in parentheses. ***/**/* indicate significance at the 1%/5%/10% level. 16,336 observations.

The coefficients of 0.366 and 0.367 for the second and third wave imply that bus ticket sales on affected routes increase by 0.44% following a one percentage point increase in the fraction of services cancelled. Perhaps most strikingly, the effect persists beyond the duration of the strike. The significantly positive post-strike coefficient of 0.155 indicates that routes which experience a one percentage point higher train cancellation rate during the strike see 0.17% higher ticket sales after the strike. This amounts to a permanent 8% increase in ticket sales on the average route. As discussed in Section 5, this estimate most likely constitutes a lower bound of the true treatment effect.

In light of the literature on habit formation and switching costs, our results provide evidence that rail customers have under-experimented before the strike. After having been forced to experiment with alternative transport modes during the strike, some passengers continue to use buses, even after rail services have resumed. Service disruptions of incumbent firms can thus have pro-competitive effects, as rivals can permanently lure away some of the incumbent's customers.[24]

[24]Travellers might have booked bus tickets after the November 2014 rail strike, because they were worried about potential future strikes. The rail strikes lasted beyond the strikes in 2014, and the labour dispute was only resolved after additional strike waves in April and May 2015. However, immediately after the strike wave

In order to investigate whether aggregate effects are driven by first-time customers, we estimate our baseline specification, using the logarithm of the number of tickets sold to first-time customers as the dependent variable. The effects of the strike are even more pronounced for first-time ticket sales, as Column (2) of Table 1 shows. A one percentage point increase in train cancellations during the first strike wave is associated with a 1.52% increase in ticket sales among first-time customers. Second and third-wave effects are also stronger for first-time customers. The estimated coefficient for the post-strike effect is not significantly different from zero. This is to be expected, as routes affected more strongly during the strike should not attract more first-time customers after the strike (only repeat customers).

Potentially the best proxy for people who first switch to the bus during the strike, is the logarithm of spontaneous first-time ticket sales. This entails all first-time bookings made no more than three days prior to departure.[25] Estimated coefficients are even larger for spontaneous sales (Column 3). A one percentage point increase in cancellation on a given route during the first wave is associated with a 1.79% increase in spontaneous ticket sales.

Perhaps surprisingly, we also find significant post-strike effects, although the effect is only one tenth of that during the strike. This indicates that MFB continues to gain more new customers on strike-affected routes. Customers might have had a negative experience with DB during the strike and make the switch only for their next trip or they might have experienced MFB already but not have made the booking themselves. The majority of bookings in our data are made for at least two travellers. In the booking process, only one person registers. If other members of the group decide to book a ticket in the post-strike period, they would appear as first-time customers in our data, despite having travelled with MFB before.

As shown in Figure 4, the absolute number of ticket sales varies across days of the week. Any such variation should be absorbed by day-of-the-week fixed effects in all of our specifications. However, it is possible that treatment effects vary across different days of the week. In particular, passengers travelling over the weekend are more likely to travel for leisure than for work. They might thus be more willing to endure longer travel times using buses in return for lower ticket prices. We would thus expect the post-strike impact to be stronger for weekend trips. If treatment effects vary across different days of the week, this could also explain the relatively large standard error of the estimated post-strike coefficient (about twice as large as the standard errors of the contemporaneous strike coefficients) observed in Column (1) of Table 1. We therefore re-run our baseline regression adding another regressor, namely the interaction

in November, GDL announced a temporary truce. It would refrain from industrial action until the new year. Even though some customers may have distrusted the truce, it is unlikely that increased bus ticket sales in this period are driven by the fear of new strikes.

[25] We observe ticket sales. We cannot observe whether the ticket was actually used or bought as insurance against potential rail cancellations.

of our post-strike treatment variable with a dummy indicating whether the bus departs on a weekend (i.e. on Friday, Saturday, or Sunday).

Regression results including this triple interaction are presented in Column (4) of Table 1. The post-strike interaction with the strike exposure measure, *trains cancelled (%)*, turns insignificant in this specification. However, the triple interaction term is highly significant and about twice as large as the post-strike coefficient in our baseline specification (Column 1). The results imply that the increase in ticket sales on strike-affected routes after the strike is driven by weekend sales. Bus travel thus seems a useful substitute for train travel mostly for leisure passengers, travelling on weekends (Friday to Sunday).

6.2 Effects over time

To further illustrate the impact of the strike over time, we regress ln ticket sales on our strike exposure variable, interacted with a dummy for each day in our sample period.[26] Estimated coefficients are depicted graphically in Figure 6. The event study serves two objectives. First, estimated coefficients are mostly insignificant before the first strike, indicating no difference in pre-treatment trends between treatment and control group. Second, the graph clearly shows that the post-strike effect is always significantly positive for Fridays, Saturdays and Sundays. It also shows significant positive effects for the two days following the third strike wave (Sunday, November 9th and Monday, November 10th). This is to be expected for two reasons: First, the strike was originally planned to last until Sunday but ended prematurely on Saturday following public pressure. Second, as indicated by Figure 1, service disruptions continued until Monday, November 10th as it took DB two days to return to its normal schedule.

