

The Benefits of Forced Experimentation: Striking Evidence from the London Underground Network*

Shaun Larcom[†]

Ferdinand Rauch[‡]

Tim Willems[§]

May 10, 2016

Abstract

We estimate that a significant fraction of commuters on the London underground do not travel on their optimal route. We show that a strike on the underground, which forced many commuters to experiment with new routes, brought lasting change in behavior. This effect is stronger for commuters who live in areas where the underground map is more distorted, thereby pointing towards the importance of informational imperfections. The information produced by the strike has improved network-efficiency. Search costs are unlikely to explain the suboptimal behavior. As commuters seem to under-experiment, constraints imposed on them can be welfare-improving.

JEL-classification: D83, L91, R41

Key words: Experimentation, Learning, Habit, Optimization, Search

1 Introduction

Do agents make first-best choices? And do their choices converge to the optimum if the underlying problem repeats itself over time? These questions address fundamental assumptions of economics across all fields, and have been subject to intense debate. We

*Many thanks to Transport for London (especially Dale Campbell and Maeve Clements) for providing us with the necessary data. We also thank Maarten Bosker, David Cox, Marco Grazzi, Vernon Henderson, Guy Michaels, Peter Neary, Henry Overman, Jörn-Steffen Pischke, Diego Puga, Collin Raymond, Kevin Roberts, Daniel Sturm, Alastair Young, seminar participants at Bern, Bologna, the UK Department for Transport, the London School of Economics, Massachusetts Institute of Technology (Sloan business school), Munich, Oxford, the Royal Economic Society Conference, St Andrews, and Toronto for useful comments and discussions. Francesca di Nuzzo provided excellent research assistance.

[†]Department of Land Economy, University of Cambridge. E-mail: stl25@cam.ac.uk.

[‡]Department of Economics, University of Oxford. E-mail: ferdinand.rauch@economics.ox.ac.uk.

[§]International Monetary Fund. E-mail: twillems@imf.org.

present evidence on these questions by analyzing a unique dataset from the London underground network (also referred to as “the tube”), which enables us to track individual commuter behavior over time. Our results provide evidence that individuals are “satisficing” (as hypothesized by Simon (1955)) and that they under-experiment in normal times.

Our dataset contains detailed information on commuters using Transport for London (“TfL”) services for four weeks. Most commuters use the electronic “Oyster Card” ticketing system to pay for their journeys. Each Oyster Card is associated with a unique identifier, which allows us to track the movement of individual commuters over time. In the middle of our sample period, the London underground transportation system was subjected to a strike. This caused major disruptions to the network on February 5 and 6, 2014. On these days, some (but not all) tube stations were closed. As a result, certain commuters were forced to experiment and explore new routes during this period, while others could go to work as usual. We analyze whether such an episode of forced experimentation produces any observable effects that last beyond the duration of the strike. When all stations were open again on February 7, did people who could not use their usual commuting route switch back to their original commuting paths? Or did some of them stick to alternative routes that they may have found during the disruption? By revealed preference, the latter possibility would suggest that they prefer the newly discovered alternative to their old habit, which would indicate that these commuters failed to find their best alternative before the strike.

Commuting itself is an important predictor of life satisfaction. Stutzer and Frey (2008) for example calculate that reducing total commuting time by 44 minutes a day, is worth about 35% of average net monthly labor income in terms of well-being.¹ But despite the sizeable stakes, we encountered various examples in the media of commuters who changed their route following the strike.² This indicates that a sizable share of people might be stuck in sub-optimal habits, even when it comes to activities that are repeated frequently.

¹Ahlfeldt *et al.* (2015) suggests a cost of commuting that is broadly consistent with the Stutzer and Frey (2008) estimate: their model suggests that a 10 minute commute reduces utility by 14 percentage points, whereas the utility-based number in Stutzer and Frey (2008) is 10 percentage points.

²See *e.g.* <http://www.bbc.co.uk/news/uk-england-london-26037534>, where it is noted that “some commuters have discovered more enjoyable ways of getting to work.” That article for example cites a commuter named Andy, who has experimented by taking the Thames Clipper water bus. He commented: “It has been fine, the boat is a great journey. I think I will get the boat back too.” Another commuter is quoted as saying that “the walk from Liverpool Street was a refreshing change from the horrors of the Circle Line. I suspect I may permanently switch so I can cut out this, the most stressful part of my journey.” After reading a first draft of our paper, Mike Spagat shared the story of his father with us: less than a year before his retirement, renovation works made him discover a new way to his job in Chicago which was superior to his old route that he had been taking for over 40 years.

This finding has more general importance.

The results of this study provide evidence on the ability of individuals to find optimal paths in networks, as well as on their approach to problem-solving. The latter issue has been subject to intense debate over many years. While the rational approach to decision-making has a long history in economics (see the contributions by Rothschild (1974a), Weitzman (1979), Roberts and Weitzman (1981), and Morgan and Manning (1985) to the literature on optimal search), others have remained skeptical of this characterization. Simon (1955) for example argued that agents are “satisficing” rather than maximizing – meaning that they stop their search-for-the-optimum once they have reached a satisfactory utility-level and apply rules-of-thumb from that point onward. It should be noted that although “satisficing” behavior could imply irrationality, this is not necessarily so. Subsequent work by Baumol and Quandt (1964) argues that such behavior may very well be rational when there are costs associated with decision-making – thereby anticipating the aforementioned “costly search literature” pioneered by Rothschild (1974a) and Weitzman (1979).³ Baumol and Quandt (1964) distinguish between “optimal” and “maximal” solutions: the latter refers to the exact solution, which would be obtained if there were no search costs, while the former takes such costs into account. Since its inception, Simon’s “satisficing-hypothesis” has proved to be difficult to test empirically. The first study to do so, was Caplin, Dean and Martin (2011): they presented experimental evidence suggesting that Simon’s “satisficing-approach” offers a good characterization of agents in a laboratory-environment. We take their study one step further by exploiting the unique setup offered by the February 2014 London tube strike, which enables us to provide the first field data-based evidence on this matter. At the same time, we are also able to analyze whether agents optimize in the sense of Baumol-Quandt, extending the analysis by Caplin, Dean and Martin (2011).

Our results also provide evidence on the inclination of individuals to experiment. After all, the new commute was already available pre-strike and could have been found beforehand through voluntary, as opposed to forced, experimentation. Many theoretical papers have pointed out that a certain degree of experimentation is optimal in settings where information is imperfect,⁴ but to the best of our knowledge there is no field data-based empirical work analyzing the incidence, as well as the effects, of experimentation in prac-

³More recently, Sims’ (2003) theory of rational inattention formalizes a similar idea: in his setup, decision makers have to allocate their scarce attention over multiple sources of uncertainty, which leads to deviations from standard maximizing behavior. Also see Matejka and McKay (2015) for an extension of the theory of rational inattention to a discrete-choice setup that characterizes our setting (should I take route A or route B?).

⁴See *e.g.* Rothschild (1974b), Aghion *et al.* (1991) and Bolton and Harris (1999).

tice.⁵ This paper is able to contribute along this dimension: we know exactly when many commuters were experimenting (namely during the strike), while the tube-environment provides us with a setting where information is very imperfect (thereby making a certain degree of experimentation optimal). The distorted nature of the schematic London tube map, which many travelers use to navigate, makes it difficult for travelers to minimize journey time (Guo, 2011).⁶ The fact that many line-characteristics (such as the line’s crowdedness and the follow-up journey to the final destination) are initially unknown, adds to the opaqueness of the situation.

Thanks to the presence of informational imperfections, our study is also able to add to the debate on the so-called “Porter-hypothesis”. Porter (1991) argued that – when information is imperfect – exogenously-imposed constraints may help agents to get closer to their optimum by triggering a period of experimentation, innovation and re-optimization. Porter originally phrased his hypothesis in the context of environmental regulation, but the underlying idea is more general and also applies to the tube-setup considered in this paper.⁷ Problematically to some, the Porter-hypothesis imposes a great deal of irrationality on the part of decision-makers: it implies that \$10 bills are waiting to be picked up from the pavement. After all, why would it take an exogenously-imposed constraint to make agents realize that they were not optimizing beforehand? Why wouldn’t they experiment voluntarily? As a result, Porter’s hypothesis has been dismissed by many scholars as being unrealistic – initially mostly on anecdotal grounds (see *e.g.* Palmer, Oates and Portney (1995) and Schmalensee (1993)). Subsequently, many studies have tried to test the theory empirically but, as noted by Porter and Van der Linde (1995) and Ambec *et al.* (2013), data limitations make it hard to put Porter’s hypothesis to a proper test

⁵Gabaix *et al.* (2006) conducted a laboratory study in which they found that a myopic strategy is a better characterization of actual search behavior than the rational approach set out in Weitzman (1979).

⁶Alan Turing once beautifully characterized the informational imperfection by describing a friend as “[thinking] of Paris like [...] I would think of a Riemann surface; he only knew the circles of convergence round every Metro station, and couldn’t analytically continue from one to another” (Hodges, 2014: 610). Similarly, in *The Guardian* of April 27 2015, it is written that: “When you first move to London it’s very common to quickly gain very detailed, even intimate knowledge of two or three locales, but not know how they are connected geographically. It’s not until there’s a Tube strike and you have to cycle or take the bus [...] that you suddenly realize that places you thought were separated by several sets of escalators and two Tube lines are only 15 minutes walk apart. It was only last week that one of us realized that Goodge Street is a short walk from Euston Station.”

⁷Porter stated that tighter environmental standards “do not inevitably hinder competitive advantage against foreign rivals; indeed, they often enhance it. Tough standards trigger innovation and upgrading”. Similarly, Porter and Van der Linde (1995: 98) claim that environmental regulations can “trigger innovation (...) that may partially or more than fully offset the costs of complying with them”. This idea goes back to the notion of “induced innovation”, developed in Hicks (1932), and has also been taken beyond Porter’s original application to environmental regulation (see *e.g.* Aghion, Dewatripont and Rey (1997) for a paper that analyzes related issues in a more general setup).

in practice. The fact that measureable progress often takes time to occur makes it for example difficult to keep “all else equal”, while it is also not clear how “an improvement” is to be defined in the first place. As a result of these complications, the literature has not settled upon a consensus with respect to this issue (see *e.g.* Gray (1987), Jaffe and Palmer (1997), Berman and Bui (2001), and Copeland and Taylor (2004)). By analyzing the behavior of commuters who were faced with a short-lived, temporary constraint on the London underground network, the present study overcomes many of these problems.

In addition, our study is informative on the existence, strength and persistence of habits. As noted by Wood and Neal (2009), research on habits is important since about 45% of people’s behavior is repeated on a daily basis. Commuter behavior is an exponent of this. Along these lines, Goodwin (1977: 95) has argued that “the traveler does not carefully and deliberately calculate anew each morning whether to go to work by car or bus. Such deliberation is likely to occur only occasionally, probably in response to some large change in the situation”.

Finally, our paper provides evidence on the effects of a public transport strike. Although there are some earlier studies analyzing disruptions in transportation networks (see Van Exel and Rietveld (2001) for an overview of this sparse literature), they tend to rely on survey data – thereby leading to small sample sizes and preventing a clean comparative analysis of travel patterns before and after the disruption (Zhu and Levinson, 2011: 19). As we will explain in greater detail in Section 4, the present study has the entire population of actual travel movements on the London underground at its disposal, which brings advantages over earlier contributions.

The remainder of this paper is structured as follows. We start by providing background information on the London underground network in Section 2. Subsequently, Section 3 describes the tube strike that took place in February 2014, after which we discuss our dataset in Section 4. To motivate certain choices in our empirical exercise, we provide some notable descriptive statistics in Section 5 (which may also be of general interest and relevance), after which Section 6 continues by describing our method. We then present our analysis of the effects of the strike in Section 7, after which Section 8 interprets our results. Section 9 concludes.

