# A New Class of Doubly Stochastic Day-to-Day Dynamic Traffic Assignment Models

Katharina Parry
*Institute of Fundamental Sciences, Massey University, Palmerston North, New Zealand*

David P. Watling *
*Institute for Transport Studies, University of Leeds, Leeds, United Kingdom*

Martin L. Hazelton†
*Institute of Fundamental Sciences, Massey University, Palmerston North, New Zealand*

---

**Abstract**
Real-life systems are known to exhibit considerable day-to-day variability. A better understanding of such variability has increasing policy-relevance in the context of network reliability assessment and the design of intelligent transport systems.

Conventional equilibrium models are ill-suited, because deterministic models such as these do not account for any kind of variability. At best, these types of models are restricted to finding a steady state of the mean flow patterns, they cannot capture the variance in flows as well. A more suitable alternative are stochastic day-to-day dynamic models studied by Cascetta (1989). These types of traffic assignment models represent the traffic flows via a Markov process, where the current route flows are modelled as a function of previous traffic conditions. Day-to-day dynamic models differ from equilibrium models in that day-to-day changes in the system are modelled dependent on the time and thus allow for a far wider representation of traveller behaviour. However, to some degree they still suffer from some of the limitations of equilibrium analyses, in that while they permit variation they are still wedded to the concept of 'stationarity'.

In this paper, we show how these Markovian day-to-day dynamic traffic assignment models can be extended by replacing a subset of the fixed parameters in the Markov model with random processes. The resulting models are analogous to Cox process models. They are conditionally non-stationary given any realization of the parameter processes. We present numerical examples that demonstrate that this new class of doubly stochastic day-to-day traffic assignment models can indeed reproduce features such as the heteroscedasticity of traffic flows observed in real-life settings.

*Key words:* Markov, transportation, network, doubly stochastic, heteroscedasticity, day-to-day

*Email: `d.p.watling@its.leeds.ac.uk`; tel. +44 113 343 6612; fax +44 113 343 5334; postal address: Institute for Transport Studies, University of Leeds, Leeds, United Kingdom

†Email: `m.hazelton@massey.ac.nz`; tel. +64 6 356 9099 x2483; fax +64 6 350 5682; postal address: Massey University PN322, Private Bag 11222, Palmerston North, New Zealand 4442

# 1 Introduction

Traffic assignment models translate origin-destination travel demands into flows and travel times on the links of a network, by modelling the interaction between traffic congestion and drivers' route choice decisions. Historically these models have been based on notions of deterministic or stochastic equilibrium, potentially time-sliced within the day as in dynamic traffic assignment approaches. Such models are not based on empirical considerations, but rather are based on idealised behaviour that economic agents might display, if fully informed and travelling in a system with no variability. Real-life systems are known to exhibit considerable day-to-day variability, and indeed understanding such variability has increasing policy-relevance in the context of network reliability assessment and the design of intelligent transport systems. It is therefore relevant to question the relationship between the models of idealised systems and the considerable real-life evidence we now have. As Cascetta (1989) notes, this is no simple task, since due to the non-linearity of the traffic assignment process, by neglecting variation equilibrium models cannot even be interpreted as unbiased estimators of *mean* traffic conditions. We believe it is therefore timely to ask: if starting instead with the kind of empirical evidence typically available for traffic networks, what kind of traffic assignment model might we develop in order to describe the underlying features and characteristics of the observed traffic flows? The present paper aims to propose such an empirically-inspired class of models.

Suppose, then, that we obtain traffic flow data from successive points in time over a period of days. We may encounter a variety of different phenomena. For example, we may observe regular systematic patterns, perhaps due to changes in route choices during the week in order to avoid roads prone to congestion on weekdays. We may also encounter irregular fluctuations in the traffic volume as a response to a variety of factors, from weather patterns to minor traffic accidents.

To illustrate these points we introduce some observed traffic flow data from New York state. The data were collected over sequences of consecutive (week)days in 2009 measured on road links as displayed in Figure 1. The data were sourced from the New York State Department of Transportation website, `https://www.dot.ny.gov/divisions/engineering/technical-services/highway-data-services`.

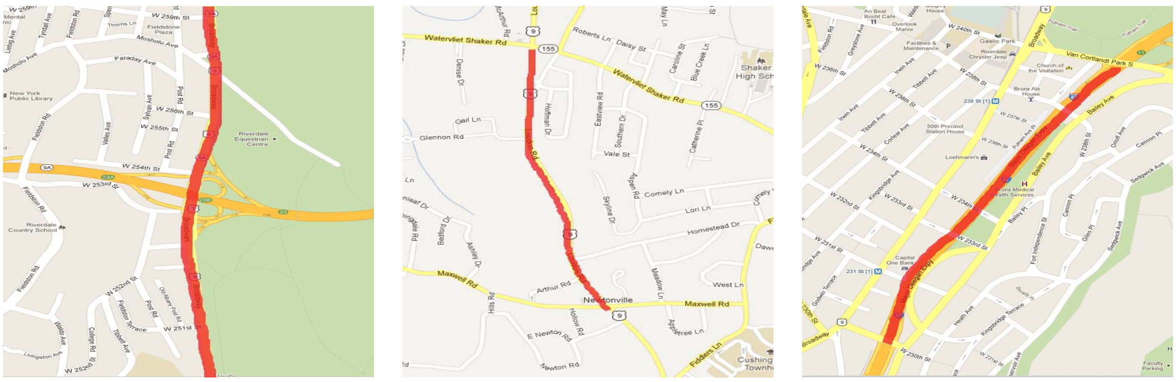

(a) Link 10012, Bronx. (b) Link 110137, Albany. (c) Link 10027, Bronx.