6.3 Duration of trip effects

Travelling by bus between cities in Germany typically takes longer than travelling by train. The relative trip duration further increases with the length of the travelled route. For each additional stop on a bus journey, the bus has to leave the highway and enter the city centre to reach the bus terminal. Figure 7 illustrates the increasing divergence in travel time. The longer the trip, the slower bus travel is compared to the train. If customers dislike travel with an increasing margin, then a modal switch from rail to bus might be more attractive on short routes. Travelling by train offers the comfort of getting up to walk around or take a meal onboard, a comfort that might increase in value with the length of the trip. When the travel

[26]We estimate the following equation: $\ln y_{ijt} = \sum_d^D \beta_d(exposure_{ij} \times day_t^d) + \mu_{ijdow} + \mu_{iw} + \mu_{jw} + \boldsymbol{X}_{it}\boldsymbol{\gamma}' + \boldsymbol{X}_{jt}\boldsymbol{\delta}' + \epsilon_{ijt}$. day_t^d is a dummy identifying each day in the sample period. We omit the week before the beginning of the strike in order to interpret the coefficient relative to a baseline.

time exceeds several hours, bus travel might simply not be a good substitute for a train trip, and rail customers might have switched to travel by car or plane instead.

Figure 6: The effect of rail strikes on bus ticket sales, by day of departure

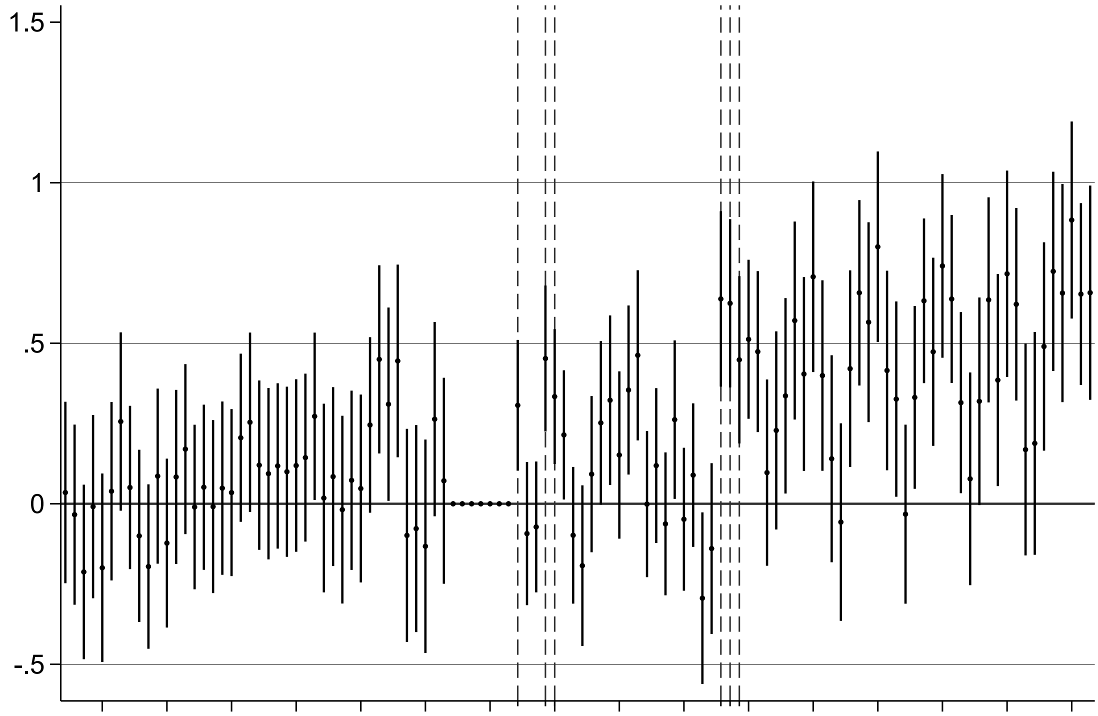

Note: Estimated coefficients on the vertical axis, days from August 27 to December 16 on the horizontal axis. The dashed vertical lines indicate strike days. The plots show point estimates with 95% confidence intervals, indicating significant effects during and after the strikes. Ticks on the x-axis indicate Sundays.

We hence introduce the duration of the bus trip as an alternative treatment variable to test the hypothesis that shorter bus routes experienced an increase in bus ticket sales during and after the strike. We define a bus ride to be bivariate relatively short if the scheduled time of travel is below the median of 265 minutes (a little over 4.5 hours). Note that shorter bus routes were not systematically more affected by the strike.[27]

Using relative trip duration as the treatment variable has another advantage compared to the fraction of services cancelled. Short routes remain relatively more attractive than long routes after the strike. As discussed in Section 5, passengers do not always travel along the same route. Customers who were affected by the strike by having travelled on a route with a high fraction of services cancelled and who decide to also travel by bus after the strike will probably do so on different (untreated) routes, too. Consequently, the estimated post-strike coefficient underestimates the treatment effect.