2 The London underground network

Over the sample period considered in this paper (January 19 to February 15, 2014), the London underground network consisted of 11 different lines, connecting 270 different sta-

tions. It is operated by London Underground Limited (which is fully owned by Transport for London, the corporation that runs most of London’s public transport services) and covers 402 kilometers of track. The London underground serves up to 4 million passenger journeys per day and is a popular mode of transportation for many people living and/or working in London.

Crucially for our paper, users of the London underground face imperfect information on several relevant features of the available alternative routes in getting from A to B. An important source for this imperfection is the London tube map – a major aid to travelers in finding their way through the network. It is a schematic transit map, showing only *relative* positions of tube and train stations along lines. Consequently, the map is geographically distorted and gives users false impressions when it comes to actual distances between two points – especially when comparing points along different tube/train lines.⁸ The distorted nature of the map gives rise to further problems of similar nature when traveling from the exit station to the final destination, which is likely to lie somewhere in between the various lines, where the map is not well-defined.

Next to commuting time, travelers are initially also uncertain on many characteristics of the various available alternatives. How crowded is a particular line at the preferred time of travel? Is the route from the exit station to the final destination convenient? There could for example be a supermarket along the way, or a place that serves good breakfast.

An important way in which these various uncertainties can be reduced, is by actually *trying* the available alternatives – i.e. through experimentation. And because of the strike that we are about to describe in the next section, many travelers were forced to do exactly that in the first week of February 2014.

3 The strike

On January 10, 2014, the Rail Maritime Transport union, the largest trade union in the British transport sector, announced a 48-hour strike of London tube workers. The strike was to begin on Tuesday evening (21:00h) 4 February. It was called for in response to the announcement of a plan by Transport for London to close ticket offices and to introduce non-compulsory redundancies for part of its workforce.

⁸Guo (2011) calculates that for the London underground map, the correlation between actual and “mapped” distances is only 0.22. He also gives several examples of actual distortions. A famous case is that of Covent Garden and Leicester Square: both stations are only 260 meters apart, but the 20 second tube ride (at £4.70) remains in high demand.

The decision to participate in the strike remained with individual workers. In the past, it has therefore sometimes been the case that unions called for a walkout, but workers did not act accordingly. For example, in December 2005, the union called for action on New Year’s Eve but, according to an official bulletin, the “strike has had little impact on London Underground’s services (...) The majority of station staff have ignored the call for industrial action and are working normally.”⁹

However, more workers participated in the February 2014 strike. Due to the resulting non-availability of staff members, 171 (out of 270) tube stations remained closed for at least part of the duration of the strike (see Figure 1 for a visualization). Which stations were closed was decided in a chaotic process heavily influenced by the availability of staff and was not known until the morning of the strike. Many of the factors were outside the control of Transport for London, such as the degree of unionization along various lines/stations and the willingness of individual employees to participate in the strike. The underlying process does not seem to have been strategic. As a result, treatment and control stations are similar along observable characteristics. Most importantly, the number of neighboring stations within walking distance (1km, 2km, and 5km) is not statistically different between treatment and control stations.¹⁰ As our main outcome is the probability to switch commuting route, this is the key dimension along which our exercise requires similarity.

There are a number of stations on the network that serve multiple lines and were only partially closed during the strike, with one or more lines still operating on them. In our econometric exercise we code these stations as closed, even though some commuters might have been able to use them. During the two strike days, many lines had fewer trains running, while there were no services on the Bakerloo line, the Circle line, and the Waterloo & City line. For a large part, these closed lines run in parallel and overlap with lines that remained open. As of Friday morning February 7, all stations were open again with services running as usual.

Insert Figure 1 about here

The previous strike affecting the London underground network as a whole took place in 2010.¹¹ In that year, certain stations were closed on the following dates (the number of

⁹See tfl.gov.uk/info-for/media/press-releases/2005/december/london-underground-service-update – 2000hrs.

¹⁰The p-values of tests of similarity are 0.76 (1km), 0.54 (2km) and 0.12 (5km).

¹¹Since 2010, there have been several minor strikes - affecting individual lines/stations only. For example, there was a minor disruption on the Bakerloo Line on January 15, 2011 due to staff protests.

closed stations follows within brackets): October 3 (100), November 2 (95), November 3 (134), November 28 (94) and November 29 (125). Before that, the network suffered from major disruptions on June 9-11, 2009 (strike), September 3-5, 2007 (strike), July 7-25, 2005 (7/7-bombings) and June 29, 2004 (strike). No individual travel data are available for the periods around these earlier disruptions, as a result of which we cannot analyze their impact.

The February 2014 strike has several desirable features that make it particularly suited for studying the question at hand. It was the first major disruption in over three years, as a result of which the sample is likely to contain many individuals who hadn't been subjected to forced experimentation on such a grand scale before.¹² Moreover, the strike was not complete: about 37 percent of all stations remained open, actually enabling travelers to experiment within the tube network (which would not be possible if *all* stations were closed). The partial nature of the strike furthermore leaves us with a sizable share of both “treated” and “non-treated” commuters, which will help our analysis below. The first full strike day (February 5) was rather rainy. According to weatheronline.co.uk there was 7mm of rain in London during the morning, which is likely to have discouraged travelers from experimenting by bike or foot, in which case they would no longer show up in our data and we no longer know whether they went to work or worked from home.¹³ Finally, with a duration of 48 hours, the strike was relatively short-lived. Thus any changes in behavior are likely to be driven by optimality-considerations – not by changes in habits, which are believed to take much longer to form (Wood and Neal, 2009).

4 Data

Data on commuting behavior are particularly well-suited to analyze decision-making and learning by individuals. First, commuters are faced with a very similar problem, to get from home to work and back, on a workdaily frequency. Our dataset shows us, in great detail, how they answer this problem at the exact same frequency. Second, many aspects of the problem (such as journey time) are quantifiable and observable to us, which allows

Other lines remained unaffected. Occasionally, technical failures and the like have had similar impacts.

¹²The fact that London attracts about 350 thousand new inhabitants per year (many of them tube-using workers), implies that about 1.2 million Londoners at the time of the most recent disruption were not living there during the previous strike in 2010. This amounts to about 25 percent of London's current working-age population. Others would have changed jobs and/or houses in the meantime – resetting any previous commute-specific information they acquired.

¹³Also see the advance warnings, for example “Weather hits trains as London tube strike begins” in *The Guardian* of February 4, 2014.

for a rich analysis. Third, the outcome matters (see the large estimates on the costs of commuting by Stutzer and Frey (2008) and Ahlfeldt *et al.* (2015)), while the solution is not trivial: there are for example no less than 13 reasonable ways to travel from King’s Cross to Waterloo. Finally, the fact that only part of the network closed down during our sample period leaves us with treated and non-treated individuals – which greatly facilitates a comparative analysis.

The dataset that we use was provided to us by Transport for London. It contains all individual travel movements on the London public transport system from Sunday January 19 to Saturday February 15, 2014. For all modes of public transportation other than bus (that is: for tube, train, tram, DLR, and boat), the dataset provides us with the station of entry for a particular journey, the station of exit, as well as the times of check-in and check-out.¹⁴ Since the February 2014-strike applied to the tube network (all boat, bus, train, tram, and DLR stations remained in operation), the focus of our study is on journeys that involve the underground.

Over our sample period, payments for individual journeys could be settled in two ways: either by purchasing a ticket that is valid for a certain time period and/or area, or by using a re-chargeable plastic card called an “Oyster Card”. The latter is used in about 80 percent of all journeys. Each Oyster Card has a unique number, as a result of which we are able to track individual travel behavior of Oyster Card users over our sample period. As we want to observe how repeat-behavior changes after a disruption, we analyze Oyster Card-using commuters. Most of these commuters face the same problem of getting from A to B and back every weekday – thereby enabling us to extract information at a reasonably high frequency.

We identify commuters as individuals who use London’s tube network during every non-strike working-day in our sample (of which there are 18) between 7am and 10am. The presence-requirement leaves us with a balanced panel of tube-users, while the time-requirement implies that we only look at the morning rush hour, which runs from about 7am to 10am, see Figure 3 below. We want to focus on commuters that are stuck in a daily routine, which is more likely in the morning than in the evening, where Londoners may pursue other activities before heading home.

We create two main datasets, which we call “unconditional” and “balanced”. In the unconditional dataset we use no information from the post-strike period to identify our

¹⁴We don’t have this information for bus journeys (we can only see whether they take place or not, as Transport for London does not record exits from buses). Along these lines, it should be noted that (next to the 270 tube stations) London hosts 366 train stations, 39 tram stations, 45 DLR stations, and 25 boat stops.

commuters. We only require commuters to be present in our dataset on every morning during the pre-strike period. The advantage of this dataset is that there is no selection of data based on outcomes. In the balanced sample, on the other hand, we require commuters to be present on every morning in both the pre- *and post-strike* period, as well as on at least one strike day.¹⁵ The advantage of this balanced panel is that it enables us to observe behavior of all commuters included in our dataset before, during and after the strike, as by construction, no one abandons the TfL-system. The disadvantage of this approach is that we require the commuter to remain within TfL-services, which is a selection on an outcome. As we will show in Section 7.1 using the unconditional sample, treated and untreated individuals are about equally likely to opt-out of the TfL-system post-strike. This suggests that the selection procedure underlying the balanced sample is not problematic. We therefore proceed by using the balanced sample as our main one – taking full advantage of the fact that this sample allows for a detailed comparison of behavior before and after the strike.

We infer the “usual” entry and exit stations of travelers by setting it equal to the station which they use most frequently during the pre-strike period (the “modal station”). A small minority of about 700 individuals, approximately 4% of the sample that we are left with at this stage, have multiple modes on either or both ends. Since it is not obvious how they are to be dealt with in our analysis (which is focused on identifying “deviations from the mode” – assuming the latter is unique), we drop them as well.

Cutting the data in this way leaves us with an unconditional sample of 60,747 Oyster Card IDs, and a balanced panel of 18,113 Oyster Card IDs. For each of these IDs we have 20 working days of observation.

Throughout this paper, we employ a rather strict definition of the concept of a “commuter” as we require them to behave in a very consistent manner. Consequently, we are definitely making some type II errors here, excluding individuals who actually are commuters. We also miss individuals who use multiple Oyster Cards, as well as those who were absent from London’s public transport system for one non-strike weekday or more over our sample period. Given the size of our data, this is not a major problem. Moreover, if anything, this strict selection procedure implies that the mode-change probabilities reported below are a lower bound, as we have selected those individuals who adhered to a rather strong routine (potentially even a habit) during the pre-strike period. In addition, this procedure is likely to select individuals who have few alternative routes

¹⁵The latter requirement serves to ensure that we analyze the behavior of individuals who were actually present on the underground during the disruptive phase (instead of working from home) – thereby making sure that they have had a chance to explore alternative routes during this period.

available other than their modal one – again implying that our estimates are going to form a lower bound.

5 Descriptive statistics

Given the novelty and level of detail that is present in our dataset, we start by providing some descriptive statistics based upon the entire data population – which may be of independent interest. Next to that, these statistics are used to motivate certain choices that we make in the econometric exercise that is to follow.

First of all, our data are informative on the dominant public transport commuting patterns within the Greater London area. Figure 2 displays stations of first entry in the morning and evening for one of the days in our sample, January 31, 2014. Circle-sizes correspond to relative station use. The morning commute is characterized by a dispersed start, often from residential areas in the outskirts of London or the large commuter railway stations on London’s periphery. The evening commute, on the other hand, is much more concentrated – starting from well-known business districts like Canary Wharf and the City.