Figure 1: Locations of road sections in New York. Reproduced courtesy of 2012 Google Map data.

The first set of counts were northbound traffic observed on Broadway Road in the Bronx, coded as link 10012, on weekdays in 2009. The second sequence of observations were vehicles travelling north on Loudon Road in Albany, coded as link 110137, over a span of seven months from 22 June until 20 November, with flows observed on the weekend removed. The third collection of vehicle flows come from cars heading north on the Major Deegan Expressway in the Bronx, the road section coded as 10027. The data were gathered from 9 March until 26 June, again with weekend data removed. All traffic flows were measured for hour-long periods, in particular from 4-5 pm for the roads located in the Bronx and 10-11pm for the road in Albany.

Figures 2, 3 and 4, are the corresponding time series (and associated autocorrelation plots) for the observed traffic flows. They demonstrate a variety of patterns of temporal dependence.

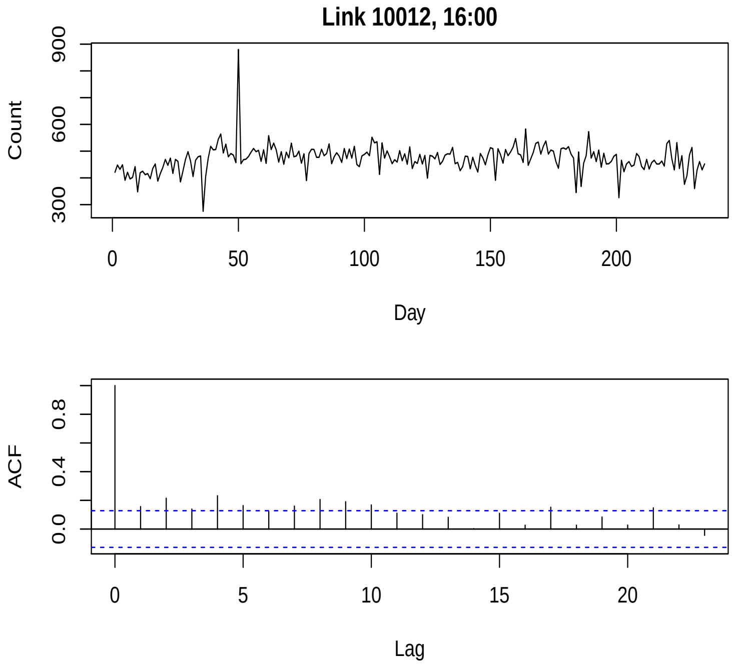


Figure 2: Time series and associated autocorrelation plot for traffic counts northbound on Broadway (route 9) in the Bronx, NY.

In Figure 2 the day-to-day variation in flows seems almost entirely haphazard, with a few extreme events, particularly at time points 36 and 50. Aside from these outliers, we might effectively be observing a sequence from a white noise series; that is, observations of entirely independent and identically distributed random variables. While the autocorrelation plot indicates that the data exhibits some mild serial dependence, this is weak enough to suggest that the series might be modelled as white noise.

Let us now consider a different kind of characteristic pattern. In Figure 3 we see a clear deviation from a white noise series. There is manifest periodicity in the time series plot, leading to strong autocorrelation at lags of multiples of five. This indicates a marked day-of-the-week effect. While the mean flow remains more or less constant, there is at least a hint of increased variability towards the end of the observational period.

Thirdly and finally, in Figure 4 we see some still different characteristic features. There is clear evidence of heteroscedasticity, with the variation being significantly larger in the second half of the period than the first. (The variance ratio between these periods is 3.40, with a nominal p-value of 0.0002 based on an F-test on 39 degrees of freedom.) In addition, the autocorrelations show no significant lags, thus no periodicity as seen in Figure 3 is evident; thus a model that

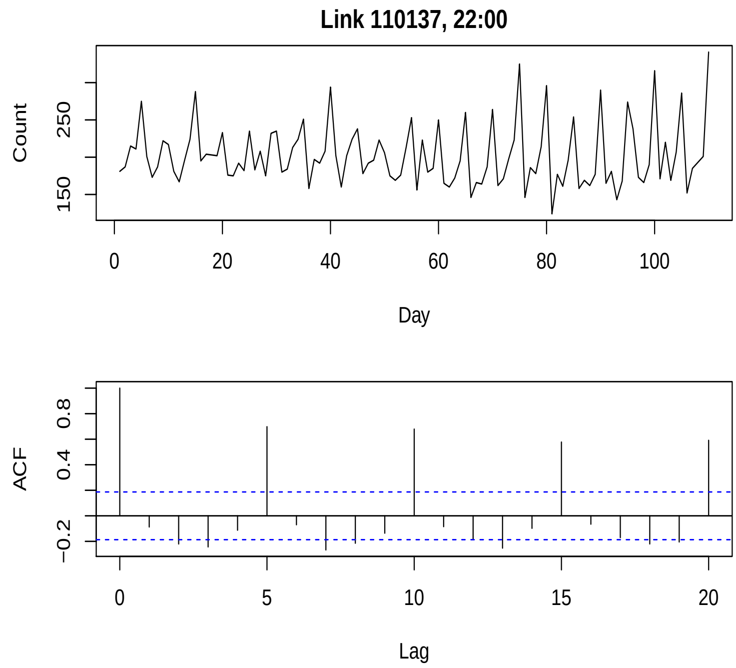


Figure 3: Time series and associated autocorrelation plot for traffic counts northbound on Loudon Road (route 9) in Albany, NY.

included a covariate for the weekday effect would not suffice in this case. The implications of this are that we are dealing with a variation in the magnitude of the flows over time that cannot be adequately modelled as white noise. The opposite seems to be the case: the behaviour of the link flows suggests that there is underlying effect that is causing the temporal variability.