[27]The correlation between the duration of the bus ride and the fraction of trains cancelled on a route is 0.06.

Relatively short bus routes, however, offer a more attractive alternative to travelling by train both during and after the strike. Consequently, one would expect post-strike bus travel to increase more strongly on shorter routes. Even during the strike, decisions by passengers to switch transport modes may partly be driven by uncertainty regarding the reliability of the emergency timetable, so passengers switched to buses even on routes less strongly affected by the strike. As we show in the robustness section, there is indeed some evidence for a general increase in bus use during the strike, even on less affected routes. Faced with this uncertainty, passengers may switch from train to bus if they perceive the bus to be a decent substitute. Finally, not having to rely on information on service cancellations means that we can include routes that are only served by regional trains (recall that for regional trains we do not have information on service cancellations). This almost doubles the sample size.

Figure 7: Travel time bus vs. rail

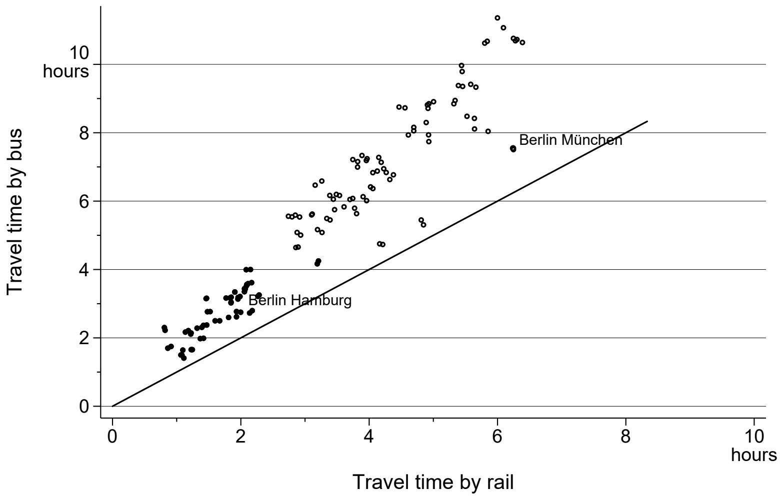

Note: Scatter of routes in duration rail and duration bus space with 45 degree line. Filled dots indicate below-median bus travel time. Routes Berlin–Munich and Berlin–Hamburg plotted as examples.

Table 2 presents the results. Column (1) shows that the number of tickets sold indeed increases more strongly on short routes. This is true during as well as after the strike. During the first wave, ticket sales increase by 37% on short routes. In the weeks after the strike, short bus routes sell up to 24% more tickets. Sales to first-time customers even increase by 69% during the first wave (Column 2).[28] In Column (4), the post-strike coefficient is now significantly positive, although effects remain stronger during the weekend. All in all, results are qualitatively similar to those presented in Table 1. Given that strike-effects are indeed stronger on shorter routes and that many short routes are excluded in the baseline regression

[28]The positive post-strike effect for first-time customers could once again be driven by group members who first travelled during the strike but did not make the booking. As in Column 3 of Table 1, the coefficient is much smaller in magnitude than the contemporaneous strike coefficients.

due to data limitations, the results imply that the baseline specification underestimates the true treatment effect.

Table 2: The effect of rail strikes on bus ticket sales, duration of trip

	(1)	(2)	(3)	(4)
Dependent variable: ln($tickets$)	all	first	spont.	all
Bus ride short × strike 1	0.312***	0.527***	0.585***	0.283***
	(0.0317)	(0.0470)	(0.0454)	(0.0316)
Bus ride short × strike 2	0.225***	0.542***	0.721***	0.290***
	(0.0238)	(0.0388)	(0.0410)	(0.0243)
Bus ride short × strike 3	0.319***	0.545***	0.722***	0.339***
	(0.0221)	(0.0322)	(0.0343)	(0.0223)
Bus ride short × post-strike	0.218***	0.150***	0.00842	0.106***
	(0.0169)	(0.0188)	(0.0173)	(0.0174)
Bus ride short × post × weekend				0.280***
				(0.0123)
R^2	0.906	0.783	0.739	0.908

Note: OLS regressions with origin-week, destination-week and route-dow fixed effects. Standard errors clustered by route in parentheses. ***/**/* indicate significance at the 1%/5%/10% level. 34,818 observations.

7 Extensions and Robustness

In the following, we explore the dynamics of ticket prices and capacity during the strike. If ticket prices at MFB were exceptionally low during the strike, this could indicate that MFB used the increased public attention during the railway strike to attract new customers with extremely competitive price offers. In this case, the estimated strike effect would be supply-side driven rather than the result of a shift in demand. We test this hypothesis by regressing the logarithm of average ticket prices at the route-day level $\ln(price_{ijt})$ on our strike exposure variables, as well as time dummies identifying the different strike waves. The results, provided by Column (1) of Table 3, indicate no such supply shock. Average ticket prices are significantly higher during the second and third strike wave. This is true for all routes, not just those affected by the strike (although prices are even higher on affected routes). We are hence confident that our results mirror an increase in demand.[29] As an additional robustness check, we re-run our baseline regression but include average prices as an additional control (Column (2) of Table 3). The price coefficient is significantly positive, while all other coefficients remain similar to

[29]Generally, MFB bus fares dynamically increase as capacity fills up.

the baseline in both magnitude and significance. The post-strike effect even increases in both magnitude and statistical significance.