Insert Figure 2 about here

Second, due to the absence of other significant events during our sample period all non-strike working days were approximately equally busy:¹⁶ the busiest day is Friday January 24, 2014 (with 19,301,730 data entries and 3,652,851 unique travel IDs) while the quietest non-strike day of our sample is Wednesday February 12, 2014 (with 18,259,114 data entries and 3,496,720 unique travel IDs). Within each day, activity followed a standard rush hour pattern, an example of which (again that of January 31, 2014) is displayed in Figure 3. As one can see from this figure, the morning commute runs from about 7am to 10am, which motivates our earlier choice along these lines.

Insert Figure 3 about here

Finally, Figure 4 shows the evolution of some key variables of interest for all weekdays in our sample period. The top-left panel shows the fraction of commuters by the balanced panel who enter at their modal station, while the top-right panel shows the same at the

¹⁶Around January 28 (corresponding to day 9 in Figure 4), there was heavy rain and flooding in southern England, which might have influenced travel decisions of some commuters. However, during our sample period, no announcements of travel disruptions on the London Underground (other than the tube strike) could be found.

exit-margin. The two strike days can be found in between the vertical lines. As one can see from the two panels, far less commuters were able to use their modal station during the strike – which implies that a substantial number of individuals were forced to explore alternative routes. Moreover, the post-strike data also suggest that the strike brought about some lasting changes in behavior, as the fraction of commuters that make use of their modal station can be seen to drop after the strike.¹⁷

The lower two panels of Figure 4 provide information on journey times: as the bottom-left panel shows, the duration of the average journey on London’s public transport system went up during the strike (by about 6%), while the bottom-right panel shows that this increase in average duration was also accompanied by an increase in dispersion.

Insert Figure 4 about here

6 Estimation Strategy

As set out before, the partial nature of the February 2014 strike conveniently leaves us with treated and non-treated commuters. This enables a difference-in-differences exercise, which is the approach that we will take in this paper. After all, we are ultimately interested in the question whether individuals who were “treated” (i.e.: forced to experiment) during the strike, went on to behave any differently from their non-treated peers (the control group) in the post-strike period. Consequently, we will typically estimate regression equations that are of the following form:

$$d_{it}^{\text{mode}} = \alpha_i + \beta \cdot d_t^{\text{post}} + \gamma (d_t^{\text{post}} \cdot d_i^{\text{treat}}) + \epsilon_{it}. \quad (1)$$

Here, d_{it}^{mode} is a dummy-variable that takes the value 1 if individual i makes his “modal journey” (i.e.: travels from his modal station of entry to his modal station of exit) on date t , d_t^{post} is a dummy-variable that takes the value 1 in the post-strike period, while d_i^{treat} is a dummy-variable that takes the value 1 if individual i was part of the treatment group, defined in different ways as described below. ϵ_{it} in equation (1) is the error-term, β measures time effects, while γ captures the treatment effect. Data from the two strike days are not used in the estimations (only to identify the treatment group). We estimate equation (1) using OLS and apply robust standard errors. When the treatment is defined at the station level we cluster standard errors at this level as well. We also include individual fixed-effects as captured by α_i in equation (1). The reason is threefold.

¹⁷Establishing this formally is the objective of the remainder of this paper.

First, fixed effects control for unobserved demographic factors, such as age, that may affect an individual’s inclination to experiment. Second, they control for station or area characteristics, which may influence the propensity to switch. Third, fixed effects correct for the fact that different individuals use their modal station with different intensity.

As we will clarify in the remainder of this section, identifying the treatment group from our dataset is non-trivial. To ensure robustness, we will therefore show results for three different definitions of treated commuters – where all measures have their specific advantages and disadvantages.

Our first measure of treatment defines treated individuals as all those who deviated from their pre-strike modal journey during the strike. This would include individuals who were forced to explore a new route due to closure of an entry, exit, or connecting station, but will also encompass those who deviated from their pre-strike mode for non-strike related reasons. The second measure of treatment takes a more direct approach: in this exercise, we take individuals to be treated if their pre-strike modal station (either entry and/or exit) closed down during the strike. After all, we can be reasonably sure that these individuals were not able to travel to or from their modal station during the strike – as a result of which they were definitely forced to explore alternatives. This measure does not allow for selection into treatment, as it seems reasonable to assume that the closure of stations is random with respect to individual characteristics. However, it suffers from the fact that it is likely to pool a significant number of treated individuals with the non-treated group. The reason is that many individuals in our dataset travel from station A to station B via at least one connecting station C. Closure of the latter would force this individual to explore alternatives, but we don’t observe connecting stations in our data, only stations of entry and exit. Consequently, our second measure of treatment is likely to lead to type II errors and underestimate the true effect. Our third measure of treatment is based on travel time: here we take individuals to be treated if their travel times during strike days were sufficiently different (i.e.: longer or shorter) from their travel times during the pre-strike period. This method identifies those commuters who had a very unusual experience during the strike as measured by time. It does not rely upon our definition of closed stations (as pointed out in Section 3, some stations were only partially closed), while it also side-steps our concept of “deviations from the modal commute” – thereby giving us some heterogeneity of treatment. This measure of treatment is however prone to errors of both the first and the second kind (i.e.: there will be both “false positives” and “false negatives”). After all: if an individual had a different journey time on strike days, that does not necessarily imply that he was actually exploring an alternative route.

It could simply be the case that his modal route took much longer due to congestion on the network, or due to the reduction in the number of trains running. Similarly, it is also possible that a commuter explored a different route, but that this did not lead to a markedly different travel time.

For all treatment measures we check that the treatment and the control group are similar for observables, such as average journey duration, and the number of stations near modal entry and exit station. Figure 5 displays trends for treatment and control groups for our outcome variables, taken from the balanced sample. The graph shows the fraction of commuters who travel using both their modal entry and exit station. Visually, the pre-strike levels and trends look similar for both groups in all cases, with the exception of the “deviate during strike” group, where levels in the pre-strike period differ somewhat.¹⁸ This suggests that commuters who switch during the strike, are more likely to deviate from their mode before the strike as well. This points to some selection into treatment for this variable, although this is partly offset by the treatment dummy in the estimation equation. It is furthermore not clear whether any residual that is left after accounting for the treatment dummy, biases our estimates in the up- or downward direction (if at all). On the one hand, it could point to easier substitution for the treatment group in which case we would expect them to have a greater propensity to switch if treated. On the other hand, it may also be the case that the treatment group experimented more pre-strike, uses better routes to start with as a result, and thus has a smaller propensity to change behavior post-treatment.

Insert Figure 5 about here

7 Findings

In this section we present our main estimation results. Section 7.1 describes outcomes of the main regressions results, Section 7.2 provides some robustness checks for these results, Section 7.3 analyzes the effects on travel time, while Section 7.4 investigates the mechanism underlying the main effect.

¹⁸The treatment definition employed in this case implies that, by construction, the fraction of commuters in the control group traveling their pre-strike modal commute jumps to 1 on both strike days.

7.1 Core results

As set out in Section 6, we rely upon difference-in-differences estimations to ask whether treated commuters were more likely to deviate from their pre-strike modal journey in the post-strike period, relative to their non-treated peers. The answer to this question can be found by looking at the sign of our estimate of the treatment effect γ in regression equation (1).

We start by considering whether the strike affected people’s decision to abandon the public transport system as a whole. Outcome s_{it} indicates presence in the TfL-services, excluding buses. It equals “1” for commuters who use the tube, train, boat or the DLR, it equals “0” for people who use all other modes of transportation or stay at home. Table 1 uses the “unconditional” sample. In column 1 the treatment group consists of those people who deviated from their usual route on either strike day. The coefficient on the post-strike dummy is large and significantly negative in all specifications (capturing natural attrition), which we would expect given that the modal station is identified using information from the pre-strike period only. The coefficient on the interaction of post-strike and treatment dummies is not statistically significantly different from zero at conventional significance levels. This suggests that for this measure of treatment, people do not have a greater propensity to opt out of the system if they had to change their commute due to the strike. The second column identifies the treatment group as those commuters whose modal commute started from or ended at a station which participated in the strike. Again the coefficient is not statistically different from zero. In the three remaining columns we measure the treatment by commuters who had a substantially longer or shorter journey on strike day (the “factor” mentioned in the top-row of Table 1 considers different definitions of “substantially”: a factor 1.2 implies that an individual was included in the treatment group if their average morning commute during the strike was at least 20% longer or shorter than their pre-strike average). Here, we find significant negative coefficients. This suggests that some commuters who had bad experiences during the strike indeed decided to abandon TfL-services. The propensity to abandon is increasing in the intensity of journey time deviation on strike day (our estimate of γ is larger in column 5 than it is in column 4, which is in turn larger than column 3). Our main measures, in the first two columns, show no effect. This suggests that selection into other means of transportation is not a great concern for these measures, which is why we proceed with the balanced sample for the rest of the analysis. For the time-based definitions of treatment, the fact that our estimate of γ is negative suggests that any treatment effects on switching behavior reported below are likely to form a lower bound (since they are based upon a

sample which excludes the group of commuters that switched radically by ceasing to use the TfL-system post-strike).

Table 1: OLS-DiD results.¹⁹

	(1: not on mode)	(2: mode on strike)	(3: factor 1.2)	(4: factor 1.5)	(5: factor 2)
	s_{it}	s_{it}	s_{it}	s_{it}	s_{it}
d_t^{post}	-0.141*** (0.001)	-0.163*** (0.00086)	-0.147*** (0.0009)	-0.146*** (0.0007)	-0.147*** (0.0006)
$d_t^{post} \cdot d_i^{treat}$	-0.024 (0.0024)	0.0023 (0.0021)	-0.021*** (0.0011)	-0.032*** (0.0011)	-0.0415*** (0.0012)
obs	1,093,446	1,093,446	1,093,446	1,093,446	1,093,446

Table 2 reports our main estimates of interest. In short, they suggest that those who were forced to explore alternatives during the strike, were less likely to use their pre-strike modal commute after the restriction was lifted. This result is robust to different estimation strategies (as we will document below), while it arises no matter how we define the treatment group: again, in Column 1 the treatment group consists of all individuals who deviated from their modal morning commute during the strike. Column 2 includes commuters in the treatment group if their pre-strike modal commute started or ended at a station which was closed during the strike. In the last three columns of Table 2, individuals are considered to be part of the treatment group if they experienced a substantially longer or shorter journey time during the strike.

Table 2: OLS-DiD results.²⁰

	(1: not on mode)	(2: mode on strike)	(3: factor 1.2)	(4: factor 1.5)	(5: factor 2)
	d_{it}^{mode}	d_{it}^{mode}	d_{it}^{mode}	d_{it}^{mode}	d_{it}^{mode}
d_t^{post}	-0.0108*** (0.00186)	-0.0466*** (0.00185)	-0.0402*** (0.00175)	-0.0464*** (0.00140)	-0.0504*** (0.00128)
$d_t^{post} \cdot d_i^{treat}$	-0.0569*** (0.00242)	-0.00860*** (0.00248)	-0.0205*** (0.00245)	-0.0201*** (0.00293)	-0.0113** (0.00451)
obs	312,156	312,156	312,156	312,156	312,156

¹⁹In this table, as well as in all tables that are to follow, * denotes significance at the 10% level, ** implies significance at the 5% level, while *** indicates significance at the 1% level.

²⁰The number of observations in the treatment group is: 243,254 for Column (1), 188,080 for (2), 185,081 for (3), 83,380 for (4) and 28,732 for (5).