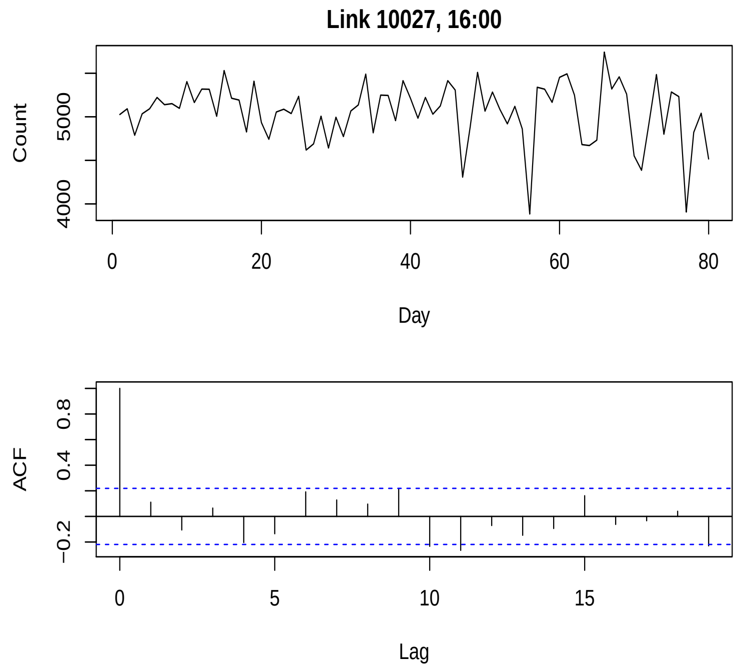


Figure 4: Time series and associated autocorrelation plot for traffic counts northbound on Major Deegan Expressway (route 87) in the Bronx, NY.

As we remarked above, conventional Wardrop equilibrium (Wardrop 1952) and stochastic user equilibrium (Daganzo and Sheffi 1977) models are ill-suited in this kind of setting, because deterministic models such as these do not account for any kind of variability. At best, these types of models are restricted to finding a steady state of the mean flow patterns, they cannot capture the variance in flows as well.

Traffic models that also account for 'second-order' properties are an appropriate alternative to consider. As we consider the case that link flows were observed over the course of successive days, we can apply day-to-day dynamic models which model the evolution over days of travel choices. In particular, we implement day-to-day dynamic models as discussed by Cantarella and Cascetta (1995). Day-to-day dynamic models differ from equilibrium models in that they analyse the evolution over days of travel choices using a learning model based on the effect of previous traffic conditions on the travellers, that is, day-to-day changes in the system are modelled dependent on the time. The framework of these types of models is very flexible as it allows for an effortless integration of different behavioural patterns and types of learning processes (Watling and Hazelton 2003). In other words, day-to-day dynamic models can concentrate to a great extent on drivers' responses to external factors such as congestion or other changes in the system.

Both deterministic and stochastic approaches to day-to-day dynamic modelling have been studied. Each can describe temporal variability in traffic flows, but for deterministic models this variation takes the form of an entirely prescriptive response to known stimuli. Stochastic models can likewise incorporate responses to observed factors, although this is typically through changes to a probability distribution describing the system. As a concrete example, both deterministic and stochastic models can describe the changes to today's traffic flow pattern due to heavy congestion on some routes on the previous day, though the former model might account for this through changes to the proportions of travellers on the various routes while the latter will change the corresponding probabilities. However, stochastic models have in addition the potential to represent system variation due to unobserved factors, an idea that we expand upon later.

An important class of stochastic day-to-day dynamic models, which are able to take into account the observed mean traffic volumes as well as day-to-day variation, was first introduced by Cascetta (1989), although some of the foundational ideas can be traced back to Horowitz (1984). These types of stochastic process assignment models represent the traffic flows via a Markov process, where the current route flows are modelled as a function of previous traffic conditions. That is, they are based on the Markovian assumption that knowledge of the preceding $m$ states effect traffic flows on a given day, independent of traffic conditions prior to the $m$ days before the current state.

However, as we shall discuss in the present paper, these type of stochastic process models are limited in terms of the types of realistic changes they may reproduce in both the mean and the variance over time. To some degree they still suffer from some of the limitations of equilibrium analyses, in that while they permit variation they are still wedded to the concept of 'stationarity' which is defined as the mean, the variance as well as the covariance all remaining constant over time. Clearly, there are a variety of ways in which this assumption may be violated; for representing realistic phenomena, day-to-day dynamic, stochastic process models of this kind are still too constrained to represent all types of non-stationary characteristics possible.