Table 3: Ticket prices, bus capacity, and alternative treatments

	(1)	(2)	(3)	(4)	(5)
Dependent variable:	ln(*price*)	ln(*tickets*)	ln(*seats*)	ln(*tickets*)	ln(*tickets*)
Treatment variable:	Trains cancelled	Trains cancelled	Trains cancelled	Above median	Train delay
Treatment × strike 1	0.0448	0.489***	0.0285	0.229***	0.00125***
	(0.0508)	(0.0571)	(0.0854)	(0.0419)	(0.000170)
Treatment × strike 2	0.163***	0.314***	0.0630	0.183***	0.000944***
	(0.0509)	(0.0527)	(0.0642)	(0.0341)	(0.000141)
Treatment × strike 3	0.117**	0.280***	0.0403	0.172***	0.000981***
	(0.0516)	(0.0459)	(0.0793)	(0.0316)	(0.000138)
Treatment × post-strike	0.0194	0.169**	0.110*	0.0853**	-0.0000479
	(0.0235)	(0.0831)	(0.0596)	(0.0368)	(0.000143)
Strike 1	0.0362		0.0889**		
	(0.0267)		(0.0421)		
Strike 2	0.134***		0.00825		
	(0.0249)		(0.0299)		
Strike 2	0.220***		0.00890		
	(0.0268)		(0.0425)		
Post-strike	-0.0952***		-0.105***		
	(0.0135)		(0.0380)		
ln(*price*)		0.137***			
		(0.0431)			
Observations	16,282	16,282	16,282	16,336	18,592
R^2	0.969	0.922	0.917	0.917	0.912

Note: OLS regressions with origin-week, destination-week and route-dow fixed effects as well as controls. Standard errors in parentheses, clustered at the route level. ***/**/* indicate significance at the 1%/5%/10% level.

An increase in ticket sales might not directly reflect an increase in demand to an equal extent if sales were capped when MFB's offer reached short-term capacity peaks. MFB did increase the number of buses running on a route, thereby the number of seats available for booking, on departure days with high travel activity. In our data, we observe capacity increases parallel to sales increases during the strike, but also on weekends and national holidays. On average during our observation period, 0.3% of all bus connections were fully booked, meaning that there were no more tickets available for this specific connection at one specific departure time. However, occasionally, several - or even all - connections are booked out on a certain day on a certain route. During the strike, the total number of fully-booked connections increases. Yet, no route exceeds a share of 37% in fully booked connections during the strike. Within the same day on the same route, there were always options for departure at a different time. It seems that while

MFB was used to adjusting capacities to demand, the strike did not present an exceptional challenge in this regard.

We re-estimate our baseline specification with the logarithm of the number of available seats per day per route as the dependent variable. As can be seen from Column (3) of Table 3, capacity increases during the strike were mostly insignificant and did not specifically affect strike-exposed routes. After the strike, at a time when overall capacities were in decline along with the seasonal trend, we observe a significant increase in capacity on strike-affected routes. This falls in line with expectations given the increase in post-strike ticket sales on these routes.

In Columns (4) and (5) of Table 3, we explore alternative measures for a route's exposure to the strike. The effect of service cancellations on passenger numbers may be non-linear. For example, passengers might not mind if a small fraction of rail services is cancelled as long as alternative connections are offered. However, if the fraction of services cancelled exceeds a certain level, passengers might switch to alternative modes of transport. Instead of using the fraction of services cancelled directly as a regressor, we thus construct a dummy that equals one if the fractions of services cancelled on a particular route is above the median (and zero otherwise). This variable is then interacted with the four strike dummies. The results, reported in Column (4) of Table 3, are qualitatively similar to the baseline, indicating increasing passenger numbers on affected routes both during and after the strike. Specifically, routes with above median exposure to the strike experience a 9% increase in ticket sales in its aftermath.

Next, we create an alternative treatment variable, *train travel delay*, that captures the additional waiting time travellers would on average have to incur to take the next train if their service is cancelled. Unlike *trains cancelled (%)*, this measure takes into account the frequency of connections on a given route. The routes Berlin–Munich and Hamburg–Berlin, for example, both experienced service cancellations (*trains cancelled (%)*) of about 75%. Yet, a customer on average had to wait for 450 minutes for the next train during the strike on the route Berlin–Munich, whereas for Hamburg–Berlin another departure was on average available within 114 minutes. This is because Hamburg–Berlin operated at a much higher frequency even in times of the strike. The average waiting time - the minutes until the next uncancelled departure - constitutes our treatment variable *train delay*. Customers who knew about the cancellation might have chosen to arrive at the train station later and take the next departure according to the emergency timetable. Hence, *train delay* is not so much a measure of actual waiting time, but rather the loss of flexibility during the strike.

