Supporting Information

Silva et al. 10.1073/pnas.1412908112

SI Text

The Data. The public transportation system in Greater London consists mainly of an interconnected network of railways and buses. In 2011, public transportation accounted for about 19 billion passenger kilometers representing 43% of total journey stages, compared with 34% of journey stages by private transportation, 21% walking, and 2% cycling (1). The main railway networks comprise the London Underground network (also called "the Tube"), the Overground rail network, and the DLR.* Fig. S1 shows the map of these railway systems. The London Underground railway dates back to 1863, when it opened serving eight stations between Paddington and Farringdon. Since then, it has grown to 11 lines totaling 402 kilometers of extension serving 270 stations. The London Overground network covers 83 stations on six lines to provide connections between areas outside of Central London. The DLR is an automated light rail system with 45 stations in seven lines, which opened in 1987 to serve the redeveloped Docklands area of London. In 2011, the volume of passenger traffic was 9.5 billion passenger kilometers in Underground, 645 million passenger kilometers in Overground, and 456 million passenger kilometers in DLR (1). The local government body responsible for operation, management, and planning of transportation in London is TfL.

In 2003, TfL introduced an automated fare collection system called the "Oyster" card. This radiofrequency identification-based smart card is a form of electronic ticket used on public transportation system within the London fare zones. Users touch the card on an electronic reader when entering the transport system (a "tap-in") and leaving the system (a "tap-out") to deduct the fare. The Oyster card records the time and place of these transactions. The use of these smart cards is encouraged by the offering of substantially cheaper fares compared with those in cash. By June 2012, TfL issued more than 43 million Oyster cards and their use accounted for more than 80% of all public transport journeys in London (2).

For this study, we obtained from TfL the data containing each single journey taken using an Oyster card for a subset of days between February 14, 2011 and February 9, 2012, a total of 70 weekdays and 25 weekend days. Each Oyster card record (a "tap") contains the date of the transaction, anonymized user ID, and the event time (measured in minutes) and event code. Event codes represent transactions such as adding more credit to the card or cancelling previous transactions. In this study, we consider the events of entering or exiting an Underground, Overground, or DLR stations only. Each tap-out record includes the location of the first entry in the journey, and the corresponding entry time. For the time period covered by our data, the combined system of Underground, Overground, and DLR consisted of 374 stations.† From the raw data, we excluded cases where journeys were not completed, that is, tap-ins without a corresponding tap-out. This was done by considering only records of tap-out events, using the corresponding fields of time and location of first entry as the information about tap-ins. The missing tap-out events with a recorded tap-in can be safely treated as missing from a population that is not of interest to this study, because they are typically due to travelers who attempt to use the card in an invalid way (as in the case where there is no minimal sufficient credit while tapping-in) or due to TfL staff members. After these exclusions, we obtained 210,764,572 journeys by 10,687,141 unique users, with an average of 1.71 journeys per user per day, 1,756,756 unique IDs per day, and 3,010,922 journeys per day. For the analyses reported in the main text, we restrict our study to weekdays only, because the pattern of traffic between weekdays and weekends are different and the effects of disruptions on weekdays are more relevant. We discuss the differences between weekdays and weekend days in *SI Text*, *Assumptions About Weekdays*.

Lines of the London Tube and the Tube Graph. The lines for the Underground and Overground were segmented according to the official TfL classification, whereas the DLR was treated as a single line. Stations that lie at two or more lines were treated as single stations (for instance, King's Cross), with the exception of a few stations that have different identifier codes in the Oyster card system (for instance, Edgware Road at Hammersmith is distinguished from Edgware Road in the Circle Line). A full list of stations and lines is available upon request. We call the undirected graph formed by the conjunction of 374 stations and lines the "Tube graph," where each station is a vertex, with an edge between each pair of stations physically joined by tracks.

Assumptions About Weekdays. Weekdays are not strictly exchangeable. Standard rank tests (we used kstest2 from MATLAB) reject the null hypothesis of pairwise exchangeability at a 0.01 level for any two weekdays using a "bag of minutes" representation, where counts for every minute and every station are pooled together for each particular day of the week, $1{,}200 \times 374$ data points in each day. For instance, in the raw data there is an increase on the overall number of journeys in late Friday evenings compared with the rest of the week.

We nevertheless adopt the assumption of exchangeability of weekdays (and exclude weekends from our analysis) for simplicity. A quick visual inspection of the histograms of the bag of minutes representation reveals no evident visual features separating weekdays during the busiest times (9:00 AM–6:00 PM) and this assumption strongly simplifies the analysis. The difference between weekdays and weekends, however, is very strong. Fig. S2 depicts a summary of our data, illustrating what we claim to be three clusters of days: weekdays, Saturday, and Sunday.

The Model for the Natural Regime. Passengers enter and exit the system at different time points, where time is an integer from $t = 1$ (5:00 AM) to $t = 1{,}200$ (1:00 AM of the next day). Each value of t corresponds to a minute. Each day consists of $|V|^2 \times \mathcal{T}$ outcomes; $|V|$ is the number of stations and $\mathcal{T}$ is the number of time points, 374 and 1,200, respectively. If dependence across time is to be modeled at all, assumptions have to be chosen carefully on how to decouple such dependencies in a model that is easy to interpret and that can be fitted without much computational burden.

Patterns at different OD pairs are expected to show degrees of similarity, particularly in geographically close places. However, we should differentiate trend similarities from stochastic dependencies. Trend similarity states that model parameters for different stations or OD pairs over time should show some proximity according to suitable similarity metrics. Stochastic dependencies are probabilistic associations among random variables modeling different locations. If we postulate that (i) each passenger decides

*In addition to these main railway systems, there is a tram system (called Tramlink) of 28 km with 39 stops serving the south London boroughs of Croydon and Merton. The traffic generated by Tramlink in 2011 was only 148 million passenger kilometers. We do not include Tramlink in our study.

†The Blackfriars London Underground station and the DLR stations of Stratford, Bank, and Canning Town were mostly closed during this period.

independently when and where to leave from, and where to arrive at, and (*ii*) there are no physical constraints on the journey, then all OD counts will be stochastically independent given time of the day. The first assumption is approximately true when most passengers do not enter the system in groups and entrance times are not jointly affected by stochastic factors (such as delays on a bus bringing two passengers planning to enter the same station). The second assumption is false, taking into account that passengers cannot move independently within the system, and that variability on train arrivals will associate the times that different passengers take to arrive at their destinations. Moreover, given the entrance times of passengers, their exiting times are random but with a degree of predictability.

Our approach is to propagate stochastic dependencies and model exit rates through three sets of processes. The first set describes how many passengers enter a station to start a journey at particular times (entrance processes). The second set describes for how long passengers remain inside the system according to their origin and time of entrance (negotiation processes). Stochastic dependence over time is accounted by these two sets. The third set of processes contains $\mathcal{O}(|V|^2)$ models for OD counts conditioned on the state of the negotiation processes, but without any conditional stochastic dependence across time (exiting processes). We ignore trend similarity in our fitting process, that is, no regularization is used to penalize differences across models for different stations or OD pairs. This estimation procedure is justified by the large amount of data available and the computational cost of methods such as probabilistic matrix factorization (3). We describe the three sets of processes as follows.

The model for the entrance processes. Let L_{it} be the number of passengers entering station S_i at time t. We define $L_{it} \equiv 0$ for $t < 1$ and model the expected values of L_{it} for $t \geq 1$ given the entire past as

$$\mathsf{E}[L_{it}|\mathsf{PAST}, L_{it} > 0] = \left(\theta_{L_{it}} + \sum_{w=1}^{W} \beta_{L_{wi}} L_{i(t-w)}\right)_+, \qquad \textbf{[S1]}$$

where $(x)_+$ means $\max(x, 0)$, and

$$\mathcal{P}(L_{it} > 0|\mathsf{PAST}) \equiv \pi_{L_{it}}. \qquad \textbf{[S2]}$$

Parameter $\theta_{L_{it}}$ represents an unconditional time-dependent mean and explains most of the variation of the data. That is, given $\theta_{L_{it}}$, present behavior is mostly independent of past behavior. Assuming passengers arrive independently given a particular time of the day and location, there should be no dependence on the past at all. Our AR component captures weaker stochastic associations (e.g., people arriving in batches from buses before entering the station), but the AR coefficients $\{\beta_{L_{wi}}\}$ have little impact.

The model does not account for closed stations or unusual behavior. Standard autoregressive models (that is, with no $\theta_{L_{it}}$) are sensitive to station closures with some delay, as illustrated in *SI Text*, *Forecasting, the AR Model, and Sensitivity Analysis for the Natural State Predictive Performance*, but they provide no basis for prediction under a disrupted regime nor are they good models for the natural regime. Here, we are interested in modeling standard behavior, because we do not aim at predicting when external shocks happen (not possible with the data we have), but at deriving what happens when an external shock affects the system.

Parameters $\theta_{L_{it}}$, $t = 1, 2, \ldots, 1{,}200$ are fitted using cubic spline smoothing. We regress L_{it} on t, where $L_{it} > 0$. MATLAB function csaps [version 8.0.0.783 (R2012b)] is used with default parameters, which automatically selects the degree of smoothness. To avoid problems with extrapolation, we set $\theta_{L_{it}}$ to zero for $t \in [0, k_{bottom}] \cup [k_{upper}, 1200]$. Position k_{bottom}, typically taking place in the first few minutes of the day, is defined as the last time point before some tap-in happened in any day in our data. Position k_{upper} is the first time point, typically taking place in the last few minutes of the day, such that no tap-in happened afterward in any day of our data.