In this paper, our goal is therefore to develop a new class of stochastic process model that is also able to exhibit non-stationary behaviour, such as the kind observed in Figure 4. To this end we aim to combine existing ideas from transportation network analysis (Cascetta-type

Markov models) with developments in statistics in the representation of 'doubly stochastic' processes. In the latter, certain parameters of the model are set to be random variables, and by incorporating such concepts with transportation system analysis we create a hybrid, stochastic, day-to-day dynamic assignment process.

The additional level of stochastic variation in these models might be interpreted as arising from unmeasured factors, like weather conditions or minor traffic accidents. Such factors can generate spatio-temporal correlations in the traffic flow pattern that will not be directly attributable to any causative effect within our model. While we do not attempt to explicitly explain the origins of these correlations, accounting for their existence is important. One reason for this is that any statistical method for model calibration will perform best if we have a good understanding of the patterns of stochastic variation in the data. Also, if we are to use models for transport planning purposes then it is critical to assess system performance not only under mean conditions, but also under the varying flow patterns that may arise due to all factors, whether explicitly modelled or not. That is to say, if we approach the problem of modelling transport networks from a statistical standpoint, then there will always be some explanatory sources of variation that are measured and some that are unmeasured. While the focus to date has been on developing relatively complex day-to-day models for capturing the measured sources, the modelling of the unmeasured sources has to date been relatively crude and simplistic, and this has significant consequences for the statistical performance of the model. Our focus in the present paper, therefore, is on improving the modelling of such unmeasured sources through more sophisticated models of the assumed random error sources. We can predict (by analogy with other statistical systems) that such modelling developments will, in turn, have potentially significant consequences not only for model fit and the explanation of historical patterns, but also for model forecasts and therefore policy decisions.

The structure of the paper is as follows. In section 2 we provide a review of Markovian day-to-day traffic assignment models. In section 3 we present our proposed extension of these type of models, and provide an illustrative example. We perform numerical experiments using this model in section 4. The paper then concludes with general comments on our doubly stochastic day-to-day models, and identifies paths for further research in this area.

## 2 Day-to-day Markov Traffic Models

We consider a transport network with $L$ directed links and $q$ origin-destination (OD) pairs. Let $n_k$ be the number of routes corresponding to OD pair $k$, $(k = 1, 2, ..., q)$, and define $n = \sum_{k=1}^{q} n_k$. Let $\boldsymbol{W}$ be the $n \times q$ path-movement incidence matrix of 0/1 elements, with $w_{rk} = 1$ if route $r$ relates to OD pair $k$, and let $\boldsymbol{A}$ be the $L \times n$ routing matrix of 0/1 elements, where $a_{lr} = 1$ if link $l$ is part of route $r$.

We seek to model the evolution of traffic flows over the network over a sequence of time intervals indexed by $t = 1, 2, ..., T$. These intervals might, for example, be morning peak hour over successive days, and we shall use the terminology 'day' to refer to each interval. Travel demand on each day is described by a $q$-vector $\boldsymbol{\mu}$ of OD flows. These OD flows distribute themselves across the available routes through the network, producing an $n$-vector $\boldsymbol{y}$. The elements of both $\boldsymbol{\mu}$ and $\boldsymbol{y}$ are lexicographically ordered, and the vectors related by $\boldsymbol{W}\boldsymbol{y} = \boldsymbol{\mu}$. We index the route flows by time period to give $\boldsymbol{y}(t)$, but assume that $\boldsymbol{\mu}$ is constant. This is not particularly restrictive, because we can introduce a single dummy route for each OD pair denoting the option of not travelling for each OD pair, thus allowing transport systems with fluctuating demand from day-to-day within our modelling framework.

Let $\boldsymbol{x}$ be the $L$-vector of link flows. For day $t$ this is related to $\boldsymbol{y}(t)$ by

$$\boldsymbol{x}(t) = \boldsymbol{A}\boldsymbol{y}(t). \tag{1}$$

These link flows generate travel costs, according to the equation

$$\boldsymbol{c}(t) = \boldsymbol{c}(\boldsymbol{x}(t)). \tag{2}$$

The corresponding route costs are then given by

$$\boldsymbol{k}(t) = \boldsymbol{A}^T\boldsymbol{c}(t). \tag{3}$$

In modelling the day-to-day dynamics of the system, we assume that the route choices made by travellers on day $t$ depend on previously experienced travel costs. To this end we define a learning process in terms of the disutility (i.e. generalized cost) of routes, $\boldsymbol{d}$. We assume that disutility and route choice processes depend only on properties of the network from the past $m$ time periods, where $m$ is a finite constant. There is little practical loss in making such assumptions, but the theoretical properties of our model become far more tractable because of the resulting Markov structure.

Writing $\boldsymbol{d}(t)$ for the vector of mean route utilities on day $t$, our model is

$$\boldsymbol{d}(t) = h(\boldsymbol{d}(t-1), \boldsymbol{c}(t-1), \ldots, \boldsymbol{d}(t-m), \boldsymbol{c}(t-m), \boldsymbol{\zeta}) \tag{4}$$

for some function $h$ parameterized by vector $\boldsymbol{\zeta}$. The evolution of the route flows is then modelled by

$$\boldsymbol{y}(t) \sim f(.|\boldsymbol{d}(t-1), \boldsymbol{y}(t-1), \ldots, \boldsymbol{d}(t-m), \boldsymbol{y}(t-m), \boldsymbol{\psi}) \tag{5}$$

where $f$ denotes some conditional probability distribution with parameter vector $\boldsymbol{\psi}$. Note that if we define the state vector $\mathbf{s}(t)$ by $\mathbf{s}(t) = (\boldsymbol{y}(t), \boldsymbol{d}(t), \boldsymbol{y}(t-m+1), \boldsymbol{d}(t-m+1)$, then $\{\mathbf{s}(t); t = 1, 2, ...\}$ is a Markov process.