Column (5) of Table 3 shows the regression results when exposure is measured as *train delay*. The independent variable runs from 0 minutes to 557 minutes of additional waiting time for potential travellers caused by rail service disruptions. An additional minute of waiting time yields 0.12% more tickets sold during the first strike wave. According to the regular schedule, the average waiting time across all our routes is 45 minutes. During the strike, the average wait increases to 116 minutes. We estimate that this average 71-minute *train delay* corresponds to 8.52% more tickets sold for MFB. Unlike in the case of trains cancelled, the additional waiting time did not induce a long-term increase in ticket sales beyond the strike. An explanation might be that the ticket sales increase in response to *train delay* is driven by customers with a preference for flexibility. They will be keen to switch to the bus if it is the next available travel option but will revert to the train as soon as it is frequently available again.

In our baseline estimation, we exploit variation across routes as well as across strike and non-strike days to estimate a treatment effect. Origin-week and destination-week effects control for overall time trends. A conventional difference-in-differences set-up would require time dummies that capture exactly the treatment period. As a robustness check, we extend the regressions shown in Table 1 with four time-variant dummies that identify the three strike periods as well as the post-strike period. The results are provided in Table A.1 in the appendix. Estimated coefficients of the three strike dummies are significantly positive in Columns (2) and (3), indicating that more first-time as well as spontaneous users travel by bus during the strike. This is not surprising, as many people were probably either not aware of the emergency timetables or did not trust them, leading them to travel by bus on unaffected routes, too. The post-strike dummy is significantly negative in almost all specifications. This is in line with our descriptive evidence that the overall number of bookings declined over the sample period. Our treatment variables remain qualitatively similar to the baseline estimates, indicating an increase in the number of passengers travelling by bus on affected routes both during and after the strike.

We perform the same exercise using our distance variable (analogous to Table 2). The results, reported in Table A.2 in the appendix, point towards the same direction. They show an overall increase in the number of bus passengers during the three strike waves and a decline in the post-strike period. Shorter bus routes experience a stronger increase in passenger numbers during the strike, which also persists in its aftermath.

An even more conservative approach would be the additional use of day fixed effects. We choose not to do this in our baseline specification as they might absorb too much variation needed to identify treatment effects. As discussed above, estimating differential effects for weekends in the post-treatment period becomes problematic because the exposure variable does not explicitly vary in the post-treatment period any more. However, we do run our main regressions including day fixed effects as a robustness check. The results are reported in Table 4. Estimated coefficients for our preferred strike and post-strike variables (fraction of services cancelled) remain qualitatively similar (Column 1). The post-strike effect even increases in magnitude and significance, probably because day fixed effects better capture the overall decline in ticket sales in the post-strike period than origin- and destination-week fixed effects employed in the baseline. The weekend coefficient indeed becomes statistically insignificant, although the post-strike coefficients now turn significantly positive (Column 2). When using short distance as the treatment variable, the significance of the estimated coefficients remains extremely robust, although they become smaller in magnitude (Columns 3 and 4).

In an alternative specification, we estimate Equation 4 with PPML instead of OLS. The results, reported in Columns (5) and (6) of Table 4, are qualitatively similar to the baseline results.

Table 4: Day fixed effects and ppml

	(1)	(2)	(3)	(4)	(5)	(6)
Dependent variable: ln(*tickets*)	OLS	OLS	OLS	OLS	PPML	PPML
Treatment variable:	Trains cancelled	Trains cancelled	Bus ride short	Bus ride short	Trains cancelled	Trains cancelled
Treatment × strike 1	0.243	0.236	0.151***	0.134***	0.393***	0.358***
	(0.156)	(0.155)	(0.0426)	(0.0424)	(0.0513)	(0.0502)
Treatment × strike 2	0.374***	0.388***	0.117***	0.147***	0.247***	0.315***
	(0.141)	(0.142)	(0.0315)	(0.0315)	(0.0401)	(0.0437)
Treatment × strike 3	0.261*	0.266**	0.237***	0.249***	0.245***	0.258***
	(0.133)	(0.134)	(0.0302)	(0.0304)	(0.0296)	(0.0291)
Treatment × post-strike	0.327**	0.304**	0.235***	0.176***	0.203***	0.0443
	(0.132)	(0.131)	(0.0171)	(0.0184)	(0.0661)	(0.0599)
Treatment × post × weekend		0.0617		0.144***		0.348***
		(0.0951)		(0.0194)		(0.0291)
Day fixed effects	YES	YES	YES	YES	NO	NO
Observations	16,336	16,336	34,818	34,818	16,336	16,336
R^2	0.923	0.923	0.913	0.913		

Note: All regressions include origin-week, destination-week and route-dow fixed effects as well as controls. Standard errors clustered by route in parentheses. ***/**/* indicate significance at the 1%/5%/10% level.

8 Conclusion

This paper exploits a novel and rich dataset to investigate the effects of the 2014 German railway strikes – the largest in German history – on the use of intercity buses. The railway strikes provide a quasi-natural experimental setting to analyse the general question of whether a temporary service disruption can have lasting effects on demand for competing products.