In all three exercises, the interaction coefficient γ (measuring the difference-in-differences) is consistently estimated to be significantly negative. This implies that individuals who were part of the treatment group (i.e.: those who were forced to explore alternative routes during the strike) were less likely to use their pre-strike modal commute in the post-strike period.²¹ By a revealed preference-type argument, this suggests that a significant fraction of commuters had failed to find their optimal journey before the strike. After all: post-strike, all routes were available again, including the pre-strike modal one, so a failure to pick the latter option suggests that the commuter has found a better alternative during the disruption.

Our results are unlikely to be driven by a formation of new habits during the two strike days. Not only do they typically take much longer to be established (Wood and Neal, 2009), but the observed behavior of commuters is also inconsistent with this hypothesis: after the strike, many of them continue to explore alternative routes (leading to a prolonged experimental phase) after which they eventually settle on a new modal choice.²²

Looking at the magnitudes of estimates across the various tables is informative as well. Doing so shows that our estimate of the treatment effect γ is significantly larger in Column 1, while that of β (the coefficient on the post-strike dummy d_t^{post}) is significantly smaller. A negative estimate for β (the coefficient on the post-strike dummy d_t^{post}) is to be expected due to natural changes in workplace or home and a definition of the modal station that is based on the pre-strike period only. However, a large negative estimate for β could also suggest that our treatment-definitions err by including treated individuals in the non-treated group. As anticipated in Section 6, it is to be expected that our second measure of treatment is particularly prone to this statistical error of the second kind – and

²¹In the specification where we identify the treatment group via the time factor, this effect is consistently stronger for those travelers who experienced a shorter commute during the strike, which is intuitive: the alternative route is likely to look less attractive to a commuter if (s)he experienced a long delay during the strike.

²²To give a random example: one commuter in our dataset consistently traveled from station S to station C in the pre-strike period (for privacy-reasons, we are not allowed to give full station names). During the strike, (s)he experiments with entering at E – using the DLR to travel to C. In the post-strike period, (s)he first alternates between both options (seemingly comparing them) after which (s)he settles for the newly-found DLR-based route. There are also more determined examples: another commuter consistently travels from R to J on every morning before the strike. Both stations however closed down during the strike, in response to which (s)he switched to traveling from N to W on the first strike day. Subsequently, (s)he sticks with this new alternative (which has a shorter duration and a lower variance) for the remainder of our sample period. As direct evidence for experimentation, we also computed a measure that identifies commuters who enter and exit at the same station as on the previous day. For all five treatment definitions, we get negative and significant treatment effects, with a magnitude between -0.01 and -0.4. This suggests that a sizeable fraction of those commuters who had to reconsider their journeys during the strike, experimented in the days that followed it.

indeed, the absolute value of the estimate of γ is lowest in this specification, while that of β is among the highest. In Column 1 on the other hand, the estimate of β is closer to zero (which makes sense from a theoretical point of view), as a result of which this table contains our preferred estimates for the treatment effect γ . The key magnitude is the estimate of γ , which equals -0.057 for this definition of the treatment group. It suggests that the probability to switch on a given journey is five percentage points greater in the post-strike period.

To understand where the results of Table 2 are coming from, one can also increase the level of detail and estimate a specification that distinguishes between entry and exit (so the LHS-variable in that regression is either $d_{it}^{\text{entry_mode}}$ or $d_{it}^{\text{exit_mode}}$). Results of this exercise, recorded in Table 3, indicate that the treatment effect γ tends to be bigger at the exit-end. This is intuitive since the exit-end of the morning commute typically lies in the city center (recall Figure 2) where station-density, and hence substitutability, is higher. Also note that Table 3 contains only two estimates (in *italics*) that are not significant at any regular level of significance (whereas all other estimates are significant at the 1% level). However, they show up at exactly those places where this is plausible, namely when we look at what closure of a modal *entry* station does to the choice of station of *exit* (and vice versa).

The coefficients in Tables 1 and 2 are fairly straightforward in their interpretation due to the probabilistic nature of our exercise. Some complication arises as commuters in our pre-strike sample only make their modal journey for about 84% of the time: an estimate for γ of -0.03 therefore implies that treated individuals will make their pre-strike modal commute with a probability that is 3 percentage points lower compared to their non-treated peers. This does however not imply that 3% is also the fraction of switchers in our sample.

Table 4, on the other hand, does produce information on the fraction of switchers – as such a number is arguably easier to interpret. This table is constructed by first identifying those commuters who made the *exact* same morning commute (as far as stations of entry and exit are concerned) during all 10 working days of our pre-strike sample. Hence, all these individuals (whom we refer to as “pre-strike habituals”) are selected so that they make their modal commute with probability 1 in the pre-strike period. We subsequently ask: how many percentage points higher is the fraction of “post-strike switchers”²³ in the

²³“Switchers” are defined as those individuals who made a different commute than their pre-strike modal journey on the last working day of our sample (Friday February 14). This exercise therefore assumes that the experimentation phase, triggered by the strike-induced forced episode of experimentation, was over by this time (also recall footnote 22). Requiring them to deviate for more than one day, yields very similar results.

treatment group relative to the fraction of switchers among non-treated commuters?²⁴

Table 3: Estimates of γ when distinguishing between entry and exit margin.²⁵

Treatment definition	(1)	(2)
	$d_{it}^{\text{entry_mode}}$	$d_{it}^{\text{exit_mode}}$
not on mode	-0.0267*** (0.00164)	-0.0470*** (0.00224)
either on strike	-0.00480*** (0.00170)	-0.00480*** (0.00170)
entry on strike	-0.00697*** (0.00190)	<i>0.000569</i> <i>(0.00247)</i>
exit on strike	<i>-0.00146</i> <i>(0.00173)</i>	-0.00748*** (0.00230)
time factor(1.2)	-0.0141*** (0.00168)	-0.0154*** (0.00227)
time factor(1.5)	-0.0111*** (0.00207)	-0.0175*** (0.00270)
time factor(2)	-0.00766*** (0.00329)	-0.0113*** (0.00411)
obs	312,156	312,156

Table 4: Fraction of switchers among pre-strike habituals.

Treatment definition	
not on mode	5.42%
mode on strike	2.64%
time factor(1.2)	1.24%
time factor(1.5)	1.86%
time factor(2)	2.81%
obs	6,946

²⁴Here, it is absolutely essential to look at results relative to a non-treated control group since this exercise is obviously prone to “regression to the mean”: given that the habituals were using their modal station with probability 1 in the pre-strike period, they can only make (weakly) less use of it post-strike. The control group of non-treated commuters allows us to correct for mean reversion.

²⁵Note that this table is based upon 14 regressions of the same form as equation (1). For space-constraints, we only report our estimates of γ . In Column 1 we report results when the dependent variable is a dummy for entry at mode, while Column 2 reports estimates for a dummy for exit at mode. Estimates of other coefficients are available upon request.

As can be seen from the table, our data suggest that (depending on whom we consider to be treated) the fraction of post-strike switchers is 1.2 to 5.4 percentage points higher in the treatment group. Since results for our last two measures of treatment (“mode on strike” and the method using the “time factor”) are again likely to be biased by type II errors, we believe that the true number lies closer to 5.4 percentage points (the number we obtain when defining the treatment group as those who deviated from their modal journey during the strike). This is a strong result given that the individuals underlying this exercise all seemed to be stuck in a very regular habit before the strike, as they were selected exactly because they were making the same commute on every single of the twelve mornings in the pre-strike sample. The selection method could furthermore imply that these commuters have only few viable alternatives available, which also biases the results against switching. Moreover, exploring a new route during a tube strike is typically not a pleasant experience, due to the associated chaos and crowdedness, while there were also fewer trains running during the February 2014-strike – causing further delays. Consequently, it is likely that numbers would have been even larger after considering voluntary experimentation under tranquil conditions. In line with our earlier findings, this again provides evidence that a substantial proportion of commuters had failed to find their maximum before the tube strike of February 2014.²⁶

7.2 Robustness

We have found our results to be very robust. While Section 7.1 has already shown this for different definitions of the treatment group, this section will show that it also holds with respect to alternative regression specifications. This can for example be seen from Table 5.

A well-known criticism of OLS-DiD panel data regressions, is that autocorrelation for individuals over time decrease estimated standard errors (Bertrand, Duflo and Mullainathan, 2004; henceforth “BDM”). In column 1, we therefore report results generated by BDM’s most conservative robustness check – namely the one where the data are collapsed to two observations for each individual: one observation pre-strike and one post-strike, and we collapse our LHS variable by computing the mean number of modal journeys before and after the strike. Given that all variables used in our specification are binary indicator variables, coefficients remain numerically identical in this exercise. They still

²⁶We are not claiming that these commuters have found their global maximum post-strike: all we are saying is that they have found something better than their pre-strike mode, but it is very well possible that further improvements are still possible.

are highly statistically significant.

Finally, column 2 shows our baseline estimates of γ when we restrict our sample to those individuals who enter and exit on the same line (“SL”). As set out before, identifying the treatment group is somewhat challenging in the full sample as many individuals make use of connecting stations during their commute. Closure of a connecting station implies that such an individual was treated during the strike (even if his entry and exit station remained open), but unfortunately we do not observe data on connections. This concern plays no role when we limit ourselves to those commuters who enter and exit on the same line (as they are unlikely to travel via a connecting station). Due to the “same line”-restriction we are left with fewer observations, but our main result continues to emerge – albeit somewhat less significantly, which is no surprise given the smaller sample size, and typically smaller in magnitude. The latter is to be expected since the scope for experimentation is substantially smaller for commuters who use only one line, as commuters who use multiple lines and connections have more dimensions along which they can deviate.

Table 5: Estimates of γ across specifications.²⁷

Treatment definition	(1: BDM)	(2: SL)
	d_{it}^{mode}	d_{it}^{mode}
not on mode	-0.0569*** (0.00311)	-0.0414*** (0.00540)
either on strike	-0.00860*** (0.00333)	-0.0208*** (0.00608)
time factor(1.2)	-0.0205*** (0.00331)	-0.0116** (0.00568)
time factor(1.5)	-0.0201*** (0.00406)	-0.0223*** (0.00664)
time factor(2)	-0.0113* (0.00615)	-0.0337*** (0.0105)
obs	34,684	47,052

In a further robustness check, we interact each day in the post-strike period with each of the treatments separately, to study the timing of the treatment effect. We find

²⁷As with Table 3, this table is based upon 15 underlying regressions of the same form as equation (1). Due to space-constraints, we only report our estimates of γ . Estimates of other coefficients are available upon request.

coefficients of similar magnitudes for all these days, and no clear trend in time in either of the specifications.

7.3 Effects on travel time

A follow-up question to ask at this stage is: what was the effect of the strike on commuting times? We do not observe the duration of the entire commute, since commuters are not on our radar before they check-in to/after they check-out of TfL-services, but we can calculate the amount of time they spent on London’s public transport network. We do not know whether the time commuters spend traveling to and from the stations correlates with the observed journey times in the TfL-systems, and thus these results need to be interpreted with caution. However, since time spent on the London underground during rush hour is well-known to be particularly unpleasant (also compared to other modes of commuting), minimizing this time is likely to receive a significant weight in the objective functions of most commuters. Moreover, these numbers should at the very least be of interest to the providers of public transport services, as they have an incentive to minimize the time people spent on their facilities. After calculating these durations (defined as the difference between the last contact with a TfL-service and the first, with the exception of buses), we estimate the following regression:

$$\ln(\text{duration}_{it}) = \alpha_i + \beta \cdot d_t^{\text{post}} + \gamma (d_t^{\text{post}} \cdot d_i^{\text{treat}}) + \epsilon_{it}$$

Note that our dependent variable is the *natural logarithm* of duration, so that coefficients can be interpreted as percentages. Once more, our main interest lies in the estimate of γ . Estimation results are shown in Table 6 below, again for the five different characterizations of the treatment group. As can be seen from the table, our estimate of γ is consistently negative which suggests that commuters who were part of the treatment group were able to cut their “time spent on public transport” by more than their non-treated peers. On average, the treatment group seems to be able to cut their journey time by about 1% more. Given that the average journey in our sample lasts approximately 32 minutes, this amounts to a time-gain of about 20 seconds on a one-way commute. Note that the 20 seconds-statistic is an average taken over *all* treated commuters – so it includes both treated commuters who found a quicker route post-strike, and those who did not (and remember that the latter group constitutes the majority with about 95% of all treated individuals). Consequently, 20 seconds is the absolute lower bound for the one-way time-gain experienced by the subset of commuters that did find a quicker commute

thanks to the strike.