The class of models defined by equations (4) and (5) is quite broad, and incorporates the seminal models of Cascetta (1989). The Markov process defined in this way will be aperiodic and ergodic under rather general conditions. For example, this is assured if we assume that each traveller assigns a non-zero probability to all routes servicing the relevant OD pair. When the process is ergodic it has a unique stationary distribution, and assuming also aperiodicity (which will hold for any practical traffic model) the marginal probability distribution of $\mathbf{s}(t)$ will converge in time to this stationary distribution, regardless of the initial state of the system. This property also implies that we may generate unbiased estimates of features of the stationary distribution from a *single* pseudo-random realisation of the process. See Cantarella and Cascetta (1995) for related comments, and Davis and Nihan (1993) for an analysis of such models in the asymptotic case, as demands become large in tandem with the network capacity.

Of course, even when the marginal distribution of the system is stationary, the traffic flow patterns will display a complex pattern of temporal dependence. Nonetheless, stationarity dictates that the spatio-temporal correlation structure is constant. Since the link flows on any given day are a deterministic function of the state vector $\mathbf{s}(t)$, it follows that these models will be unable to reproduce the kind of temporal change in variance that was observed in our third numerical example in the Introduction.

**Illustrative example**
A natural starting point for the examination of time-homogoneous models are a subclass of Markovian models as described in Hazelton and Watling (2004). We denote by $\boldsymbol{c}(\boldsymbol{x})$ the d-vector

of travelling costs for each of the links in the network. We suppose that the cost functions $c_i(\cdot)$ for each link $i$ are quadratic, parametrised as $c_i(\boldsymbol{x}) = a + (x_i/b)^2$ so that $b$ is proportional to link capacity. The traveller learning process is based on linear filters of past costs with exponentially decreasing weights. That is, we assume in our route choice model that the measured disutilities for day $t$ (based upon the states of the transport system up to and including day $t-1$) is given by

$$\boldsymbol{d}(t-1) = s(\lambda)^{-1} \sum_{j=1}^{m} \lambda^{j-1} \boldsymbol{k}(t-j), \tag{6}$$

where $s(\lambda) = \sum_{j=1}^{m} \lambda^{j-1}$ for $0 < \lambda < 1$ and $\boldsymbol{k}(t-j)$ is the n-vector of route costs as defined in Equation 3.

The parameter $\lambda$ measures the degree to which traveller's previous experiences impact on travel decisions. A high value means that all past experiences are fairly evenly taken into account, while a value of $\lambda$ close to zero implies that traffic conditions in the past are hardly if at all considered.

On any given day each traveller selects the route with smallest personal disutility. We implement the route-choice mechanism via a standard random utility model, the logit model. In practice one might work with an adaption that overcomes the overlapping routes problems for the logit model (e.g. Cascetta et al. 1996), but for our present intents and purposes we prefer to use the simplest version of this model. We define $p\big(i|\boldsymbol{d}(t-1)\big)$ to be the probability of a traveller choosing route $i$ when travelling between a given OD pair, conditional on the $n$-vector of measured disutilities as

$$p\big(i|\boldsymbol{d}(t-1)\big) = \frac{e^{-\theta d_i}}{\sum_{r \sim i} e^{-\theta d_r}} \tag{7}$$

where $r \sim i$ indicates that routes $r$ and $i$ serve the same OD pair and $\theta$ is a measure of the degree to which travellers react to current conditions after changes in the system. We assume that travellers operate independently conditional on the disutilities, so that the number of trips on each of the routes follow a multinomial distribution. Specifically, if $\boldsymbol{y}_{\{\ell\}}$ and $p_{\{\ell\}}$ are respectively the vectors route flows and route choice probabilities for OD pair $\ell$, then $\boldsymbol{y}_{\{\ell\}}$ follows a $\mathrm{Mn}(\mu_\ell, p_{\{\ell\}})$ distribution conditional on the disutilities.

## 3 An Extended Class of Models

So far as we are aware, all day-to-day dynamic models that have appeared in the literature to date are members of the models defined by equations (4) and (5). There is considerable flexibility within these models for representing the processes by which travellers learn from past travel experiences, and how these experiences influence future route choice. Nevertheless, despite the apparent scope offered by these models, they are incapable of reproducing the kind of non-stationary behaviour seen in Figures 3 and 4.

The issue of being unable to reproduce non-stationary processes is well known in the time series and spatial-point process literature in Statistics. In those areas, a variety of models have been developed to address this matter. These include Generalised Autoregressive Conditional Heteroscedasticity (GARCH) models in financial time series modelling (Thavaneswaran et al. 2006, Bollerslev 1986), and Cox processes in spatial statistics (Møller and Waagepetersen 2007, Brix and Diggle 2001). The essential idea behind these methods is to exchange some of the fixed parameters in the models with random variables that are allowed to change over time (as is the case for GARCH models) and/or space (as for Cox process models). Such models are often termed 'doubly stochastic' to reflect the randomness in the parameter vectors and the random behaviour of the process conditional on those parameter vectors.