Parameters $\pi_{L_{it}}$, $1 \leq t \leq 1{,}200$, are fitted using decision trees by classifying Bernoulli variables $\{L_{it} > 0\}$ on t using R function rpart (version 4.1-8).

These estimates are then plugged into a constrained least-squares regression problem to fit the set $\{\beta_{L_{wi}}\}$ to the residuals $L_{it} - \hat{\theta}_{L_{it}}$, subject to positivity on the expected value of each training point. We use pcls, a function in the mgcv package (version 1.8-2) for nonparametric generalized additive models in R (4).

The fitted model is dominated by the parameters $\theta_{L_{it}}$, with fitted AR coefficients $\{\beta_{L_{wi}}\}$ having little impact, as expected.

The model for the negotiation processes. For each station S_i and time t, we have a (compressed) representation of the number of passengers who entered the system via vertex S_i and have not left the system by time $t-1$. This representation is called a presence table. The presence table is the empirical distribution of such passengers by the amount of time they have been inside in the system. This empirical distribution is given in seven coarsened time brackets of [1,10] min, [11,20] min, …, [50,60] min, and more than 60 min. For each station S_i and time t, we have the vector $\mathbf{M}_{it} \equiv (M_{it}^1, \ldots, M_{it}^7)$ representing counts at these seven levels. For example, M_{it}^2 represents the number of passengers inside the system at time $t-1$, who have started at station S_i, and who have entered the system during the interval $[t-20, t-11]$.

The temporal evolution of the entries of $\mathbf{M}_{it}$ is modeled through a cascade of nonparametric binomial regression models. We assume that, given time t, variation on M_{it}^k $(1 \leq k \leq 7)$ depends only on $M_{i,t-1}^k$ and $M_{i,t-1}^{k-1}$, being conditionally independent from the more distant past. The model for M_{it}^k for any given day is

$$M_{it}^k|\mathsf{PAST} \sim \mathsf{Binomial}\left(M_{i,t-1}^k + M_{i,t-1}^{k-1}, p_{it}^k\right), \qquad \textbf{[S3]}$$

where we define $M_{i,t-1}^0 \equiv L_{i,t-1}$. Parameter p_{it}^k is a different parameter for each station S_i and time of the day t. The model reflects the fact that whoever stays in the kth time bracket came from either the previous cohort of people in the same bracket k, or from bracket $k-1$ (which for $k=0$ corresponds to those who entered the system the minute before). For each station and bracket k, we fit the 1,200 $\{p_{it}^k\}$ parameters with the mgcv package for nonparametric binomial regression using t as the covariate. Although one should expect very weak dependence[‡] between entrance counts L_{it} and $L_{it'}$, we should expect much stronger (marginal) dependence within $\{M_{ik}^k\}$ across time, for the physical reasons explained above.

The model for the exiting processes. Let N_{ijt} be the number of passengers exiting (tapping-out) at station S_j at time t, having started at station S_i. Let R_{ijt} be the sum of the number of passengers in presence table $\mathbf{M}_{it}$, but only for the brackets within 10 min of the median commute time from S_i to S_j, the median estimate being the empirical median of the data. For instance, if the median time is 35 min, the corresponding brackets are {[21,30], [31,40], [41,50]} and $R_{ijt} = M_{it}^3 + M_{it}^4 + M_{it}^5$. The model for N_{ijt} is then

$$E\left[N_{ijt}|\mathsf{PAST}\right] = R_{ijt} \times q_{ijt}. \qquad \textbf{[S4]}$$

Regression is done separately for each of the $|V|^2$ models, by fitting N_{ijt}/R_{ijt} as a nonparametric function of t and using the least-squares cost function. MATLAB's csaps is used again.

Notice that, in principle, the data could allow for $N_{ijt} > R_{ijt}$, because measurement error or misuse of cards leads to tap-outs

‡Again, we are speaking of probabilistic dependence here: There is strong evidence that the means should vary smoothly over time. However, given the model, the L_{it} counts at different times should be essentially independent.

not being matched to some tap-ins. Nevertheless, in practice, this never happens, because R_{ijt} is usually far larger than N_{ijt}. R_{ijt} should therefore be considered a convenient summary of the presence table. Notice also that the dynamic models for the negotiation processes do not explicitly depend on N_{ijt}. The raw observed data for all exits is still used to fit the model for $\mathbf{M}_{it}$. $M_{i,t+1}^k$ is a function of the set $\{L_{it'}\}$ for $t' < t$ and the observable counts $\{N_{idt}^{t'}\}$, the number of passengers leaving the system who were inside the system for t' minutes, and who left the system at time t via exit point S_d. Even though $N_{idt}^{t'}$ is calculated to derive the presence tables to fit the negotiation process, we do not model $N_{idt}^{t'}$ directly, but only via the aggregates N_{idt} and M_{it}^k. Therefore, part of the observable information in the data are lost when compressing it into such models. We nevertheless believe this is a fine enough degree of modeling for our purposes, with $N_{idt}^{t'}$ playing no particular role in the sequel.

A model for N_{ijt} automatically gives a model for $N_{jt} \equiv N_{\cdot jt}$ (number of people exiting station S_j at time t) and $N_{i\cdot t}$ (number of people leaving at time t who started at station S_i). Parameters such as the presence table evolution p_{it}^k (Eq. **S3**) provide information about rate of evasion for passengers who started at S_i, and q_{ijt} (Eq. **S4**) can be used to compare destinations by how they absorb passengers from different origins. Some exploratory analysis can be done by clustering such curves, which we leave for future work.

In the next section, we will also use this model to test the assumption of "clumpiness" in the exiting process: given R_{ijt}, tap-out count N_{ijt} does not follow a binomial process with parameter q_{ijt}. The predictive coverage of N_{ijt} given R_{ijt} is not good using the binomial variance. An explanation is the fact that passengers arrive jointly in trains, so there will be a common source of variability because of this quantization effect on the time of departure. This effect is not negligible because we are looking at individual OD levels. Even if we aggregate OD pairs to predict the overall exit counts for a particular station, we must take this into account. For each destination S_j we also introduce the parameter ϕ_j, which does not vary over time. The parameter regulates the covariance of any pair of Bernoulli variables $\{X_{ijt}^{(m)}, X_{ijt}^{(n)}\}$, where $X_{ijt}^{(m)}$ is the binary indicator that a passenger m is leaving at S_j at time t, having started at S_i. Their correlation is given by $\phi_j \times q_{ijt} \times (1 - q_{ijt})$, with $\phi_j = 0$ in the binomial case. We estimate ϕ_j by a method of moments.

Effect of Traffic Stickiness. Another aspect of the model not captured by the blackbox models is that our model relates entrance behavior to exit behavior. This is reflected by the station-level parameter ϕ_j, which dictates the association level between individual exit events from the network. Our assumption that there is a sizeable level of dependence on how people leave stations can be explained by the fact that people are grouped within trains. One way of testing the hypothesis that $\phi_j > 0$ is by the predictive coverage of a model with our estimated $\{\phi_j\}$ parameters.

At any time point t, conditional on the presence table variables $\{M_{it}^k\}$, we generate predictive confidence intervals of three different magnitudes (90, 95, and 99%) and compare them against the intervals under the assumption $\{\phi_j = 0\}$. We generated aggregated confidence intervals for each N_{jt} by summing the means and variances of the predicted N_{ijt}, which are all independent in the model across stations and time once we conditioned on $\{M_{it}^k\}$. We then use a normal approximation to define the (90, 95, and 99%) predictive confidence intervals.

Let $X_{ijt}^{(m)}$ be the binary event of a particular passenger m leaving station S_j at time t, given the passenger is counted as being in the presence table summary R_{ijt} for N_{jt}. Referring to the model for N_{ijt}, we have

$$E\left[X_{ijt}^{(m)} | R_{ijt}\right] = q_{ijt},$$

where the association among passengers is given by

$$Corr\left(X_{ijt}^{(m)}, X_{ijt}^{(n)} | R_{ijt}\right) = \phi_j.$$

This implies

$$Var(N_{ijt} | R_{ijt}) = R_{ijt} q_{ijt} (1 - q_{ijt}) \left(1 + (R_{ijt} - 1)\phi_j\right). \quad \textbf{[S5]}$$

To estimate ϕ_j, we first estimate q_{ijt} as before. Then, for a fixed S_j, we calculate the empirical variance of the corresponding Bernoulli trials, averaged over all days and time points, and solve the average of Eq. **S5** for ϕ_j. When the estimate is negative (which is possible, because q_{ijt} was estimated separately), we set ϕ_j to zero.

The predictive intervals obtained under the dependent model averaged over the five folds were $(0.87, 0.91, 0.95)$. For the (binomial) model with $\phi_j = 0$, the intervals were severely underdispersed, with a coverage of $(0.66, 0.73, 0.83)$ (all SE under 0.0001). This provides strong evidence for a need to include a dependence structure among the Bernoulli trials, which in our case has physical explanations.