Table 6: OLS-DiD results for travel time.

	(1: not on mode)	(2: mode on strike)	(3: factor 1.2)	(4: factor 1.5)	(5: factor 2)
	$\ln(duration_{it})$	$\ln(duration_{it})$	$\ln(duration_{it})$	$\ln(duration_{it})$	$\ln(duration_{it})$
d_t^{post}	0.00711*** (0.00164)	0.00113 (0.00158)	0.00125 (0.00132)	0.000670 (0.00108)	-0.000698 (0.00103)
$d_t^{\text{post}} \cdot d_i^{\text{treat}}$	-0.0124*** (0.00206)	-0.00518** (0.00204)	-0.00548*** (0.00198)	-0.00977*** (0.00261)	-0.0121*** (0.00430)
obs	312,103	312,103	312,103	312,103	312,103

7.4 Mechanism

Given that the previous sections have established that treated commuters were more likely to switch stations and cut travel time in the post-strike period than their non-treated peers, a logical follow-up question is: why? In the remainder of this section, we will provide evidence which suggests that this is due to the existence of informational imperfections. To make this point we use information on two characteristics of the London underground system that are not easily observed by commuters, namely map distortion (Section 7.4.1) and line speed (Section 7.4.2).

7.4.1 Map distortion

As noted before, an important source of imperfect information is the fact that the London tube map provides a distorted picture of reality. For the exercise in this subsection, we quantify these distortions in the following way. For each station on the map (S) we list those stations that lie within a 2km radius (which is about a 25 minute walk) from S .²⁸ We subsequently correlate the true distance between these stations, with the distance on the tube map, which we have digitized. Subtracting the resulting correlation from 1, gives our measure for distortion. Other measures where we take 5km radius, or consider the closest 10 stations give measures of distortion that correlate highly with the 2km-radius measure, with correlation coefficients of 0.96 and 0.94 respectively.

Map distortions are not constant across London: some people live in areas where the tube map is more distorted than others, the general rule being that distortion increases

²⁸Very similar results are obtained if we use a radius of 5 kilometers.

with distance to central London. Thanks to this spatial variation, we are able to ask: do commuters who live in areas that are more distorted on the London tube map, have greater difficulty in finding their preferred route? And do they learn more from the strike as a result? To answer this question, we estimated the following difference-in-differences-in-differences regression:

$$d_{it}^{j\text{-mode}} = \alpha_i + \beta \cdot d_t^{\text{post}} + \gamma (d_t^{\text{post}} \cdot d_i^{\text{treat}}) + \zeta (d_t^{\text{post}} \cdot dist_i^j) + \theta (d_t^{\text{post}} \cdot d_i^{\text{treat}} \cdot dist_i^j) + \epsilon_{it}, \quad (2)$$

where “ $dist_i^j$ ” is our measure of map distortion around individual i ’s modal station of entry or exit (with $j \in \{\text{entry, exit}\}$). Note that this exercise again explicitly distinguishes between the station of entry and exit, since map distortions are likely to be different at both ends. Tables 7-9 report our results. In this regression, a negative estimate for θ would suggest that treated commuters who live in (or travel to) more distorted areas, are less likely to use their pre-strike modal journey in the post-strike period. This would provide evidence in favor of the hypothesis that commuters who live in more distorted areas, have greater difficulty in finding their optimal commute. And as can be seen from Tables 7-9, this indeed seems to be the case: our estimate of θ tends to be significantly negative across specifications, thereby pointing towards the importance of informational imperfections in explaining our findings. The estimates of γ tend to remain negative and significant, which suggests that map distortion can not explain the full effect, but the fact that the absolute value of our estimate tends to go down suggests that it does explain part of it.

Table 7: OLS-DiD results when treatment group is identified as individuals deviating from pre-strike mode during strike.

	(1)	(2)
	$d_{it}^{\text{entry_mode}}$	$d_{it}^{\text{exit_mode}}$
d_t^{post}	-0.00440 (0.00364)	-0.00317 (0.00511)
$d_t^{\text{post}} \cdot dist_i^j$	0.00141 (0.0250)	-0.0435 (0.0338)
$d_t^{\text{post}} \cdot d_i^{\text{treat}}$	-0.0152*** (0.00478)	-0.0407*** (0.00661)
$d_t^{\text{post}} \cdot d_i^{\text{treat}} \cdot dist_i^j$	-0.0675** (0.0327)	-0.0263 (0.0438)
obs	267,588	267,588

Table 8: OLS-DiD results when treatment group is identified as individuals traveling to or from affected stations pre-strike.

	(1: entry on strike)	(2: exit on strike)	(3: either)	(4: either)
	$d_{it}^{\text{entry_mode}}$	$d_{it}^{\text{exit_mode}}$	$d_{it}^{\text{entry_mode}}$	$d_{it}^{\text{exit_mode}}$
d_t^{post}	-0.0196*** (0.00272)	-0.0186*** (0.00612)	-0.0210*** (0.00359)	-0.0350*** (0.00510)
$d_t^{\text{post}} \cdot \text{dist}_i^j$	0.00306 (0.0207)	-0.115*** (0.0431)	0.0103 (0.0254)	-0.00363 (0.0350)
$d_t^{\text{post}} \cdot d_i^{\text{treat}}$	0.0110** (0.00482)	0.00335 (0.00838)	0.00927* (0.00494)	0.00278 (0.00677)
$d_t^{\text{post}} \cdot d_i^{\text{treat}} \cdot \text{dist}_i^j$	-0.160*** (0.0376)	-0.142** (0.0612)	-0.0996*** (0.0340)	-0.0971** (0.0455)
obs	226,404	184,482	267,588	267,588

Table 9: OLS-DiD results when treatment group is identified by travel time.

	(1)	(2)	(3)	(4)	(5)	(6)
	$d_{it}^{\text{entry_mode}}$	$d_{it}^{\text{exit_mode}}$	$d_{it}^{\text{entry_mode}}$	$d_{it}^{\text{exit_mode}}$	$d_{it}^{\text{entry_mode}}$	$d_{it}^{\text{exit_mode}}$
d_t^{post}	-0.0101*** (0.00342)	-0.0230*** (0.00479)	-0.0158*** (0.00278)	-0.0265*** (0.00382)	-0.0177*** (0.00256)	-0.0317*** (0.00351)
$d_t^{\text{post}} \cdot \text{dist}_i^j$	-0.0428* (0.0235)	-0.0782** (0.0320)	-0.0323* (0.0189)	-0.0820*** (0.0254)	-0.0316* (0.0174)	-0.0712*** (0.0233)
$d_t^{\text{post}} \cdot d_i^{\text{treat}}$	-0.0103** (0.00490)	-0.0183*** (0.00668)	0.000678 (0.00597)	-0.0257*** (0.00793)	0.0224** (0.00942)	-0.0124 (0.0188)
$d_t^{\text{post}} \cdot d_i^{\text{treat}} \cdot \text{dist}_i^j$	-0.0111 (0.0333)	0.0207 (0.0444)	-0.0712* (0.0402)	0.0540 (0.0525)	-0.208*** (0.0633)	0.0109 (0.0784)
obs	267,588	267,588	267,588	267,588	267,588	267,588
factor	1.2	1.2	1.5	1.5	2	2

7.4.2 Line speed

Even if the London underground network were to adopt an undistorted tube map, this still would not solve all informational problems. The reason is that many characteristics of various lines, such as crowdedness, nature of the follow-up journey to work, etc. remain

unknown until that line is actually tried. One such characteristic that is easily quantified, is line speed. As shown in Table 10, speed differs considerably across lines.²⁹ Consequently, two journeys that look equally far on an undistorted map, are still not equivalent if they are made in trains that travel at different speeds.

Table 11 therefore reports results that were obtained after estimating the following difference-in-difference-in-differences regression:

$$d_{it}^{\text{mode}} = \alpha_i + \beta \cdot d_t^{\text{post}} + \gamma (d_t^{\text{post}} \cdot d_i^{\text{treat}}) + \zeta (d_t^{\text{post}} \cdot \text{speed}_i) + \theta (d_t^{\text{post}} \cdot d_i^{\text{treat}} \cdot \text{speed}_i) + \epsilon_{it} \quad (3)$$

Since speed varies across lines, we now limit ourselves to the sample of commuters who stay on the same underground line for their entire commute (the same sample that was used in Column 2 of Table 5). Consequently, our speed-variable becomes individual i -specific. The “same line”-restriction reduces sample size, as a result of which our estimates become less significant (like in Table 5).

Table 10: Average speed of trains along various tube lines

Line	Speed (km/h)
Bakerloo	27.84
Central	40.12
Circle	22.82
District	29.71
Hammersmith & City	15.04
Jubilee	36.61
Metropolitan	32.92
Northern	32.28
Piccadilly	32.69
Victoria	39.38
Waterloo & City	47.40

However, if anything, our exercise suggests that treated individuals are more likely to change their journey in the post-strike period if they were commuting on a relatively slow line before the strike. Because this regression includes speed, which is inversely

²⁹This table draws upon own calculations (based upon TfL-information) and contains the average speed attained by the various trains in between stations. Consequently, our measure is not distorted by the density of stations on a particular line, which is a characteristic that is easily observed from the tube map.

related to slowness, a *positive* estimate for θ now provides evidence in favor of the idea that switchers move away from slower lines. The reason seems to be that the episode of forced experimentation during the strike makes slow-line commuters aware of the fact that their usual train is rather slow-paced, which induces them to reconsider their options post-strike. This is again consistent with the hypothesis that informational imperfections drive our main results.

Table 11: OLS-DiD results when interacting with line speed.

	(1: not on mode)	(2: mode on strike)	(3: factor 1.2)	(4: factor 1.5)	(5: factor 2)
	d_{it}^{mode}	d_{it}^{mode}	d_{it}^{mode}	d_{it}^{mode}	d_{it}^{mode}
d_t^{post}	-0.0133 (0.0286)	-0.0448 (0.0449)	-0.0250 (0.0335)	-0.0821*** (0.0235)	-0.0894*** (0.0210)
$d_t^{\text{post}} \cdot \text{speed}_i$	0.0196 (0.0520)	0.0476 (0.0786)	-0.00320 (0.0601)	0.0977** (0.0418)	0.106*** (0.0371)
$d_t^{\text{post}} \cdot d_i^{\text{treat}}$	-0.163*** (0.0392)	-0.0645 (0.0502)	-0.129*** (0.0423)	-0.0874* (0.0463)	-0.168** (0.0773)
$d_t^{\text{post}} \cdot d_i^{\text{treat}} \cdot \text{speed}_i$	0.210*** (0.0695)	0.0797 (0.0882)	0.207*** (0.0750)	0.112 (0.0807)	0.233* (0.134)
obs	47,052	47,052	47,052	47,052	47,052

8 Interpretation

Our paper has presented evidence that a significant fraction of commuters in our dataset failed to optimize their journey due to the existence of informational imperfections. As a result, a disruption was able to bring about lasting changes in behavior and associated time-gains. How should we interpret this result? Broadly speaking, there are two competing hypotheses that could explain our findings.