Motivated by such models, we propose generalising Cascetta-type traffic models through replacement of some elements of the parameter vectors $\boldsymbol{\zeta}$ and $\boldsymbol{\psi}$ (from equations 4 and 5 respectively) by random processes. This provides a means of replicating the kind of non-stationarity experienced in some of our observed data examples, even when information on the underlying cause is not known. In principle one can postulate explanations as to what variations in the parameter vectors might represent. For example, autocorrelated temporal variation in the parameters describing traveller learning and route choice might be interpreted as responses to different weather conditions. Given a particular state of the system, perhaps travellers react differently (as described by a different conditional probability distribution) depending on whether it is raining or dry. Nonetheless, such a specific interpretation of the variation in model parameters is not necessary, and indeed may not be helpful. It is sufficient to accept that the properties of the system will differ according to a host of *unmeasured* factors.

There is unlimited scope in the manner in which the parameters $\boldsymbol{\zeta}$ and $\boldsymbol{\psi}$ might vary. In order to impose some structure, we suggest that these vectors be modelled as Gaussian random processes. That is, at any spatial location $\boldsymbol{z}$ and time $t$ the vectors are a realization of a joint random process $\{(\boldsymbol{\zeta}(\boldsymbol{z},t)^\mathsf{T}, \boldsymbol{\psi}(\boldsymbol{z},t)^\mathsf{T})^\mathsf{T}\}$. This leads to a doubly stochastic model, analogous to a Cox process, which can allow for patterns of spatio-temporal correlations in both the utilities $\boldsymbol{u}$ and traffic flow vectors $\boldsymbol{y}$.

**Illustrative example continued**

We present a specific modification of the Cascetta-type model described in the illustrative example previously in order to incorporate variation in the temporal dependence structure. We redefine the parameter $\theta$, that is, the measure of sensitivity, in equation 7 as

$$\log\big(\theta(t)\big) = \omega(t) + \nu \tag{8}$$

at time point $t$ where $\nu$ is a constant. The values of $\theta$ are kept on a log scale in order to ensure non-negativity. We can then describe the parameter $\omega$ as a first order autoregressive (AR(1)) process

$$\omega(t) = \phi\omega(t-1) + \epsilon(t) \tag{9}$$

where $\epsilon_t \sim \mathcal{N}(0, \sigma^2)$. The marginal distribution of $\omega$ is then given by

$$\omega(t) \sim \mathcal{N}\Big(0, \frac{\sigma^2}{1-\phi^2}\Big). \tag{10}$$

On the natural scale this implies that the mean of $\theta$ is given by

$$E(\theta) = \exp\{\nu\}\exp\left\{\frac{1}{2}\text{var}(\omega)\right\} = \exp\{\nu\}\exp\left\{\frac{\sigma^2}{2(1-\phi^2)}\right\}. \tag{11}$$

The full model is now defined by equations (4), (5), (6) and (7) (which define the system conditionally given any value of $\theta$) and equations (8) - (10).

In this illustrative example, and indeed more generally, the random process for the parameters is itself stationary. It follows immediately that the marginal process for the state vector $\mathbf{s}(t)$ is also stationary. This might at first seem to contradict our desire to replicate the kind of non-stationary behaviour observed in our real data examples. However, it is critical to recognize that our doubly stochastic models are non-stationary *conditional* on any (non-trivial) realization of the random process on the parameters. This means that our models can reproduce stochastic trends and changes in volatility in a manner which is impossible when the parameter vectors are fixed. We illustrate this through synthetic numerical examples in the next section.

# 4 Empirical study of Modelling Non-stationarity

We investigate properties of traffic flows for the illustrative Markov assignment model that was introduced in section 2 and extended in section 3. We make direct comparisons between: (a) traffic flows simulated under the classic Cascetta-type Markov model with a fixed level of sensitivity $\theta$ and (b) the characteristics of traffic flows that result from a modified version of the same model, where the parameter $\boldsymbol{\theta}$ is instead a random variable. For several variations of the parameters in the model we examine the extent to which the models can reproduce the general characteristics seen in the real traffic flows.

The results in this section are intended to illustrate the range of stochastic behaviour in the models. We have made no attempt to model variation due to measured factors like day of the week, for example. For real-world modelling it would be important to explicitly account for such effects, but for illustrative purposes it is helpful to focus primarily on the patterns of random variation.

Computation was doing using the software R (R Core Team 2013) running on a MacBook laptop with 1GB memory and a 2GHz processor. Results on the simple network from section 4.1 required approximately 0.25 seconds of CPU time per simulation run. Each simulation for the large network in section 4.2 took about 2.5 seconds of CPU time.

## 4.1 Temporal variability

We begin our analysis with a simple two-zone network with only one interzonal movement, in which the zones are connected by two parallel non-overlapping routes. We ran simulations for a range of combinations of three factors, first, the level of inertia of travellers, second, the degree to which travellers learn from their experiences over time and thirdly the memory length (i.e. the number of days taken into consideration by travellers when making route choices). In general, we found that memory length was not very influential in terms of the resulting simulated link flows, other than through the smoothing effect of averaging over a larger number of days. That is, if experiences up to 20 days in the past were taken into account, then we found that noticeable changes over time were only produced for relatively high levels of sensitivity.