Conditioning on the internal scale of the system (variables $\{M_{it}^k\}$) helps interpretability, because this conditioning allows one to separate variability owing to fluctuations in the entrance numbers at the origin from the degree of dependence between underlying Bernoulli trials of the exit events. Using the physical distance between each station and Oxford Circus as a surrogate to how frequently trains depart a station, we noticed a positive Spearman rank correlation of 0.32 between our estimates of $\{\phi_j\}$ and the physical distance, for our universe of 374 stations.

Fig. S3 shows the average 99% predictive confidence interval for a set of 14 test days independent of 56 d used for fitting the parameters.

Forecasting, the AR Model, and Sensitivity Analysis for the Natural State Predictive Performance. In the main text, we describe the use of a plain AR model for blackbox prediction of aggregated exit counts. The model is simply

$$E\left[N_{jt} | \mathsf{PAST}\right] = \beta_{\mathrm{AR}_0} + \sum_{w=1}^{30} \beta_{\mathrm{AR}_w} N_{j(t-w)}. \quad \textbf{[S6]}$$

The model is analogous to the entrance process in *SI Text*, *The model for the entrance processes*, except that no smoothing parameter $\theta_{L_{it}}$ is used. The method of least-squares is used to fit this model.

For all models, including the plain AR model, step-ahead forecasts are done by propagating means. This ignores the truncation at zero from Eq. **S1** and similar equations, as positivity is nearly always satisfied and expected values then become linear functions of past expected values for all models. For instance, given tap-out counts observed up to time t_0, we forecast N_{jt} for $t = t_0 + 1, t_0 + 2, \ldots, t_0 + 30$ using the corresponding estimated AR model as follows:

i. Let $C_w = N_{j(t_0+1-w)}$ for $w = 1, 2, \ldots, 30$
ii. For i in $1, 2, \ldots, 30$
  iii. Let $\hat{N}_{j(t_0+i)} = \hat{\beta}_{\mathrm{AR}_0} + \sum_{w=1}^{30} \hat{\beta}_{\mathrm{AR}_w} C_w$
  iv. For w in $30, 29, \ldots, 2$, let $C_w = C_{w-1}$
  v. Let $C_1 = \hat{N}_{j(t_0+i)}$.

This is just an application of iterated expectations to Eq. **S6**.

The blackbox spline model used as a competitor is a regression function from time index t to expected outcome N_{jt},

$$E[N_{jt}] = f_j(t),$$

for some unknown function $f_j(\cdot)$, where $1 \le t \le 1{,}200$. A different spline model is fit to each station. MATLAB function csaps for cubic spline fitting is used, as in *SI Text*, *The Model for the Natural Regime*.

To provide further evidence that our model for the natural regime is robust to overfitting, despite estimating every single OD pair traffic, we did further experiments at a coarser resolution of aggregation.

The task is to predict the aggregated exits for all stations in zones 1 and 2, the busiest zones in London, for traffic originated only in zones 3–9. A new blackbox spline model has to be fitted, because the one in the previous section was exclusively for the full aggregation. Our model, however, is exactly the same, but now aggregating different ODs.

With the same fivefold cross-validations setup, the average RMSE difference per load amounts to 0.001 (SE 0.002), providing more evidence that the OD model is robust.

Fig. S4 illustrates a comparison among the proposed tracking model, the blackbox spline model, and the AR model for Oxford Circus station on Monday, February 14, 2011. Oxford Circus is one of the Tube stations with the highest traffic.

The Probabilistic Flow Model. Although more sophisticated network tomography models are available for accurately estimating traffic volumes from and to pairs of stations (e.g., refs. 5–9), they are in general computationally infeasible at the scale of our massive and complex system, and to the best of our knowledge there is no available software applicable to this study. The approach we take highly simplifies computation of the estimators by relying on simple structural features of the network along a rich source of survey data measuring routes taken by actual passengers.

The RODS is a survey of passenger destinations and the routes chosen (10). For each respondent, his or her origin S_O and destination S_D are recorded, along with change points taken during the passenger's journey. We used the 2012 and 2013 surveys, in which 50,410 and 49,253 distinct routes were observed, respectively, with a total of 8,822,636 journeys.

We need an estimate of $\pi^{\text{OD}}_{h,i,l}$, the probability of passing first through S_h then S_i at line l, during a journey from S_O to S_D. In our context, given a RODS entry, the most important piece of route choice information for a S_O, S_D journey will be the last point of change.[§] From this, given a line closure event that takes place in the sequence $\mathcal{K}^l \equiv \{S_{k(1)}, \ldots, S_{k(n-1)}, S_{k(n)}, S_{k(n+1)}, \ldots, S_{k(M)}\}$ in line l, we obtain $\pi^{\text{OD}}_{k(n),v,l}$, for $S_v \in \{S_{k(n-1)}, S_{k(n+1)}\}$. Let the shortest path in line l between two stations be defined as the shortest of all paths taken with respect to the subgraph of the Tube graph given by stations from line l only, with the respective edges. Given the last point of change S_X for a trip starting at some arbitrary S_O and ending at some $S_D \in \mathcal{K}^l$, we define

$$\mathit{Aligned}(X, k(n), v, D, l) = \begin{cases} 1, & \text{if } S_X \text{ is in } l \text{ and the sequence} \\ & S_X - \cdots S_{k(n)} - S_v - \cdots S_D \\ & \text{is the shortest path in } l \text{ between } S_X \text{ and } S_D. \\ 0, & \text{otherwise.} \end{cases} \quad \textbf{[S7]}$$

The above definition allows for the cases $S_X = S_{k(n)}$ and $S_v = S_D$.

[§]Notice the last point of change is S_O itself if no changes are made.

The idea is to express $\pi^{\text{OD}}_{k(n),v,l}$ as a function of $\pi^{\text{OD}}_{(X,l)}$, defined as the probability of some S_X at l being the last point of change in a journey from S_O to S_D. The relationship is

$$\pi^{\text{OD}}_{k(n),v,l} = \sum_{S_X} \pi^{\text{OD}}_{(X,l)} \times \mathit{Aligned}(X, k(n), v, D, l). \quad \textbf{[S8]}$$

That is, we sum over the probabilities of all possible last points of change for the (S_O, S_D) journey, but including only those S_X such that[¶] the final leg of the journey is $S_X \to \cdots S_{k(n)} \to S_v \to \cdots S_D$ in line l.

To estimate $\pi^{\text{OD}}_{(X,l)}$, we assign it a prior and use RODS data to calculate its posterior, using the posterior expected value as our estimate. Each RODS record specifies the last point of change for particular OD pairs, and the number of passengers who made that choice. Because RODS data do not specify the line of the last change point S_X, which might share multiple lines with S_D, we assume it is the line with the least number of hops between S_X and S_D.

The prior for $\pi^{\text{OD}}_{(X,l)}$ for all possible pairs of stations S_X and line l is a Dirichlet distribution on pairs (X, l) with hyperparameter entries $n \times \alpha(\pi^{\text{OD}}_{X,l})$, where $n = 10$ is an effective sample size parameter. Each hyperparameter $\alpha(\pi^{\text{OD}}_{X,l})$ quantifies a prior choice for the corresponding probability of pair (X, l) with $\sum_{S_X}\sum_l \alpha(\pi^{\text{OD}}_{X,l}) = 1$, the sum going over all choices of station S_X and line l.

Hyperparameters $\alpha(\pi^{\text{OD}}_{X,l})$ are set as follows. Let c_1, c_2, c_3 be three auxiliary hyperparameters of our prior. Let $\mathcal{X}_L$ be the set of triplets (station, line, and cost) defined as follows:

i) If S_O shares a line l with S_D, add (S_O, l, d) to $\mathcal{X}_L$ where d is the distance in hops between S_O and S_D on l.

ii) If S_O does not share any line with S_D, but one can move from S_O to S_D with exactly one line change $l' \to l$ (where S_O is on l' but not l, and S_D is on l but not l'), add to $\mathcal{X}_L$ the triplet $(S_X, l, d + c_2)$ if the distance d for S_X (number of hops from S_O to S_X along l' plus hops from S_X to S_D along l) is the smallest among all stations on l. In case of ties, add all tied pairs.

iii) If moving from S_O to S_D requires at least two line changes, add $(S_X, l, d + c_3)$ to $\mathcal{X}_L$ if (*i*) S_X and S_D are on l, (*ii*) S_X minimizes the sum $d \equiv h_1 + h_2$, where h_1 is the number of hops between S_X and S_D on l and h_2 is total number of hops between S_O and S_X in the Tube subgraph given by the union of all lines other than l.

iv) If S_X fails all three criteria above but it is present in some RODS entry as being the last point of change between S_O and S_D, add (S_X, l, ∞) to $\mathcal{X}_L$ for all l containing both S_X and S_D.

With these criteria, we first set $\alpha(\pi^{\text{OD}}_{X,l})$ to zero for any (X, l) not in $\mathcal{X}_L$. Notice that the implied prior is partially empirical, because the fourth item above looks at RODS data. For all (X, l) that enters $\mathcal{X}_L$ via condition iv above, we set (for now, unnormalized) $\alpha(\pi^{\text{OD}}_{X,l}) = 1/374$.