Under Hypothesis I, agents in our dataset were acting rationally and followed the optimal search rule, but due to the presence of search costs they (rationally) aborted their exploration for the best alternative before they had found their global maximum. Along these lines, Aghion *et al.* (1991) formally show that following the optimal search strategy does not necessarily imply that the global maximum will be found. Using the language of Baumol and Quandt (1964), Hypothesis I implies that although agents were not maximizing, they were optimizing (i.e.: behaving optimally given the existence of search costs).

Under Hypothesis II, on the other hand, agents were not adhering to the optimal search rule and experimented too little relative to the prescription of the standard-rational model.³⁰ Using Baumol-Quandt terminology, this hypothesis implies that agents were neither maximizing nor optimizing; instead, this hypothesis implies that agents were “satisficing” in a way that is harder to rationalize (as in Simon (1955)).

To investigate which of these two hypotheses is in the best position to explain our results, it is useful to see what the optimal search strategy looks like for this problem (taking into account that search is costly; hence this strategy can be described as “optimizing” in the language of Baumol-Quandt). The optimal strategy for such an environment has been characterized independently by Gittins (1979) and Weitzman (1979). Using Weitzman’s formulation and notation, the optimal strategy is to continue trying new alternatives until:

$$c_i = e^{-rt_i} \int_z^\infty (x_i - z) dF_i(x_i) - (1 - e^{-rt_i})z, \quad (4)$$

where c_i is the cost of trying a new alternative i , r is the discount rate, and t_i is the time lag at which the value of a new alternative is learned (when learning is instantaneous upon trying a new alternative, $t_i = 0$). The parameter z is the present discounted value of the alternative that is currently chosen, while x_i represents the present discounted value of the most attractive unexplored alternative i . This value is distributed according to a c.d.f. $F_i(\cdot)$.

Given that the value of an alternative route is learned soon (if not immediately) after trying it, it seems reasonable to advance with $t_i = 0$ (such that $e^{-rt_i} = 1$). Equation (4) then simplifies to:

$$c_i = \int_z^\infty (x_i - z) dF_i(x_i) \quad (5)$$

Based upon our findings in Section 7.3, we approximate the average daily welfare-gain $\int_{z^{daily}}^\infty (x_i^{daily} - z^{daily}) dF_i(x_i)$ realized by commuters who were forced to experiment because of the strike by setting it equal to the monetary equivalent of 40 seconds per day (twice

³⁰These two competing hypothesis can also be found in the debate on the Porter-hypothesis. In their contribution, Jaffe *et al.* (1995: 156) for example write that “one must be careful when claiming that firms are not operating on their production frontiers: if there are managerial costs to investigating new production technologies, then firms may be efficient even if they do not realize that new, more efficient processes exist until regulations necessitate their adoption. In other words, there may be many efficiency-enhancing ideas that firms could implement if they invested the resources required to search for them. If firms do successfully search in a particular area for beneficial ideas, it will appear ex post that they were acting suboptimally by not having investigated this area sooner. But with limited resources, the real question is not whether searching produces new ideas, but whether particular searches that are generated by regulation systematically lead to more or better ideas than searches in which firms would otherwise engage.”

the average time-gain on a one-way commute).³¹ This is a rather conservative number since this time-gain does not capture unmeasured characteristics of the commute (like line-crowdedness), along which the new alternative is possibly preferred over the old one. Using the results of Stutzer and Frey (2008), we calculate the present discounted value of reducing the commute by 40 seconds per day. In this calculation we work with an annual discount rate of 4% and assume that the 40 second-gain lasts for a period of 4 years (which seems a reasonable number given that average job tenure in the UK is 9 years, while the average time that London households live in their home is 11.8 years³² – and note that these two changes often coincide). Given that the average net yearly income in London is about £26,176,³³ and taking commute-free weekends into account when calculating the present-discounted value, this implies that $\int_z^\infty (x_i - z)dF_i(x_i) \approx \text{£}380$.

If commuters were adhering to the optimal search strategy (prescribed by equation (5)), this implies that the cost of trying the most attractive untried alternative would have to be greater than £380. Or stated otherwise: under the assumption that our data were generated by optimizing searchers, one would have had to offer a commuter more than £380 in order to induce him to try the most attractive untried alternative *for just one day* (after which he is free to go back to his status quo again). This strikes us as implausibly high and suggests that agents underestimate the value of experimentation (possibly because their subjective beliefs on the distribution F_i are too pessimistic; cf. footnote 31) and experiment too little as a result. This calculation is back-of-the-envelope which is why we use a rather conservative calibration.³⁴

³¹Here, we assume that the subjective c.d.f. F_i coincides with the objective distribution (the one we observe in our data).

³²See http://www.cipd.co.uk/binaries/megatrends_2013-job-turnover-slowed-down.pdf and own calculations based upon <https://www.gov.uk/government/statistics/english-housing-survey-2013-to-2014-household-report>. Our calculations account for the fact that London has a significant rental sector (where turnover tends to be much higher).

³³See http://www.ons.gov.uk/ons/dcp171778_385428.pdf, where it is reported that the average gross yearly income in London is £34,346. According to <http://www.incometaxcalculator.org.uk>, this corresponds to a net income of about £26,176.

³⁴As pointed out before, our calculation misses the unmeasured advantages of the new alternative (like traveling in a quieter environment), while the underlying Stutzer-Frey calculation already makes several conservative choices too (which may contribute to the fact that Ahlfeldt *et al.* (2015) report a bigger number; recall footnote 1). In addition, measured effects are likely to be larger if agents had experimented voluntarily in a more tranquil environment (as opposed to the chaotic strike environment considered in our paper), while our selection strategy (requiring commuters to behave in a very consistent manner pre-strike) stacks the cards against switching as well. Setting the annual discount rate to 4% is a cautious choice too. However, just to illustrate the robustness of our finding, pushing the annual discount rate up from 4% to 80%, would only validate the rational search rule for $c_i \approx \text{£}100$, which still seems implausibly high. Using the quasi-hyperbolic discount function $\{1, \varphi\delta, \varphi\delta^2, \varphi\delta^3, \dots\}$ (where δ is the daily discount factor, which we base upon an annual discount rate of 4%) and setting φ equal to the Laibson, Repetto and Tobacman (2008)-estimate of 0.7, yields $c_i \approx \text{£}270$. Whether risk aversion can

But what can we say about the underground network as a whole? As set out in Section 7.1, only a subset of about 5% of commuters (henceforth: “the beneficiaries”) found a better route to work thanks to the strike. The remaining 95% however, did not make such a discovery: they only suffered from delays on February 5 and 6. Looking at the tube network as a whole, an important question therefore is: has efficiency (in the sense of Kaldor-Hicks) improved thanks to the strike? To make this calculation, we need to compare the costs imposed on all treated commuters during the strike, with the benefits the strike has brought to the subset of beneficiaries. Again abstracting from unmeasured characteristics, we express both costs and benefits in terms of travel time.

As far as costs are concerned, our data indicate that average travel time in the treatment group (as defined by those commuters who deviated from their modal journey during the strike) went up by 4.5 minutes (270 seconds) for a one-way commute on strike days. Again using a 4% annual discount rate, and assuming that the strike taught about 5% of all treated commuters a better route to work (which they can continue to use for 4 years), this implies that the strike improved efficiency if it brought a gain to the subset of beneficiaries of at least 37 seconds per day. Given that the *average* treatment effect (where the average is taken over both beneficiaries ($\approx 5\%$) and non-beneficiaries ($\approx 95\%$)) is already a saving of 40 seconds per day, it seems that the strike has improved efficiency along the lines of Kaldor-Hicks.

Together, these calculations provide evidence for the hypothesis that agents are “satisficing”, rather than maximizing. Moreover, in the language of Baumol and Quandt (1964), agents seem to “satisfice” in a way that is not optimal (i.e.: they do not seem to optimize, even if we take into account that search is costly as in Weitzman (1979)). This suggests that commuters in our dataset were not acting along conventional rational lines. Herewith, our field data seem to support the results reported in the experimental study of Caplin, Dean and Martin (2011), who also present evidence in favor of “satisficing” behavior (although they do not analyze whether the behavior of subjects can be characterized as “optimal” in the Baumol-Quandt sense).

Moreover, our findings suggest that agents in our dataset were experimenting less during tranquil times than prescribed by the standard-rational model. This is consistent with laboratory evidence surveyed and reported in Anderson (2012), but to the best of our knowledge our study is the first to present evidence in favor of this hypothesis based upon detailed field data. Our results furthermore allow for the (highly controversial) idea,

rationalize the puzzle, is not obvious: as shown in Willems (2016), making decision makers risk averse might actually make them *more* eager to experiment (since doing so also produces information on the environment, thereby reducing future risk).

advocated by Porter (1991), that imposing a constraint on an economic system (which forces agents to experiment), can enhance efficiency over time.

Finally, one could think that the shock of the strike moves the system from one solution in a coordination game to another, such that all users of the transport system slowly adapt to a new equilibrium with new conditions. We do not think that this is what we observe. The correlation of relative station use in the morning hours between the first and the last day is 0.997, and for no two non-strike days it is below 0.995. Consequently, the system as a whole looks fairly similar from the perspective of one individual commuter, and her decision is a partial one.

9 Conclusion

We have presented evidence which suggests that a significant fraction of commuters in London failed to find their optimal route to work. This failure seems to be driven by informational imperfections. We have furthermore shown that the observed behavior is unlikely to be rationalized by search costs: for that, the observed deviations from optimality are too large and too prevalent. Instead, our results seem to provide the first field data-based evidence in favor of the hypothesis that agents are “satisficing” (in a non-optimizing way) and that they underestimate the value of experimentation. As a result, they experiment less than what is prescribed by the standard-rational model which contributes to their failure in finding the optimum.

Because of the “satisficing” nature of decision-making by agents in our dataset, an exogenously-imposed constraint (the tube strike of February 2014) was able to bring about lasting changes in behavior among a significant fraction of commuters. So, as an unintended consequence of transportation strikes, people may find better commuting routes. The time gains subsequently achieved by this group, seem to outweigh the time-losses incurred by all commuters during the strike. It therefore appears that the tube network was operating so far away from its optimum, that the February 2014-strike managed to improve efficiency of the system as a whole in the sense of Kaldor-Hicks.

Despite the fact that a substantial share of travelers are likely to have received help from online journey planners, from previous disruptions to the network (calling for earlier experimentation), as well as from the experiences of others, it seems that they were still not maximizing. Given that the challenges faced by businesses are arguably more complex than the commuter-problem analyzed in this paper,³⁵ it seems a real possibility

³⁵After all, an important part of the answer to this question (travel time) can just be found at jour-

that many firms are not operating efficiently either. Consequently, the Porter-hypothesis (which states that the imposition of constraints can bring about efficiency-enhancing dynamic effects by triggering a period of experimentation and re-optimization) might be less implausible than its critics, such as Palmer, Oates and Portney (1995) and Schmalensee (1993), have argued. In the context of the London Underground, this implies that commuters could be made better off if given an external encouragement to experiment. Since partial closure of the network is a rather radical way to achieve this, it is worth investigating whether clever use of journey planner apps can “nudge” travelers to experiment more.