We find a route-specific effect of railway cancellations leading to more bus travel. More specifically, the number of bus tickets sold increases during the strike on routes that are more strongly affected (on average by 32% during the first strike wave). Most strikingly, we show that the effect persists beyond the duration of the strike, with ticket sales on affected routes being on average 8% higher. Passengers continue using the bus even when trains become regularly available again. The effect is driven by weekend travellers, implying that more price-sensitive leisure travellers continue to use the bus after having been introduced to this new mode of transport during the railway strikes. Since passengers who permanently switch to buses most likely do so not only on routes that were strongly affected by the strike, our estimates for the post-strike effect constitute a lower bound of the true treatment effect. Even during strike times, our results indicate that some customers divert from DB services regardless of whether their planned travel is covered by the DB emergency timetables. In this case, our estimates, which rely on a route-based identification, present the lower bound of contemporary strike effects.

Shorter bus routes also see an increase in passenger numbers. Since both absolute and relative travel time differences between bus and train increase with distance, we interpret this as evidence that switching is more prominent on routes where buses are closer substitutes for train travel. The advantage of this measure is that it remains the same even after the strike. The post-strike effects, which are highly statistically significant, therefore do not suffer the downward bias that might plague our baseline estimation. They provide further evidence for a persistent switch in transport mode. Positive price effects suggest that the increase in ticket sales is the result of an increase in demand due to the disruption of rail services, not driven by a positive supply shock induced by MFB in response to the strike.

Our results demonstrate that the strike at the German railway company DB had a positive and lasting effect on the number of passengers travelling with DB's rival MFB. Service disruptions may thus cause passengers to persistently switch transport modes (in our case, from rail to intercity buses). Mobility transition away from fossil-fuel-powered individual transport is among the central challenges in combating the global climate crisis. From a policy perspective, our findings imply that carbon-neutral forms of transport need to be reliable in order to persistently attract passengers. In addition, as habit and search costs prevent people from switching transport modes, governments may need to find ways to lower the barriers to experimenting with sustainable modes of transport.

Beyond transportation, our findings provide evidence that persistent service reliability is an essential part of maintaining a customer base. Even relatively short outages can make customers experiment with rivalling products, overcoming search costs and other frictions that may previously have stopped them from making optimal choices. Temporary disruptions can thus have lasting impacts on both the affected company and its rivals. This result is applicable to many industries, from hairdressers to streaming providers. While interruptions may have surprisingly positive welfare effects for consumers by encouraging them to experiment with alternatives, companies interested in keeping their customers should do everything they can to avoid them.

References

Adda, J. (2016). Economic activity and the spread of viral diseases: Evidence from high frequency data. *The Quarterly Journal of Economics*, 131(2):891–941.

Adler, M. W., Liberini, F., Russo, A., and Ommeren, J. N. (2021). The congestion relief benefit of public transit: evidence from rome. *Journal of Economic Geography*, 21:397–431.

Adler, M. W. and van Ommeren, J. N. (2016). Does public transit reduce car travel externalities? quasi-natural experiments' evidence from transit strikes. *Journal of Urban Economics*, 92:106–119.

Allen, T. and Arkolakis, C. (2022). The welfare effects of transportation infrastructure improvements. *The Review of Economic Studies*, 89(6):2911–2957.

Anderson, M. L. (2014). Subways, strikes, and slowdowns: The impacts of public transit on traffic congestion. *American Economic Review*, 104:2763–2796.

Bataille, M. and Steinmetz, A. (2013). Intermodal competition on some routes in transportation networks: The case of inter urban buses and railways. Technical Report 84, Düsseldorf Institute for Competition Economics (DICE).

Bauernschuster, S., Hener, T., and Rainer, H. (2017). When labor disputes bring cities to a standstill: The impact of public transit strikes on traffic, accidents, air pollution, and health. *American Economic Journal: Economic Policy*, 9:1–37.

Baumol, W. J. and Quandt, R. E. (1964). Rules of thumb and optimally imperfect decisions. *American Economic Review*, 54(2):23–46.

Beestermöller, M. (2017). Striking evidence? demand persistence for inter-city buses from german railway strikes. *Munich Discussion Paper No. 2017-2.*

Ben-Akiva, M. and Morikawa, T. (1990). Estimation of switching models from revealed preferences and stated intentions. *Transportation Research Part A: General*, 24(6):485–495.

Bergemann, D. and Välimäki, J. (2006). Dynamic pricing of new experience goods. *Journal of Political Economy*, 114(4):713–743.

Böckers, V., Haucap, J., Heimeshoff, U., and Thorwarth, S. (2015). Auswirkungen der fernbusliberalisierung auf den schienenpersonenfernverkehr. *List Forum für Wirtschafts-und Finanzpolitik*, 41(1):75–90.

Chen, Y. and Whalley, A. (2012). Green infrastructure: The effects of urban rail transit on air quality. *American Economic Journal: Economic Policy*, 4(1):58–97.