For all (X, l) in $\mathcal{X}_L$ via i, ii, or iii above, we define

$$s_{Xl} = \max_{\mathcal{X}_L}(c) - c(X, l) + 1,$$

where $c(X, l)$ is the corresponding cost entry of (S_X, l, c) in $\mathcal{X}_L$, and $\max_{\mathcal{X}_L}(c)$ is the maximum of all costs in set $\mathcal{X}_L$. We set hyperparameter $\alpha(\pi^{\text{OD}}_{(X,l)})$ to s_{Xl} if $\max(s_{Xl}) - s_{Xl} \le c_1$, or set it to baseline value $1/374$ otherwise. This thresholding adds an extra

[¶]This is a simplifying assumption, because it discards the possibility of passengers' mistakes or irrational behavior such as passing through the destination of interest and then having to come back. However, we do not consider it worthwhile to assign positive probabilities to these events, because it considerably complicates the problem while having no obvious advantage.

penalty to (X,l) choices that are too far from the optimal choice given by $\max(s_{Xl})$. Finally, we normalize $\alpha(\pi^{\mathrm{OD}}_{X,l})$ so that it sums to 1. Hyperparameters (c_1, c_2, c_3) are set as 5,3,7 by trying different hyperparameter values and checking whether the resulting probabilities $\{\alpha(\pi^{\mathrm{OD}}_{X,l})\}$ were plausible according to the background knowledge of the authors.

For the station closure case, where we observe some station $S_{\mathcal{K}}$ closing but lines around it remaining open, we estimate $\pi^{O\mathcal{K}}_{h,\mathcal{K}}$ similarly, except that the sum in Eq. **S8** is now over all lines.

Notice that we could refine our notion of missing outflow by ignoring flows that go through some closed segments in $\mathcal{K}_l$. That is, we would redefine $\pi^{\mathrm{OD}}_{k(n),v,l}$ for events $S_O \to \cdots \to S_{k(n)} \to S_v \to \cdots S_D$ to be nonzero only where $S_{k(n)}$ is the first station in $\mathcal{K}_l$ in this route. The definition of $\phi^{\mathrm{OUT}}(n)$ (Eq. **3** of main text) includes trajectories that are not possible under disruption (that is, all those of the type $S_O \to \cdots \to S_{k(n\mp 1)} \to S_{k(n)} \to S_{k(n\pm 1)} \to \cdots S_D$). We chose not to account for this refinement to avoid complicating the definitions of inflow and outflow because otherwise we would need different definitions of $\pi^{\mathrm{OD}}_{k(n),v,l}$ for each case. In our sensitivity analyses in *SI Text*, *Sensitivity Analyses* we discuss models for stations only at the endpoints of a disruption, where this is not an issue.

Extracting Disruption Information from the TfL Logs. For the 70 d covered by our smart-card data we also obtained logs of recorded disruption events. Unplanned closures for the line segment events were selected as the ones labeled as "Part Suspended" or "Suspended" in the logs. For the station closure events, the only relevant label was "Closed." This resulted in 3,037 raw entries of line disruptions and 1,335 raw entries of station events. Each raw data point is characterized by the line of the disruption event, its two endpoints (if a line event), the starting and ending time, plus some extra textual information that records other relevant pieces of information, such as the cause of the event. Many events are represented by multiple entries. We merged entries if they had the same endpoints, happened in the same line, and the end time of an event was within 10 min from the start time of the next event. We excluded station events if the corresponding station participated in a line event at the same time. We also excluded Overground events, a service which on average is much less frequent than the Underground and DLR. After merging, we also excluded events less than 10 min long. This resulted in 180 line events, generating 786 data points corresponding to different stations in the ROI of each event. This also resulted in 96 station events, resulting in 191 data points corresponding to the neighbors of each affected station. Table S1 shows the distribution of line events broken down by line.

We want to filter and classify the entries in the available TfL logs in two ways: (*i*) Disruptions that take place in a single direction only are excluded, and (*ii*) events that have delays happening elsewhere in the line are marked as such.

To filter line segment closures that took place in only one direction, we searched for the presence of the substrings "bound", "w/b," and "clock," meant to detect the presence of the keywords "west-/east-/north-/south-bound" in the description, or "clockwise"/ "counterclockwise" (used for the Circle Line of the Underground).

For classifying line segment closures as being accompanied or not by delays, we consider an event as a delay if the word "delay" was included in the textual description of the event. Many events were distinguished as "severe delays," but in our definition of the delay indicator we do not distinguish between severity levels.

A Note on Data Fitting. It should be noted that the data under the natural regime and the data under disruptions are not completely independent. We exclude a day of L_{jt}, $\mathbf{M}_{jt}$, and $\{N_{ijt}\}$ records when fitting the respective entrance, negotiation, and exiting processes if there is any disruption happening at S_j in the particular day. However, we do not exclude any records for the other processes. Recall that the negotiation processes for all stations are functions of all exit counts, and that there are weak but nonzero stochastic dependencies between time points within and time points outside disruptions. As a result, a minor degree of dependence between the natural regime and disruption data exists. However, we do not observe any impact of this dependence in our analysis. In particular, we repeated our analysis without excluding any records from the natural regime and observed no qualitative difference. To illustrate this, the two models obtained from fitting the line disruption model without distance covariates are now

$$E_0\left(\overline{N}^{\mathcal{S}[k_{(n)}]}_{t_1:t_F}\right) = 1.14\phi^{\mathrm{NAT}} - 1.25\phi^{\mathrm{IN}} + 0.16\phi^{\mathrm{OUT}}, \qquad \textbf{[S9]}$$

$$E_1\left(\overline{N}^{\mathcal{S}[k_{(n)}]}_{t_1:t_F}\right) = 1.24\phi^{\mathrm{NAT}} - 1.21\phi^{\mathrm{IN}} + 0.08\phi^{\mathrm{OUT}}. \qquad \textbf{[S10]}$$

SDs for the no-delay case are 0.02, 0.11, and 0.02. SDs for the delay case are 0.02, 0.07, and 0.02. Compared with Eqs. **6** and **7** in the main text, it is clear that no significant difference exists.

Visualization of Distance Functions. In Fig. S5*A* we show a visualization of the quadratic distance functions $f_0(\phi^{\mathrm{DIST}})$ and $f_1(\phi^{\mathrm{DIST}})$ as given by fitting the models for events without delays $(f_0(\cdot))$ and with delays $(f_1(\cdot))$ for line disruptions, as discussed in the main text. The functions are evaluated at observed ϕ^{DIST} points present in the data.

However, recall that the quadratic coefficient for $f_1(\cdot)$ had no strong statistical significance. To perform some sensitivity analysis on how relevant the quadratic term is for both functions, we added yet another nonlinear transformation of ϕ^{DIST} to generate the functions

$$g_x\left(\phi^{\mathrm{DIST}}\right) \equiv \gamma_{3x} + \gamma_{4x}\phi^{\mathrm{DIST}} + \gamma_{5x}\phi^{\mathrm{DIST2}} + \gamma_{6x}\log\left(\phi^{\mathrm{DIST}}\right).$$

The corresponding plot is shown in Fig. S5*B*. Whereas the shape for the no-delay curve has not been dramatically affected, the curve for the delay case confirms that the nonmononicity of $f_1(\cdot)$ is not strongly supported. Recall that we hypothesize that distance functions should be decreasing to penalize outflows, because passengers who are far from their destination are expected not to tap-out earlier, but to look for an alternative route inside the system. The estimated function for the events with delays conforms to this hypothesis. As explained in the main text, there are explanations of why this is not the case for the events without delays. For a final perspective, Fig. S5*C* shows the case when $f_0(\cdot)$ and $f_1(\cdot)$ are constrained to be linear. Once more, the overall conclusion is that the evidence for the no-delay case points to an increasing function, whereas the delay case points to a (more intuitive) decreasing function.

Visualization of Raw Data and Predictions. Fig. S6 provides scatterplots of the outcome variable and selected covariates under the two cases of full ROI or endpoints only. Fig. S6*A* may suggest that ϕ^{NAT} alone provides a good model for observed exit counts under disruption. Although the fit is good, including inflows and outflows in the model improves its predictive abilities compared with a model with ϕ^{NAT} only as shown in Fig. 3 (main text) and Fig. S7*A* and discussed in *SI Text*, *Sensitivity Analyses* below (in particular Table S4). Also, models with inflow and outflow covariates such as Eqs. **6** and **7** from the main text,

$$E_0\left(\overline{N}_{t_1:t_F}^{S[k_{(n)}]}\right) = 1.15\phi^{\mathrm{NAT}} - 1.28\phi^{\mathrm{IN}} + 0.16\phi^{\mathrm{OUT}},$$

$$E_1\left(\overline{N}_{t_1:t_F}^{S[k_{(n)}]}\right) = 1.24\phi^{\mathrm{NAT}} - 1.23\phi^{\mathrm{IN}} + 0.09\phi^{\mathrm{OUT}}$$

explain scenarios where the expected outcome might be less than the expected natural outcome ϕ^{NAT}, depending on the magnitude of ϕ^{IN} with respect to ϕ^{NAT} and ϕ^{OUT}. Without the ability of explaining decreases in expected outcome under disruption, a theory of disruption effects is incomplete. Comparing the two models above against the models with ϕ^{NAT} only,

$$E_0\left(\overline{N}_{t_1:t_F}^{S[k_{(n)}]}\right) = -0.29 + 1.10\phi^{\mathrm{NAT}},$$

$$E_1\left(\overline{N}_{t_1:t_F}^{S[k_{(n)}]}\right) = -0.22 + \phi^{\mathrm{NAT}},$$

it is clear that these do not account for cases where stations can have fewer than ϕ^{NAT} tap-outs under disruptions. Consider the case of Ealing Broadway station (Fig. 4 in main text; see also Fig. S1), which has fewer tap-outs if either its neighborhood in the Central line or its neighborhood in the District line closes down. A model with ϕ^{NAT} only cannot account for this. A theory of system-wide transportation behavior under shocks should allow for this context-sensitive variability, which we achieve using the fundamental concepts of inflows and outflows. Our framework of inflows and outflows interacting with counterfactual outcomes (Eq. **1**, main text) follows the philosophy of making a model simple, but no simpler than it should be, unlike the model with ϕ^{NAT} only.