Some particularly interesting cases are considered here, where the contrast between the more stationary-looking flows for the fixed $\theta$ case and the more non-stationary appearance of the flows for the random $\theta$ case is very obvious. In Figure 5 we see simulated link flows on one of the two links, when the sensitivity to costs is relatively high ($\theta = 1$) and the users take into account experiences from the past two days with very little consideration of less recent events ($\lambda = 0.01$). The flows show no signs of changes over time, the mean remains around 35 and the variance appears constant. The autocorrelation plot in the second panel of Figure 5 shows no strongly significant autocorrelations.

In Figure 6 we see flows from the modified day-to-day model where we now allow the parameter $\theta$ to vary over time, according to equation (8). Here the parameters $\phi$ and $\sigma$ governing the random process $\{\theta(t)\}$ have been selected so that the mean value of $\theta$ (as given in equation (11)) matches the fixed value $\theta = 1$ considered above. We see a change in both the mean flows over time as well as obvious heteroscedasticity. In particular the manner in which the link flows display a change in variability is similar to what we see in Figure 4, albeit in a more extreme fashion.

In the second panel of Figure 6 we can see a plot of the realized sequence $\{\theta(t)\}$ over time.

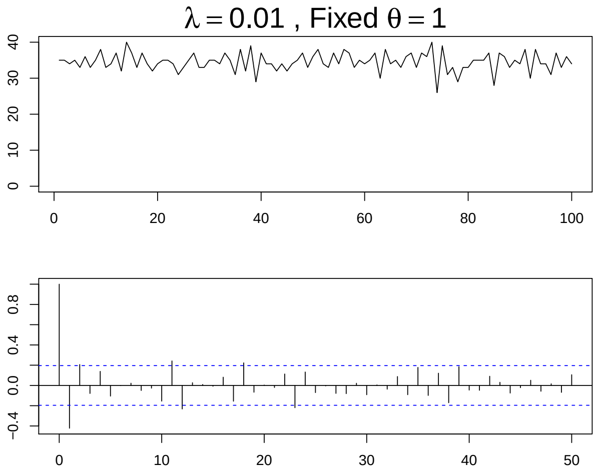


Figure 5: Simulated Time series and associated autocorrelation plot for traffic counts on two-node network with memory length 2.

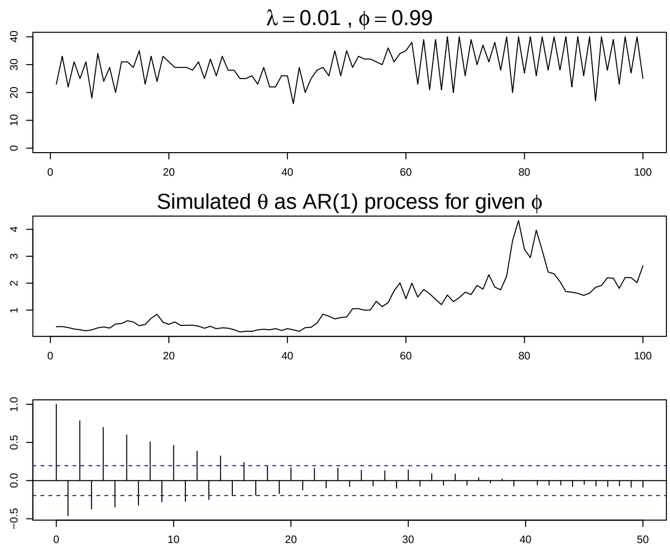


Figure 6: Simulated Time series and associated autocorrelation plot for traffic counts on two-node network with memory length 2.

For this particular instance of simulated $\theta$ values, we see that just after time point 60 the level of sensitivity is below 2, and that shortly after it surpasses this threshold the behaviour of the function of the link flows changes. For smaller values of sensitivity the corresponding link flows shows smaller oscillations and a slightly lower mean, then later on as the simulated $\theta(t)$ approaches higher values, the average link flows increase as well as the magnitude of changes in the flow levels overall.

We now consider the case that the users take into account experiences from the past two days, where the value of $\lambda$ is now 0.8, that is, all previous experiences carry a noticeable weight in the route choice process, while keeping the level of sensitivity to costs experienced high at $\theta = 1$. The flows look very similar to the flows we see in Figure 5 and again the autocorrelation plot displays no significant lags. Changes in the values of the parameter $\lambda$, as a measure of the extent to which previous experiences are taken into account, do not take much effect in this case.

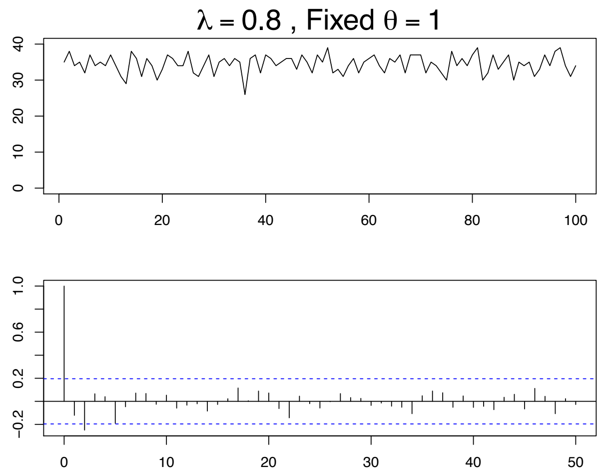


Figure 7: Simulated Time series and associated autocorrelation plot for traffic counts on two-node network with memory length 2.