Correia, S. (2016). Linear models with high-dimensional fixed effects: An efficient and feasible estimator. Technical report. Working Paper.

Correia, S., Guimarães, P., and Zylkin, T. (2020). Fast poisson estimation with high-dimensional fixed effects. *The Stata Journal*, 20(1):95–115.

Donna, J. D. (2021). Measuring long-run gasoline price elasticities in urban travel demand. *The RAND Journal of Economics*, 52(4):945–994.

Durr, N., Heim, S., and Hüschelrath, K. (2015). Deregulation, competition, and consolidation: The case of the german interurban bus industry. *ZEW Discussion Papers*, (15-061).

Evangelinos, C., Mittag, M., and Obermeyer, A. (2015). Die ökonomischen risiken einer zu naiven marktliberalisierung–der fall des deutschen fernbusmarktes. *Zeitschrift für Verkehrswissenschaft*, 86(1):65–90.

Friebel, G. and Niffka, M. (2009). The functioning of inter-modal competition in the transportation market: Evidence from the entry of low-cost airlines in germany. *Review of Network Economics*, 8(2).

Fung, C. M., McArthur, D. P., and Hong, J. (2021). Examining the effects of a temporary subway closure on cycling in glasgow using bike-sharing data. *Travel behaviour and society*, 25:62–77.

Gagnepain, P., Ivaldi, M., and Muller-Vibes, C. (2011). The industrial organization of competition in local bus services. In *A handbook of transport economics*. Edward Elgar Publishing.

Giulietti, M., Price, C. W., and Waterson, M. (2005). Consumer choice and competition policy: a study of uk energy markets. *The Economic Journal*, 115(506):949–968.

Goodwin, P. B. (1977). Habit and hysteresis in mode choice. *Urban studies*, 14(1):95–98.

Guihéry, L., Environnement, M. R., and Pontoise, F. (2016). *New interurban coach services in Germany and France: which European perspective? The key role of bus station... but where to locate them?*, chapter 14. University of Bamberg Press.

Janssen, A. (2023). Generic and branded pharmaceutical pricing: Competition under switching costs. *The Economic Journal.*

Kennan, J. (1986). The economics of strikes. *Handbook of labor economics*, 2:1091–1137.

Klemperer, P. (1987a). The competitiveness of markets with switching costs. *The RAND Journal of Economics*, pages 138–150.

Klemperer, P. (1987b). Entry deterrence in markets with consumer switching costs. *The Economic Journal*, 97(Supplement):99–117.

Klemperer, P. (1987c). Markets with consumer switching costs. *The quarterly journal of economics*, 102(2):375–394.

Klemperer, P. (1995). Competition when consumers have switching costs: An overview with applications to industrial organization, macroeconomics, and international trade. *The review of economic studies*, 62(4):515–539.

Larcom, S., Rauch, F., and Willems, T. (2017). The benefits of forced experimentation: Striking evidence from the london underground network. *The Quarterly Journal of Economics*, 132:2019–2055.

McDonald, J. T. and Bloch, H. (1999). The spillover effects of industrial action on firm profitability. *Review of Industrial Organization*, 15(2):183–200.

Morgan, P. and Manning, R. (1985). Optimal search. *Econometrica: Journal of the Econometric Society*, pages 923–944.

Moser, C., Blumer, Y., and Hille, S. L. (2018). E-bike trials' potential to promote sustained changes in car owners mobility habits. *Environmental Research Letters*, 13(4).

Porter, M. (1996). America's green strategy. *Business and the environment: a reader*, 33.

Riordan, M. H. (1986). Monopolistic competition with experience goods. *The Quarterly Journal of Economics*, 101(2):265–279.

Shapiro, C. (1983). Optimal pricing of experience goods. *The Bell Journal of Economics*, pages 497–507.

Simon, H. A. (1955). A behavioral model of rational choice. *The quarterly journal of economics*, pages 99–118.

Small, K. A. and Verhoef, E. T. (2007). *The economics of urban transportation*. Routledge.

Train, K. (1978). A validation test of a disaggregate mode choice model. *Transportation Research*, 12:167–174.

van Exel, N. J. A. and Rietveld, P. (2001). Public transport strikes and traveller behaviour. *Transport Policy*, 8:237–246.

Villas-Boas, J. M. (2006). Dynamic competition with experience goods. *Journal of Economics & Management Strategy*, 15(1):37–66.

Von Weizsäcker, C. C. (1984). The costs of substitution. *Econometrica: Journal of the Econometric Society*, pages 1085–1116.

Weeds, H. (2016). Tv wars: Exclusive content and platform competition in pay tv. *The Economic Journal*, 126(594):1600–1633.

Weitzman, M. L. (1979). Optimal search for the best alternative. *Econometrica: Journal of the Econometric Society*, pages 641–654.

Yang, Y., Beecham, R., Heppenstall, A., Turner, A., and Comber, A. (2022). Understanding the impacts of public transit disruptions on bikeshare schemes and cycling behaviours using spatiotemporal and graph-based analysis: A case study of four london tube strikes. *Journal of Transport Geography*, 98:103255.