Fig. S7 *A* and *B* compare the true outcomes under line disruption, for each of the affected stations (768 cases), against the leave-one-out predictions as given by the model combining delayed and nondelayed events. Fig. S7 *C* and *D* perform the analogous comparison for single-station disruption events.

Sensitivity Analyses. In *SI Text*, *Delays and endpoint models* we evaluate how delays interact with outcomes and under which conditions. A more in-depth look at predictive results is discussed in *SI Text*, *Predictive results*. The effect of distance measures and its evaluations is further explored in *SI Text*, *Distance model*. Finally, in *SI Text*, *Temporally fine-grained predictions* we comment on predictive results for the 1-min resolution setup.

Delays and endpoint models. Consider an alternative model for the line segment disruption problem, where the effect of delays is additive, as opposed to having two separate models for each state of the delay variable:

$$E\left[\overline{N}_{t_1:t_F}^{S[j]}\middle|\mathsf{PAST}\right] = \beta_0 + \beta_1\phi^{NAT} + \beta_2\phi^{IN} + \beta_3\phi^{OUT} + \beta_4\phi^{DELAY}.$$

The fit of this model is summarized in Table S2. It has a good R^2 fit compared with the individual models for $\phi^{DELAY}=0$ and $\phi^{DELAY}=1$, but the linear coefficient of the delay covariate does not significantly contribute to the model. However, consider the case where given a disrupted segment $\mathcal{K}=\{S_{k(1)},\ldots,S_{k(M)}\}$, the ROI is composed only of the endpoints $\{S_{k(1)},S_{k(M)}\}$ (again, excluding those that are not connected to any station outside $\mathcal{K}$—which will be the case for stations at the end of a line). A priori, this particular ROI suggests behavior that differs from the average station in $\mathcal{K}$, as they receive passengers from line l (as opposed to midpoints $S_{k(n)}$ in $\mathcal{K}$ with external connections; in that case, passengers would be planning to change lines at $S_{k(n)}$). We fit three other models for this subset of stations, again ignoring distance covariates to provide a set of models easier to interpret. These other three models are shown in Table S2. These models in general follow the theoretical structure of having positive outflow and negative inflow contributions, and no strong evidence for intercepts. There is evidence of different behavior between the models with and without delays—in particular, on the contribution of outflows, which is precisely the flow measure we believe to be most affected by delays. Differences in the contribution of outflows were also detected for the model regulated by average distance covariates and all stations, as discussed in the main text. At the same time, the additive contribution of the delay indicator is not significant (Table S2).

Predictive results. Fig. 1 in the main text shows predictive results for number of tap-outs per minute. To give a sense of scale for the RMSE using Underground stations as examples, the RMSE for the 30-min step-ahead problem is 24.9 ± 0.54 for Oxford Circus (average daily traffic per minute at the order of 60) and 22.35 ± 1.69 for King's Cross (average daily traffic per minute at the order of 50).

Table S3 shows predictive results corresponding to Fig. 3 in the main text. The table also addresses how the result changes under different delay conditions. Each of the four panels shows results for a different model: top left is for our model in Eq. **5** (main text) applied to all data, as in Fig. 3*A* (main text); top right and bottom left is our model in Eq. **5** (main text) applied to subsets of data classified according to ϕ^{DELAY}; bottom right is for our model in Eq. **8** (main text) applied to all data, and corresponds to the graph in Fig. 3*B* (main text).

The columns in Table S3 are as follows:

- Filter: indicates which test points are being used in the calculation of the respective statistics. A value of n for Filter means that only test points with an outcome variable of size n or more are used in the calculation of the remaining items in the respective row ($n=0$ being the complete set of test folds);
- Sample: the number of test points that satisfy the filter condition;
- Error: average "absolute error," which for each test point is the absolute value of the difference between the test tap-out and the tap-out predicted by the respective model. This is averaged over the selected test points. For line disruptions, the model is the one in Eq. **5** of the main text, whereas for station disruptions the model is the one in Eq. **8** of the main text;
- Diff_N: difference between the absolute error of the model with ϕ^{NAT} as the only covariate, and the absolute error of our respective model, averaged over the selected test points;
- p_N: *P* value for the signed paired *t* test, null hypothesis $\mathrm{Diff}_N=0$ against the alternative hypothesis $\mathrm{Diff}_N>0$;
- Diff_U: difference between the absolute error of the model with uniform probabilities for the passenger flows, and the absolute error of our respective model, averaged over the selected test points;
- p_U: *P* value for the signed paired *t* test, null hypothesis $\mathrm{Diff}_U=0$ against the alternative hypothesis $\mathrm{Diff}_U>0$.

Concerning the results for station disruption (bottom right of Table S3), although there is no statistical difference with respect to the uniform probability case, overall our model shows a consistent advantage over this competitor. Flow probabilities seem to matter less in this problem. In particular, without distance covariates for simplicity the model for station exits[‖] is $1.10\phi^{\mathrm{NAT}} + 0.21\phi^{\mathrm{OUT}} + 0.25\phi^{\mathrm{OUT}'}$. That is, the contributions of ϕ^{OUT} and $\phi^{\mathrm{OUT}'}$ are approximately the same in this case.

[‖]All coefficients significant at a 0.01 level, SEs approximately 0.08 for the two outflow covariates.

Distance model. The main motivation for including distance covariates in our analysis is to provide some insight of the impact of outflows as distances to other stations in a disrupted segment change. However, although the contribution of the distance covariates to the model structure is strongly statistically significant, it provides no predictive gain. Table S4 illustrates this by comparing the predictions between our full model with distance covariates and the simpler model, which uses covariates ϕ^{NAT}, $\phi^{\mathrm{IN}}, \phi^{\mathrm{OUT}}$ in the $\phi^{\mathrm{DELAY}} = 1$ case. The distance covariates only give a small but not statistically significant advantage at the cases with the larger stations.

To assess more complex uses of distance covariates, consider the case where, instead of averaging over distances with respect to all stations in the affected line, we weight each flow contribution by destination before aggregating them. In particular, we downweight $\pi^{\mathrm{OD}}_{k(n),v,l} \times \mu^{0}_{\mathrm{OD}t;t_1}$ by a function $g_u(dist(S_D, S_{k(n)}))$, where $g_u(\cdot)$ is some nonlinear transform of a normalized Euclidean distance $dist(\cdot,\cdot)$ between the stations.** Moreover, we once more avoid "self-exits" by summing over $S_O \neq S_D$ only. Therefore, for each function $g_u(\cdot)$, we define the covariate $\phi^{\mathrm{OUT}}_u(n)$ as

$$\phi^{\mathrm{OUT}}_u(n,t) \equiv \sum_{S_D \in \mathcal{K}_l \setminus S_{k(n)}} \sum_{S_O \neq S_D} \sum_{S_v \in \mathcal{N}_{\mathcal{K}_l}(n)} \frac{\pi^{\mathrm{OD}}_{k(n),v,l} \times \mu^{0}_{\mathrm{OD}t;t_1}}{g_u\left(dist(S_D, S_{k(n)})\right)}$$
$$\phi^{\mathrm{OUT}}_u(n) \equiv \sum_{t=1}^{t=t_F} \phi^{\mathrm{OUT}}_u(n,t)/F. \qquad \text{[S11]}$$

We use a set of four different functions, $g_1(x)=1$, $g_2(x)=\log(x)$, $g_3(x)=x$, and $g_4(x)=x^2$, to characterize four distinct covariates. As shown by results in Table S4 (columns Diff_2 and p_2), this extra complication did not provide a measurable payoff, whereas the model described in the main text provides an easier interpretation of the role of distance in our outflow measures.

Temporally fine-grained predictions. One final sensitivity analysis experiment covers the case where our model is not for the average exit count $\overline{N}^{S_{k(n)}}_{t_1:t_F}$, but for each individual minute: $N^{S_{k(n)}}_t$. As we mentioned in the main text, the model and the procedure for fitting it is directly applicable to any subset of $t_1:t_F$, because it relies on the same counterfactual exit counts generated for the entire period. Treating each time point as a separate training point, we have a sample of 17,844 measurements for the case without delays and 21,953 for the case with delays. We fit models without the distance covariates, obtaining

$$E_0\left[N^{S_{k(n)}}_t\right] = 0.49 + 1.13\phi^{\mathrm{NAT}}_t(n,t) - 1.23\phi^{\mathrm{IN}}(n,t) + 0.14\phi^{\mathrm{OUT}}(n,t),$$
$$E_1\left[N^{S_{k(n)}}_t\right] = -0.03 + 1.19\phi^{\mathrm{NAT}}_t(n,t) - 1.02\phi^{\mathrm{IN}}(n,t) + 0.10\phi^{\mathrm{OUT}}(n,t),$$

where all parameters are significant ($P < 10^{-15}$) except the intercept for the model with delays, assuming independence of the time points. Table S5 shows cross-validated predictive results for these two models using the same criteria as in the previous predictive evaluations. Cross-validation is performed by using as the test set the whole time series of one station at one disruption event. Errors are averaged over test folds, each fold averaged over time.