Real-life examples of behavior that is similar to that of commuters in our dataset abound:

- It was only because of an exogenous conflict with France that the British discovered port: at the beginning of the 18th century, the Royal Navy blocked French harbors – thereby stopping the export of French wines to Britain. This left British consumers in search for an alternative, which is how they came across (and fell in love with) port.³⁶
- In the 1960s, an ambitious high-jumper from Portland, Oregon faced a serious constraint – namely a lack of talent. As this deficiency became more apparent over the years, he saw himself forced to experiment with a new technique. This soon enabled him to improve his personal best by half a foot in one day and eventually made him win the Olympic gold medal in the 1968 Olympics. We are of course talking about Dick Fosbury and his “flop” is still considered to be one of the most significant innovations in sports (Hoffer, 2009).
- In August 2015, a police strike in The Netherlands implied that they were not able to supervise fans around matches in the Dutch professional soccer league. Some matches were cancelled, but others went ahead nevertheless. To the surprise of many, the matches that went ahead were completed peacefully. This taught authorities that a police-presence around these events is not always necessary – thereby opening the door to substantial future cost-savings.³⁷

neyplanner.tfl.gov.uk/. To the best of our knowledge, no such website exists for many of the everyday problems that businesses are faced with.

³⁶See <http://www.theguardian.com/lifeandstyle/2010/dec/30/port-wine-food-and-drink>.

³⁷See <http://www.dutchnews.nl/news/archives/2015/08/police-union-welcomes-trouble-free-football-this-weekend/>.

From all this, one gets the impression that decision-making is difficult in a world where information is imperfect. In addition, our findings illustrate that people might get stuck with suboptimal decisions because of “satisficing behavior” and under-experimentation. As a result, the imposition of constraints can improve long-run efficiency, while our results also highlight the importance of implementing occasional routine breaks to explore efficiency at the margins.

10 References

Aghion, Philippe, Patrick Bolton, Christopher Harris and Bruno Jullien (1991), “Optimal Learning by Experimentation”, *Review of Economic Studies*, 58 (4): 621-654.

Aghion, Philippe, Mathias Dewatripont and Patrick Rey (1997), “Corporate Governance, Competition Policy and Industrial Policy”, *European Economic Review*, 41 (3-5): 797–805.

Ahlfeldt, Gabriel M., Stephen J. Redding, Daniel M. Sturm, and Nikolaus Wolf (2015), “The Economics of Density: Evidence from the Berlin Wall”, *Econometrica*, 83 (6): 2127-2189.

Ambec, Stefan, Mark A. Cohen, Stewart Elgie and Paul Lanoie (2013), “The Porter Hypothesis at 20: Can Environmental Regulation Enhance Innovation and Competitiveness?”, *Review of Environmental Economics and Policy*, 7 (1): 2-22.

Anderson, Christopher M. (2012), “Ambiguity Aversion in Multi-Armed Bandit Problems”, *Theory and Decision*, 72 (1): 15-33.

Baumol, William J. and Richard E. Quandt (1964), “Rules of Thumb and Optimally Imperfect Decisions”, *American Economic Review*, 54 (2): 23-46.

Berman, Eli and Linda T.M. Bui (2001), “Environmental Regulation and Productivity: Evidence from Oil Refineries”, *Review of Economics and Statistics*, 83 (3): 498-510.

Bertrand, Marianne, Esther Duflo and Sendhil Mullainathan (2004), “How Much Should We Trust Differences-in-Differences Estimates?”, *Quarterly Journal of Economics*, 119 (1): 249-275.

Bolton, Patrick and Christopher Harris (1999), “Strategic Experimentation”, *Econometrica*, 67 (2): 349-374.

Caplin, Andrew, Mark Dean, and Daniel Martin (2011), “Search and Satisficing”, *American Economic Review*, 101 (7): 2899-2922.

Copeland, Brian R. and M. Scott Taylor (2004), “Trade, Growth, and the Environment”, *Journal of Economic Literature*, 42 (1): 7-71.

Gabaix, X., D. Laibson, G. Moloche and S. Weinberg (2006), "Costly Information Acquisition: Experimental Analysis of a Boundedly Rational Model", *American Economic Review*, 96 (4): pp. 1043-1068.

Gittins, John C. (1979), "Bandit Processes and Dynamic Allocation Indices", *Journal of the Royal Statistical Society*, 41 (2): 48-177.

Goodwin, Phil B. (1977), "Habit and Hysteresis in Mode Choice", *Urban Studies*, 14 (1): 95-98.

Gray, Wayne B. (1987), "The Cost of Regulation: OSHA, EPA, and the Productivity Slowdown", *American Economic Review*, 77 (5): 998-1006.

Guo, Zhan (2011), "Mind the Map! The Impact of Transit Maps on Path Choice in Public Transit", *Transportation Research Part A*, 45 (7): 625-639.

Hicks, John R. (1932), *The Theory of Wages*, London: Macmillan.

Hodges, Andrew (2014), *Alan Turing: The Enigma*, London: Random House.

Hoffer, Richard (2009), *Something in the Air*, New York, NY: Free Press.

Jaffe, Adam B., Steven R. Peterson, Paul R. Portney and Robert N. Stavins (1995), "Environmental Regulation and the Competitiveness of U.S. Manufacturing: What Does the Evidence Tell Us?", *Journal of Economic Literature*, 33 (1): 132-163.

Jaffe, Adam B. and Karen Palmer (1997), "Environmental Regulation and Innovation: A Panel Data Study", *Review of Economics and Statistics*, 79 (4): 610-619.

Laibson, David, Andrea Repetto and Jeremy Tobacman (2008), "Estimating Discount Functions with Consumption Choices over the Lifecycle", NBER Working Paper No. 13314.

Matejka, Filip and Alisdair McKay (2015), "Rational Inattention to Discrete Choices: A New Foundation for the Multinomial Logit Model", *American Economic Review*, 105 (1): 272-298.

Morgan, Peter and Richard Manning (1985), "Optimal Search", *Econometrica*, 53 (4): 923-944.

Palmer, Karen, Wallace E. Oates and Paul R. Portney (1995), "Tightening Environmental Standards: The Benefit-Cost or the No-Cost Paradigm?", *Journal of Economic Perspectives*, 9 (4): 119-132.

Porter, Michael E. (1991), "America's Green Strategy", *Scientific American*, 264 (4): 168.

Porter, Michael E. and Claas van der Linde (1995), "Toward a New Conception of the Environment-Competitiveness Relationship", *Journal of Economic Perspectives*, 9 (4): 97-118.

- Roberts, Kevin and Martin L. Weitzman (1981), "Funding Criteria for Research, Development, and Exploration Projects", *Econometrica*, 49 (5): 1261-1288.
- Rothschild, Michael (1974a), "Searching for the Lowest Price When the Distribution of Prices Is Unknown", *Journal of Political Economy*, 82 (4): 689-711.
- Rothschild, Michael (1974b), "A Two-armed Bandit Theory of Market Pricing", *Journal of Economic Theory*, 9 (2): 185-202.
- Simon, Herbert A. (1955), "A Behavioral Model of Rational Choice", *Quarterly Journal of Economics*, 69 (1): 99-118.
- Sims, Christopher A. (2003), "Implications of Rational Inattention", *Journal of Monetary Economics*, 50 (3): 665-690.
- Schmalensee, Richard (1993), "The Costs of Environmental Protection", mimeo, MIT.
- Stutzer, Alois and Bruno S. Frey (2008), "Stress that Doesn't Pay: The Commuting Paradox", *Scandinavian Journal of Economics*, 110 (2): 339-366.
- Van Exel, N. Job. A and Piet Rietveld (2001), "Public Transport Strikes and Traveller Behaviour", *Transport Policy*, 8 (4): 237-246.
- Weitzman, Martin L. (1979), "Optimal Search for the Best Alternative", *Econometrica*, 47 (3): 641-654.
- Willems, Tim (2016), "Optimal Experimentation under Risk Aversion", mimeo.
- Wood, Wendy and David T. Neal (2009), "The Habitual Consumer", *Journal of Consumer Psychology*, 19 (4): 579-592.
- Zhu, Shanjiang and David M. Levinson (2011), "Disruptions to Transportation Networks: A Review", in: David M. Levinson, Henry X. Liu and Michael G.H. Bell (eds.), *Network Reliability in Practice*, New York: Springer.

11 Figures

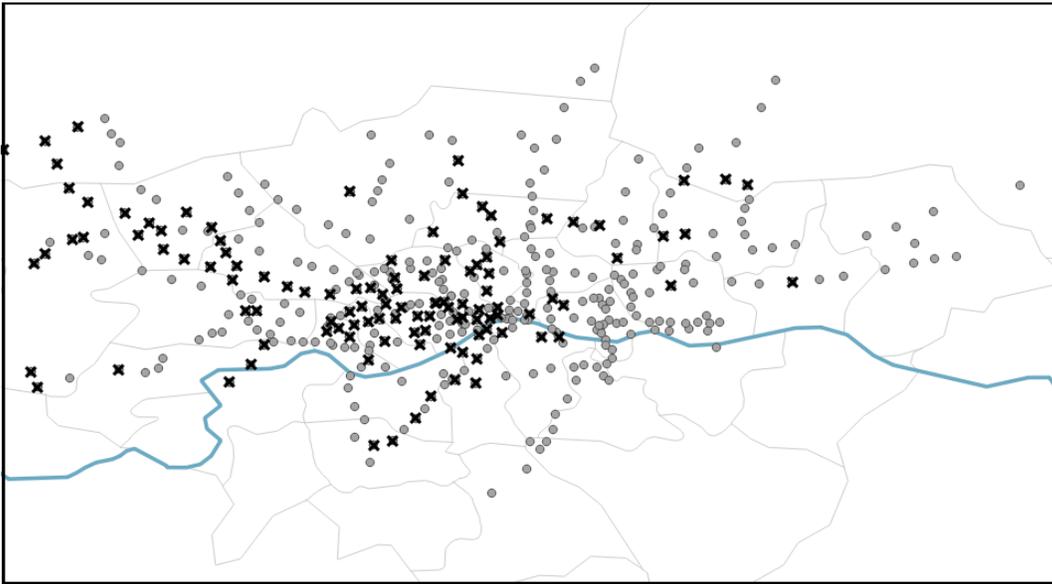
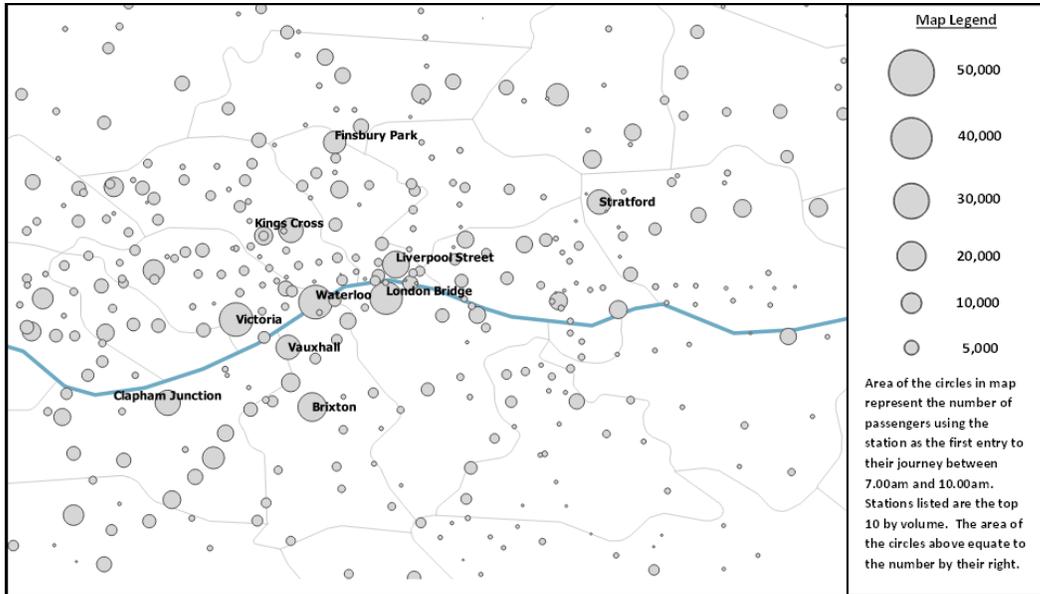
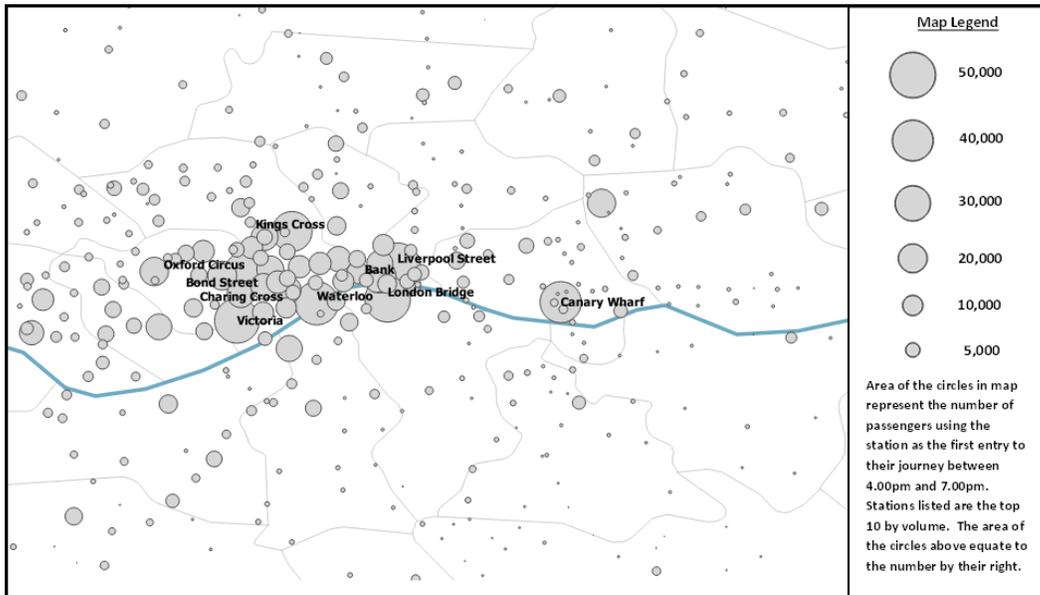