Yeung, T. Y.-C. and Zhu, D. (2022). Intercity ridesharing to the rescue: Capacity flexibility and price stability of blablacar during the 2018 french railway strike. *Transportation Research Part A: Policy and Practice*, 164:270–290.

Zephirin, M. G. (1994). Switching costs in the deposit market. *The Economic Journal*, 104(423):455–461.

Appendix

Table A.1: Difference-in-differences: Trains cancelled

	(1)	(2)	(3)	(4)
Dependent variable: ln(*tickets*)	all	first	spont.	all
Strike 1	0.144*	0.322***	0.356***	0.150*
	(0.0808)	(0.114)	(0.123)	(0.0809)
Strike 2	0.00329	0.214*	0.427***	0.00815
	(0.0744)	(0.123)	(0.131)	(0.0746)
Strike 3	0.0223	0.187**	0.672***	0.0257
	(0.0735)	(0.0916)	(0.0922)	(0.0735)
Post-strike period	-0.288***	-0.285***	-0.115	-0.228***
	(0.0774)	(0.0896)	(0.0715)	(0.0752)
Trains cancelled (%) × strike 1	0.250	0.346	0.384	0.206
	(0.155)	(0.217)	(0.235)	(0.155)
Trains cancelled (%) × strike 2	0.375***	0.544**	0.611***	0.460***
	(0.140)	(0.216)	(0.232)	(0.141)
Trains cancelled (%) × strike 3	0.257*	0.320*	0.134	0.293**
	(0.131)	(0.166)	(0.163)	(0.132)
Trains cancelled (%) × post-strike	0.370***	0.312**	0.205**	0.209
	(0.133)	(0.136)	(0.0958)	(0.130)
Trains cancelled (%) × post × weekend				0.371***
				(0.0272)
R^2	0.919	0.825	0.790	0.920

Note: OLS regressions with origin-week, destination-week and route-dow fixed effects as well as controls. Standard errors clustered by route in parentheses. ***/**/* indicate significance at the 1%/5%/10% level. 16,336 observations.

Table A.2: Difference-in-differences: Duration of trip

	(1)	(2)	(3)	(4)
Dep. variable: ln(*tickets*)	all	first	spont.	all
Strike 1	0.199***	0.380***	0.432***	0.204***
	(0.0301)	(0.0466)	(0.0514)	(0.0301)
Strike 2	0.149***	0.397***	0.636***	0.157***
	(0.0237)	(0.0369)	(0.0439)	(0.0236)
Strike 3	0.0340	0.254***	0.643***	0.0405*
	(0.0221)	(0.0293)	(0.0313)	(0.0221)
Post-strike period	-0.293***	-0.256***	-0.0954**	-0.254***
	(0.0286)	(0.0399)	(0.0382)	(0.0275)
Bus ride short × strike 1	0.153***	0.226***	0.255***	0.120***
	(0.0425)	(0.0636)	(0.0664)	(0.0424)
Bus ride short × strike 2	0.117***	0.230***	0.200***	0.173***
	(0.0314)	(0.0505)	(0.0574)	(0.0317)
Bus ride short × strike 3	0.238***	0.292***	0.185***	0.259***
	(0.0301)	(0.0408)	(0.0425)	(0.0303)
Bus ride short × post-strike	0.236***	0.157***	-0.00520	0.123***
	(0.0171)	(0.0191)	(0.0177)	(0.0176)
Bus ride short × post × weekend				0.272***
				(0.0123)
R^2	0.907	0.785	0.746	0.909

Note: OLS regressions with origin-week, destination-week and route-dow fixed effects as well as controls. Standard errors clustered by route in parentheses. ***/**/* indicate significance at the 1%/5%/10% level. 34.818 observations.

Table A.3: Dates and duration of railway strike waves in 2014-2015

Nr.	Strike Begin:	—	Strike End:	Duration (in hours):
1	Mon. 01/09/2014, 18:00	—	Mon. 01/09/2014, 21:00	3*
2	Sat. 06/09/2014, 06:00	—	Sat. 06/09/2014, 09:00	3*
3	Tue. 07.10.2014, 21:00	—	Wed. 08.10.2014, 06:00	9*
4	**Wed. 15/10/2014, 14:00**	**—**	**Thu. 16/10/2014, 04:00**	**14**
5	**Sat. 18/10/2014, 02:00**	**—**	**Mon. 20/10/2014, 04:00**	**50**
6	**Thu. 06/11/2014, 02:00**	**—**	**Sat. 08/11/2014, 18:00**	**64**
7	Wed. 22/04/2015, 02:00	—	Thu. 23/07/2015, 21:00	43
8	Tue. 05/05/2015, 02:00	—	Sun. 10/05/2015, 09:00	127
9	Wed. 20./05/2015, 02:00	—	Thu. 21./05/2015, 19:00	41

Note: Bold rows indicate waves studied in this paper. Strikes in 2015 are disregarded, because they coincide with the merger of MFB and rival competitor Flixbus in January 2015. * indicates warning strikes. Warning strikes are ignored because they only lasted a few hours and were announced with many days' advance warning.