By comparing Table S5 to the results in Table S3, where errors varied between 3 and 11 persons per minute, it is evident that prediction at individual minutes is more difficult than averages over $t_1:t_F$. We also improve the simple baseline based on $\phi^{\mathrm{NAT}}(n,t)$ only by a small margin (all stations, 0.3–1 in the delayed case, 0.2–3.7 in the no-delays case) because we lose predictive power at this resolution. We should, however, notice that for the purposes of transportation management and policy making (such as providing timely alternative transportation under disruption and long-term planning for expansions) the reliable average prediction over the disruption period is valuable information.

**We use latitude and longitude as coordinates. The normalization factor is the Euclidean distance between King's Cross and Heathrow Terminals 1–2–3, to make distances more interpretable.

1. Transport for London (2012) Travel in London, Report 5. Available at www.tfl.gov.uk/cdn/static/cms/documents/travel-in-london-report-5.pdf. Accessed May 26, 2013.
2. Transport for London (2012) Join in the celebrations across the capital this Summer with a limited edition Summer Oyster card. Available at https://www.tfl.gov.uk/info-for/media/press-releases/2012/june/join-in-the-celebrations-across-the-capital-this-summer-with-a-limited-edition-summer-oyster-card. Accessed May 26, 2013.
3. Mnih A, Salakhutdinov R (2008) Probabilistic matrix factorization. *Advances in Neural Information Processing Systems*, eds Platt JC, Koller D, Singer Y, Roweis ST (Curran Associates, Down, UK), Vol 20, pp 1257–1264.
4. Wood S (2006) *Generalized Additive Models: An Introduction with R* (Harvard Univ Press, Cambridge, MA).
5. Vardi Y (1996) Network tomography: Estimating source-destination traffic intensities from link data. *J Am Stat Assoc* 91:365–377.
6. Tebaldi C, West M (1998) Bayesian inference on network traffic using link count data. *J Am Stat Assoc* 93(442):557–573.
7. Cao J, Davis D, Van Der Viel S, Yu B (2000) Time-varying network tomography: router link data. *J Am Stat Assoc* 95:1063–1075.
8. Airoldi EM, Faloutsos C (2004) Recovering latent time-series from their observed sums: Network tomography with particle filters. *Proceedings of the Tenth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (ACM, New York), pp 30–39.
9. Airoldi EM, Blocker AW (2013) Estimating latent processes on a network from indirect measurements. *J Am Stat Assoc* 108(501):149–164.
10. Transport for London (2014) *Rolling Origin and Destination Survey: The Complete Guide, 2003*. Revised October 2010, March 2012, and January 2014 (London Underground Limited, UK).

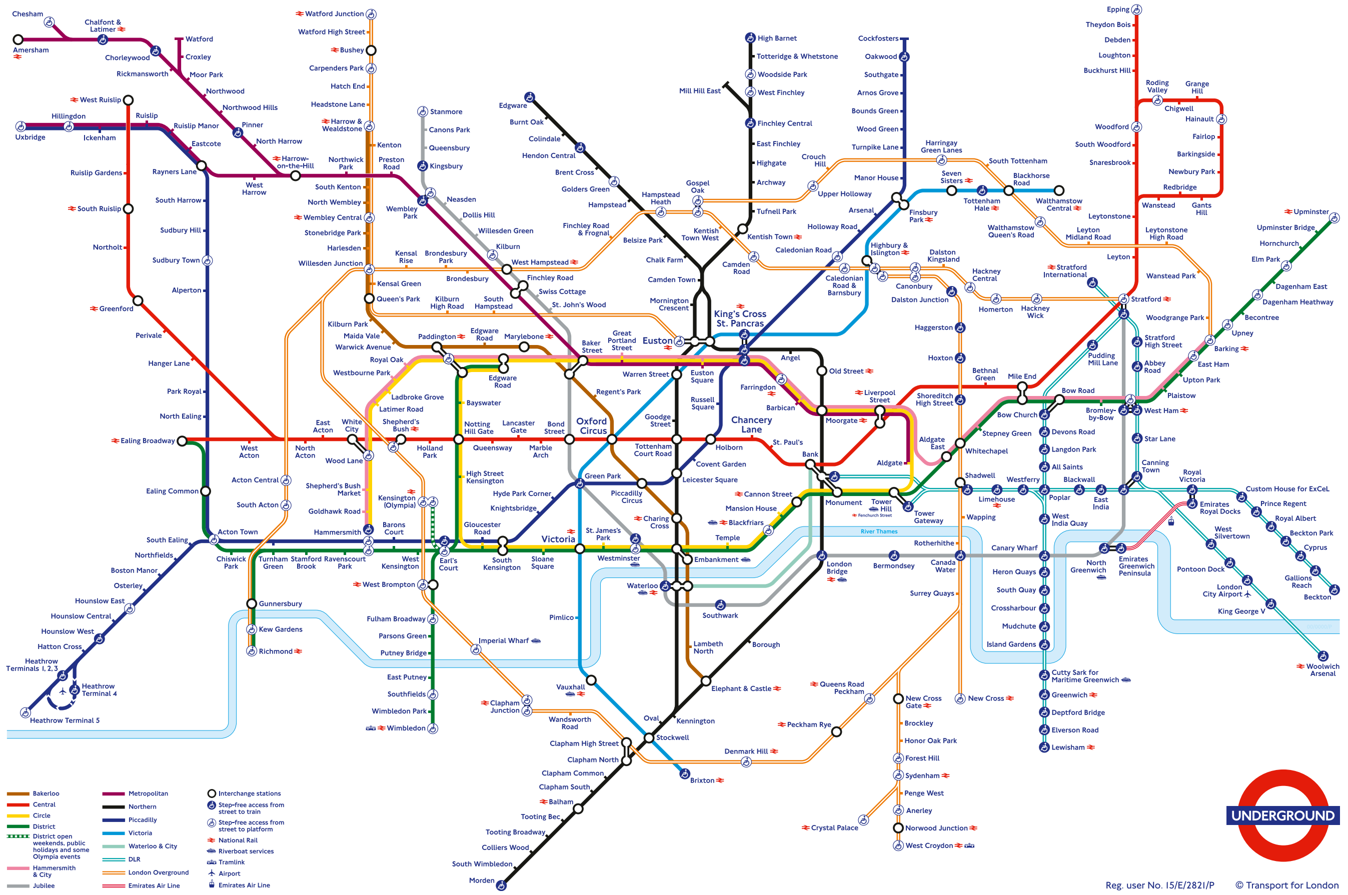

Fig. S1. Tube map including Underground, Overground, and DLR. Reproduced by kind permission of Transport for London, © TfL.

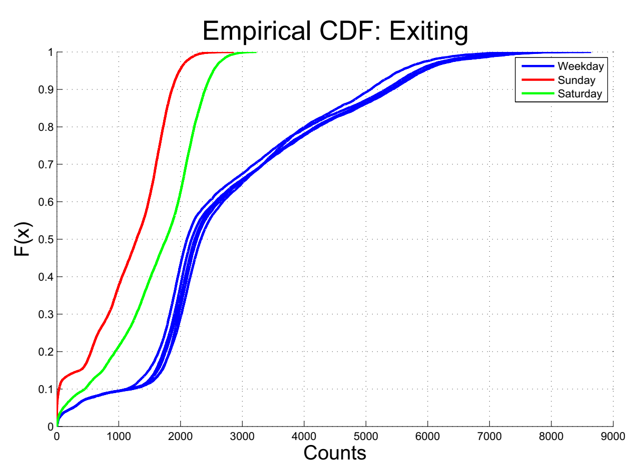

Fig. S2. Cumulative distribution function of exit counts aggregated per day for weekdays and weekends.

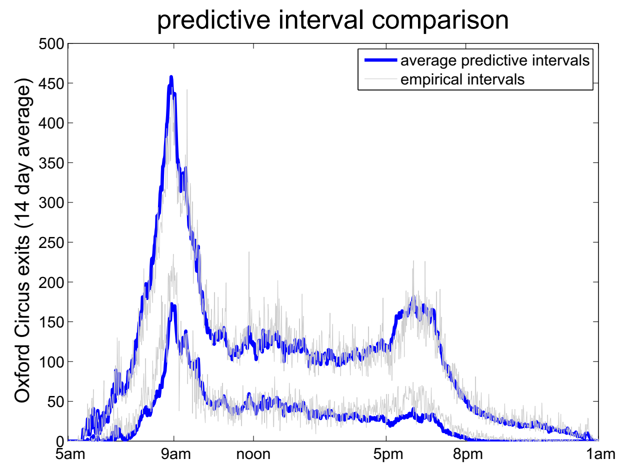

Fig. S3. Comparison between predictive 95% confidence interval and empirical intervals of exits from Oxford Circus station. Average predictive intervals are given by one-step-ahead predictions, which then are averaged over 14 test days. Empirical intervals are given as the empirical quantiles for the 5% symmetric tails.