Figure 1: Impact of the February 2014 tube strike. Circles represent stations on a standard tube map (includes Overground and DLR) with GPS coordinates used to locate position. Crosses represent stations that were fully closed during the strike period.



(a) Morning (7am-10am)



(b) Evening (4pm-7pm)

Figure 2: Stations of first entry in the morning and evening of January 31, 2014.

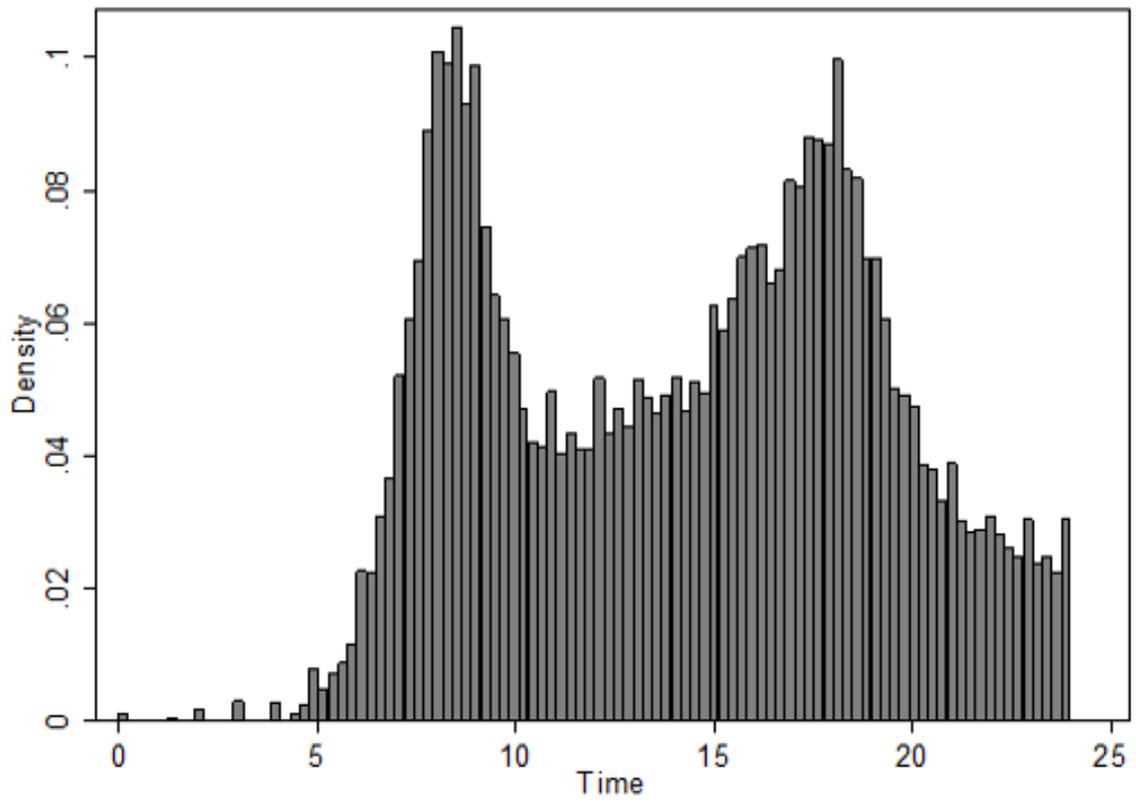


Figure 3: Travel pattern of January 31, 2014.

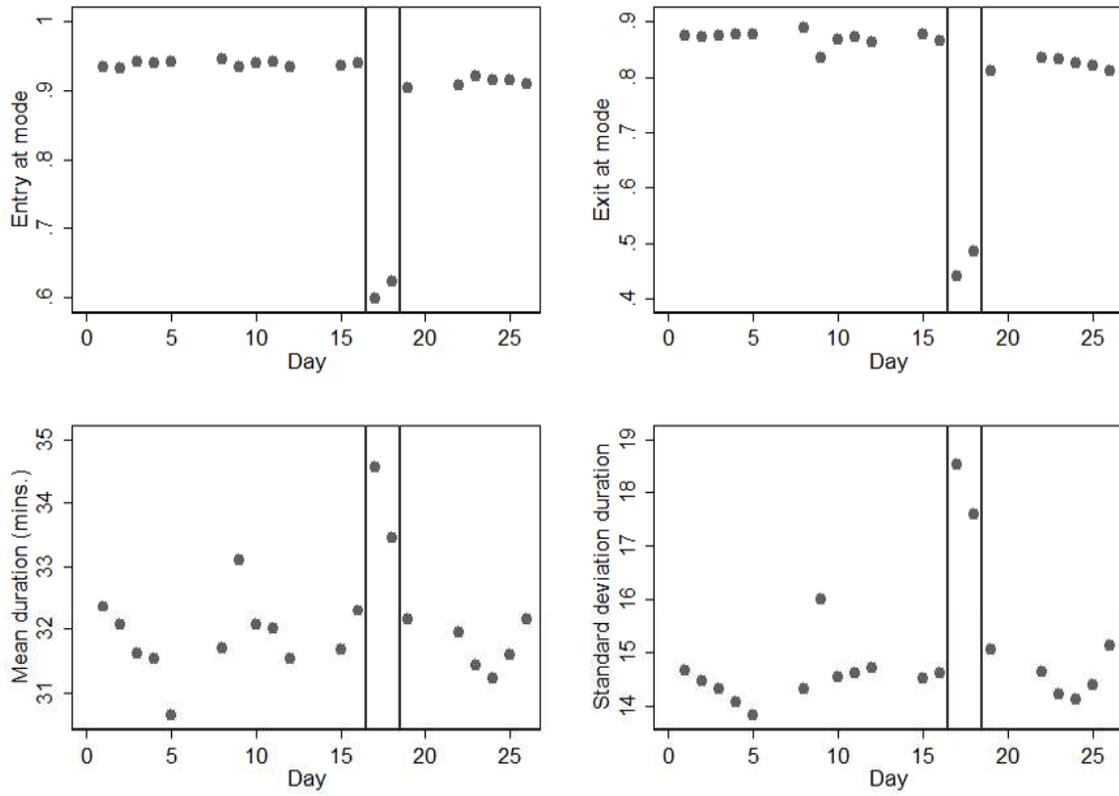


Figure 4: Summary statistics.

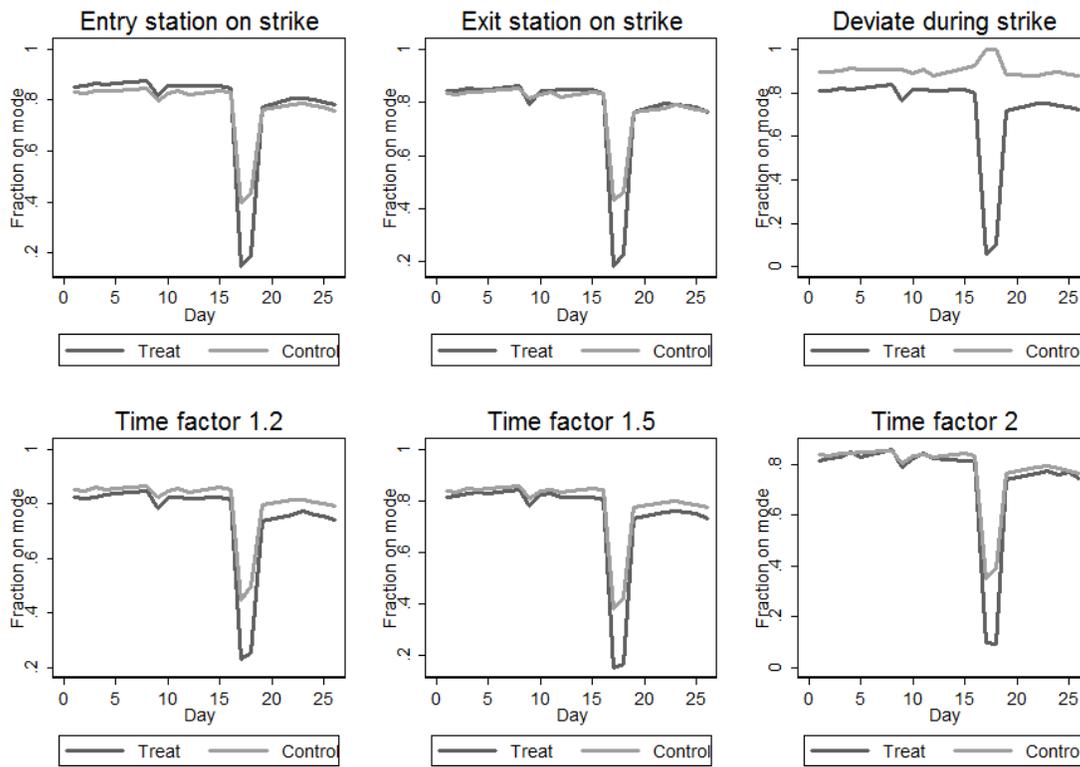


Figure 5: Time trends for treatment and control group.