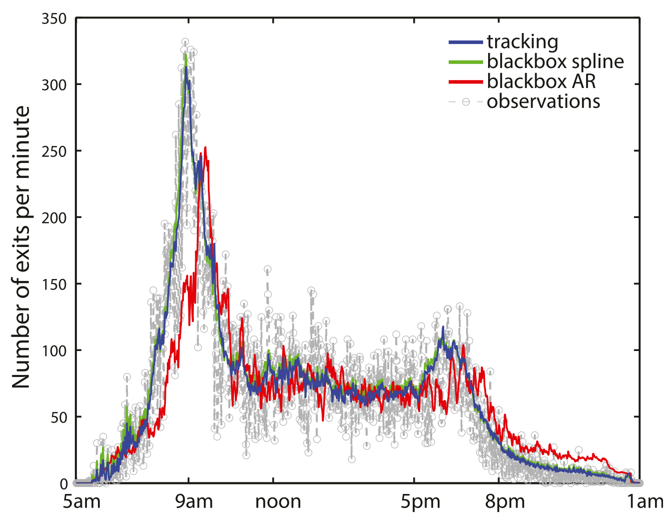

Fig. S4. Thirty-minutes-ahead prediction of the overall number of exits per minute at Oxford Circus station on Monday, February 14, 2011, given past observations of the day.

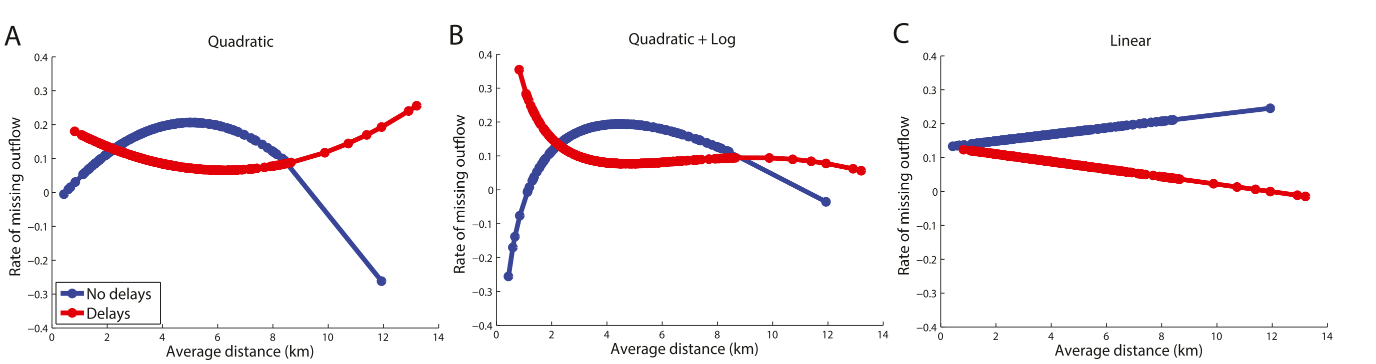

Fig. S5. Effect on ϕ^{DIST} (horizontal axis) on the weighting of ϕ^{OUT} under no delays and under delays using three variations of the distance function. (*A*) Effect for quadratic model. (*B*) Effect for model with quadratic and logarithmic transformations. (*C*) Effect for linear model.

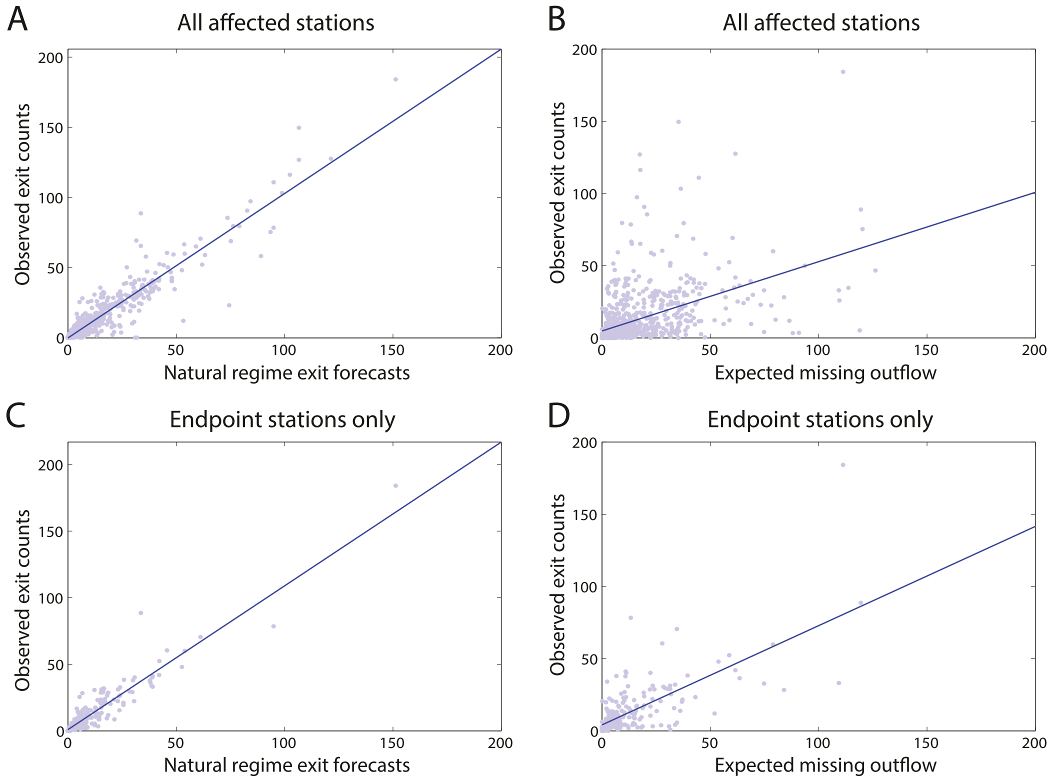

Fig. S6. Association between observed exit counts in stations affected by line disruptions, and some covariates used in prediction. (*A*) Association with respect to the predicted expectations of exits in overall region of interest under the natural regime. (*B*) Association with respect to the missing outflows ϕ^{OUT}. (*C*) Association with respect to natural regime at line segment endpoints only. (*D*) Association with respect to ϕ^{OUT} at line segment endpoints only. The lines in *A* and *C* are the fit given by least squares.

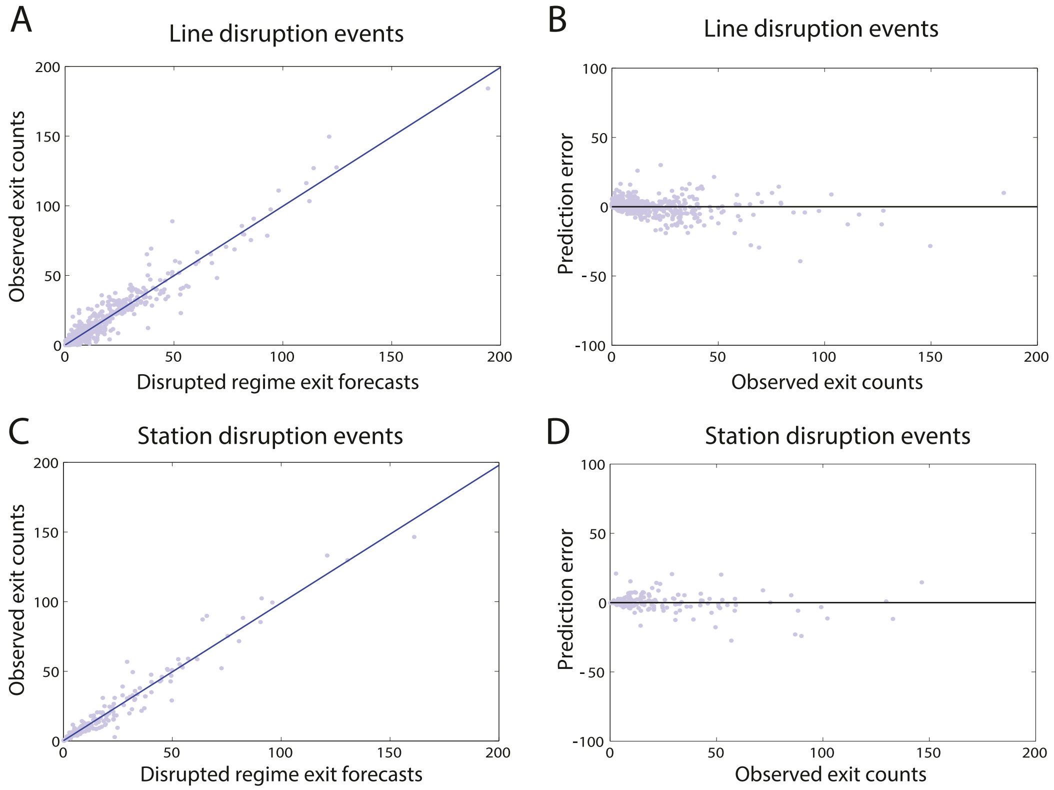

Fig. S7. Visual comparison of predicted and real outcomes for disruptions events. (*A*) All 768 cases of stations affected by line disruptions. (*B*) Corresponding residuals (difference between truth and predicted). (*C* and *D*) The analogous information for station disruption events. The lines in *A* and *C* are the fit given by least squares.

Table S1. Distribution of line events

Line name	No. of events
DLR	11
Bakerloo	12
Central	9
Circle	3
District	29
Hammersmith and City	6
Jubilee	25
Metropolitan	22
Northern	15
Piccadilly	26
Victoria	12
Waterloo and City	0

Table S2. Estimates of model for exit counts in affected line segments

	Linear delay effect (all stations) (N = 768, R^2 = 0.92)		Linear delay effect (endpoint stations only) (N = 204, R^2 = 0.91)	
Parameter	Estimate ± SE	P value	Estimate ± SE	P value
Intercept	0.05 ± 0.44	0.91	1.56 ± 0.87	0.07
ϕ^{NAT}	1.21 ± 0.01	$<10^{-15}$	1.02 ± 0.03	$<10^{-15}$
ϕ^{IN}	−1.22 ± 0.06	$<10^{-15}$	−0.76 ± 0.15	$<10^{-12}$
ϕ^{OUT}	0.10 ± 0.01	$<10^{-3}$	0.20 ± 0.03	$<10^{-6}$
ϕ^{DELAY}	−0.004 ± 0.09	0.96	−0.23 ± 0.16	0.15
	$\phi^{DELAY} = 1$ (endpoint stations only) (N = 96, R^2 = 0.85)		$\phi^{DELAY} = 0$ (endpoint stations only) (N = 108, R^2 = 0.91)	
Intercept	1.64 ± 0.90	0.07	0.57 ± 0.53	0.29
ϕ^{NAT}	0.99 ± 0.06	$<10^{-15}$	1.03 ± 0.03	$<10^{-15}$
ϕ^{IN}	−0.71 ± 0.21	0.001	−0.81 ± 0.15	$<10^{-12}$
ϕ^{OUT}	0.11 ± 0.03	0.001	0.20 ± 0.03	$<10^{-7}$

Table S3. Comparison of prediction errors of the full models against the model with ϕ^{NAT} only (Diff_N and p_N) and the model with flows given by uniform probabilities (Diff_U and p_U)

	Lines, all data						Lines, $\phi^{DELAY} = 0$					
Filter	Sample	Error	Diff_N	p_N	Diff_U	p_U	Sample	Error	Diff_N	p_N	Diff_U	p_U
0	768	3.0	0.4	0.00	0.4	0.00	344	2.7	0.3	0.03	0.2	0.11
5	392	4.6	0.7	0.00	0.6	0.00	153	4.5	0.7	0.02	0.4	0.06
10	272	5.5	0.8	0.00	0.8	0.01	97	5.3	1.1	0.01	0.6	0.05
15	200	6.2	0.9	0.02	0.8	0.02	68	6.2	1.4	0.01	0.7	0.08
20	165	6.9	0.7	0.06	0.8	0.04	51	7.1	1.5	0.03	0.8	0.09
25	130	7.3	0.8	0.07	0.9	0.04	41	7.6	2.1	0.01	1.2	0.01
30	95	8.0	0.8	0.13	1.0	0.07	33	8.2	2.2	0.02	1.4	0.01
35	65	8.7	1.4	0.07	1.8	0.02	24	8.8	2.2	0.07	2.0	0.01
40	49	9.3	2.3	0.02	2.7	0.01	15	10.2	3.7	0.05	2.9	0.01
45	36	9.6	4.4	0.00	4.7	0.00	12	10.8	4.5	0.05	3.5	0.01
50	31	10.1	5.3	0.00	5.5	0.00	10	10.7	6.5	0.01	4.2	0.00
55	27	10.8	5.0	0.00	5.3	0.00	9	11.5	5.7	0.03	3.8	0.01
60	22	11.6	4.4	0.00	4.9	0.00	8	11.9	6.1	0.04	4.3	0.01
65	20	12.1	4.6	0.00	5.3	0.00	8	11.9	6.1	0.04	4.3	0.01
70	16	10.6	4.8	0.01	5.6	0.00	6	10.2	5.5	0.11	3.8	0.03
75	15	11.0	4.9	0.01	6.2	0.00	5	11.9	6.2	0.13	3.7	0.06
80	11	12.2	5.3	0.02	6.1	0.00	4	14.5	7.0	0.17	4.6	0.05
85	11	12.2	5.3	0.02	6.1	0.00	4	14.5	7.0	0.17	4.6	0.05
90	9	10.0	4.4	0.06	5.4	0.00	2	6.3	8.3	0.31	2.3	0.14
	Lines, $\phi^{DELAY} = 1$						Stations, all data					
0	424	3.4	0.5	0.00	0.5	0.00	191	3.6	0.4	0.02	0.0	0.42
5	239	4.7	0.7	0.01	0.8	0.01	140	4.5	0.5	0.06	0.1	0.41
10	175	5.6	0.9	0.02	1.0	0.01	97	5.4	0.6	0.08	0.0	0.44
15	132	6.2	0.8	0.08	0.8	0.07	76	5.8	1.1	0.02	0.2	0.27
20	114	6.8	0.7	0.13	0.8	0.10	58	6.5	1.2	0.04	0.3	0.23
25	89	7.1	0.6	0.20	0.8	0.14	47	6.7	0.8	0.14	0.1	0.45
30	62	7.9	0.6	0.28	0.8	0.20	43	6.6	1.1	0.07	0.5	0.08
35	41	8.4	1.9	0.07	2.2	0.05	32	7.4	1.6	0.02	0.8	0.04
40	34	8.5	3.0	0.02	3.3	0.01	29	7.5	1.8	0.02	0.9	0.04
45	24	8.5	5.8	0.00	6.2	0.00	25	8.3	1.5	0.06	1.0	0.06
50	21	9.2	6.3	0.00	6.8	0.00	21	8.5	2.1	0.03	1.2	0.05
55	18	9.7	6.2	0.00	6.8	0.00	16	9.3	3.0	0.01	1.5	0.03
60	14	10.6	4.7	0.00	5.5	0.00	11	9.9	2.2	0.08	0.7	0.25
65	12	11.4	4.9	0.01	5.6	0.00	11	9.9	2.2	0.08	0.7	0.25
70	10	10.4	4.6	0.03	5.5	0.01	11	9.9	2.2	0.08	0.7	0.25
75	10	10.4	4.6	0.03	5.5	0.01	10	10.0	2.3	0.10	0.5	0.31
80	7	9.9	6.1	0.02	6.4	0.02	9	11.1	1.6	0.18	−0.1	0.55
85	7	9.9	6.1	0.02	6.4	0.02	9	11.1	1.6	0.18	−0.1	0.55
90	7	9.9	6.1	0.02	6.4	0.02	5	8.4	1.0	0.36	1.5	0.06

See text for details. In particular, absolute errors for our models are shown in the error column.

Table S4. Comparison of prediction errors of the full model (Eq. 5, main text) with distance covariates against the model without distance covariates, for $\phi^{\mathrm{DELAY}} = 1$

Filter	Sample	Error	Diff_1	p_1	Diff_2	p_2
0	424	3.4	−0.0	0.75	0.0	0.20
5	239	4.7	0.0	0.45	0.1	0.05
10	175	5.6	−0.0	0.55	0.1	0.25
15	132	6.2	−0.0	0.63	0.0	0.37
20	114	6.8	−0.0	0.56	0.1	0.32
25	89	7.1	−0.0	0.66	0.1	0.29
30	62	7.9	0.0	0.43	−0.1	0.85
35	41	8.4	0.0	0.45	−0.1	0.67
40	34	8.5	0.1	0.28	0.0	0.46
45	24	8.5	0.0	0.41	0.2	0.13
50	21	9.2	−0.1	0.64	0.1	0.29
55	18	9.7	−0.0	0.57	0.1	0.25
60	14	10.6	−0.2	0.78	0.0	0.49
65	12	11.4	−0.3	0.89	−0.0	0.55
70	10	10.4	−0.4	0.91	−0.2	0.81
75	10	10.4	−0.4	0.91	−0.2	0.81
80	7	9.9	−0.3	0.78	−0.1	0.65
85	7	9.9	−0.3	0.78	−0.1	0.65
90	7	9.9	−0.3	0.78	−0.1	0.65
95	6	11.2	−0.2	0.68	−0.0	0.52

Prediction error is defined by the absolute difference between the true number of tap-outs in an event of interest and the predicted number of tap-outs, averaged over test points in an LOOCV procedure. Column Diff_1 compares the difference in prediction error between the two models, positive numbers indicating an advantage for the full model. Column p_1 is the P value of a one-sided t test under the alternative hypothesis that our model (Eq. **5**, main text) is better than the competing model. Columns Diff_2 and p_2 are measured with respect to yet another way of using distance covariates, as explained in the text.

Table S5. Prediction errors of the models for all cases (with and without delay) for each individual minute

Filter	Error (no delay)	Error (delay)
0	6.8	8.1
5	11.6	11.8
10	14.1	13.9
15	16.0	15.4
20	17.4	16.3
25	18.5	17.0
30	20.3	18.9
35	22.8	21.6
40	26.8	23.0
45	28.6	26.0
50	30.3	27.3
55	32.1	28.9
60	32.2	30.4
65	32.2	32.4
70	35.0	34.3
75	36.6	34.3
80	37.7	31.3
85	37.7	31.3
90	43.3	31.3
95	43.3	31.6