We now introduce a random process for $\theta$, calibrated (using $\phi$ and $\sigma$) to have mean $\mathsf{E}(\theta) = 1$. As before we can see in Figure 8 that by letting the parameter $\theta$ be a random process we introduce heteroscedasticity of the type observed in the Bronx traffic flow data in Figure 4.

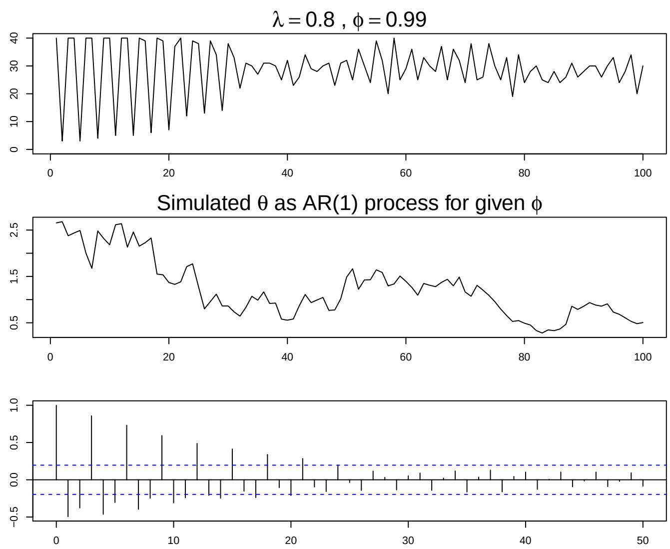


Figure 8: Simulated Time series and associated autocorrelation plot for traffic counts on two-node network with random $\theta$.

Both these artificial examples demonstrate well to what extent the kinds of non-stationary observed in real traffic data can potentially be reproduced using our doubly stochastic models.

### 4.2 Temporal-spatial correlation structure

So far, we have focused on demonstrating the kind of changes in simulated traffic flow patterns we see over time when allowing parameters in our model to be random rather than fixed. In this section, the aim is to show that these kind of temporal changes occur in larger networks as well, but more importantly, spatial correlation structures can also be shown to change when we apply our extended models. Our study uses a 21 node network based on a section of the road system in the English city of Leicester. An abstracted version of this network is depicted in Figure 9. As before, our comparisons between fixed and random processes for $\theta$ are calibrated so that the mean value of the latter matches the former.

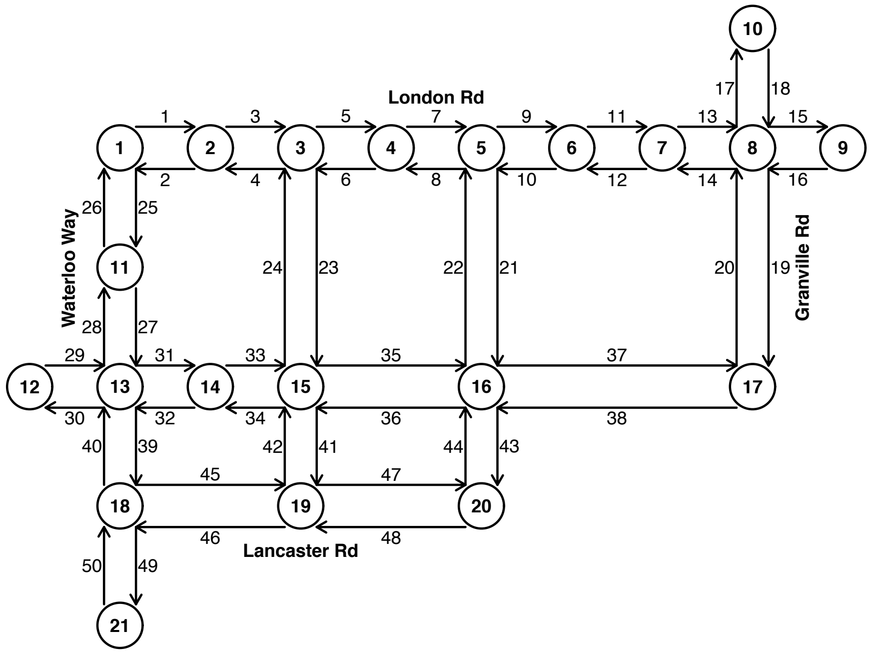


Figure 9: Test network taking from the English city of Leicester.

For brevity we present the flows on only a selected number of the road sections, specifically the links 22, 31, 36, 38, 42 and 46. When contrasting the link flows in Figure 10, we can again see the familiar feature of stationary flows in the fixed $\theta$ case, and changes in the pattern of the oscillations in the flows over time when $\theta$ is a random variable.

This same phenomenon is something we observed for a variety of different values in key parameters, for example, adjusting the length of memory, $m$, to longer periods of time delivers similar results. This demonstrates that this effect is a genuine one, in that the extended models, in a similar manner to their counterpart Cox processes, can allow for a change over time in the link flow patterns, in contrast to the standard models.

We now examine the potential spatio-temporal correlation patterns that may arise on a set of interconnected links in the Leicester example. One reason this may be of interest is rooted in the following: from a transport planning point of view if a section of road appears to be having problems with the volumes of demand, a reasonable solution may be to widen the road in response. This alteration of the traffic system typically has a broader effect. A transport planner may well expect smoother flows on the road section in question to lead to better flow patterns on other parts of the network in the vicinity, and what is more, might imagine that this effect would persist under most, if not all, network conditions. An implicit assumption being

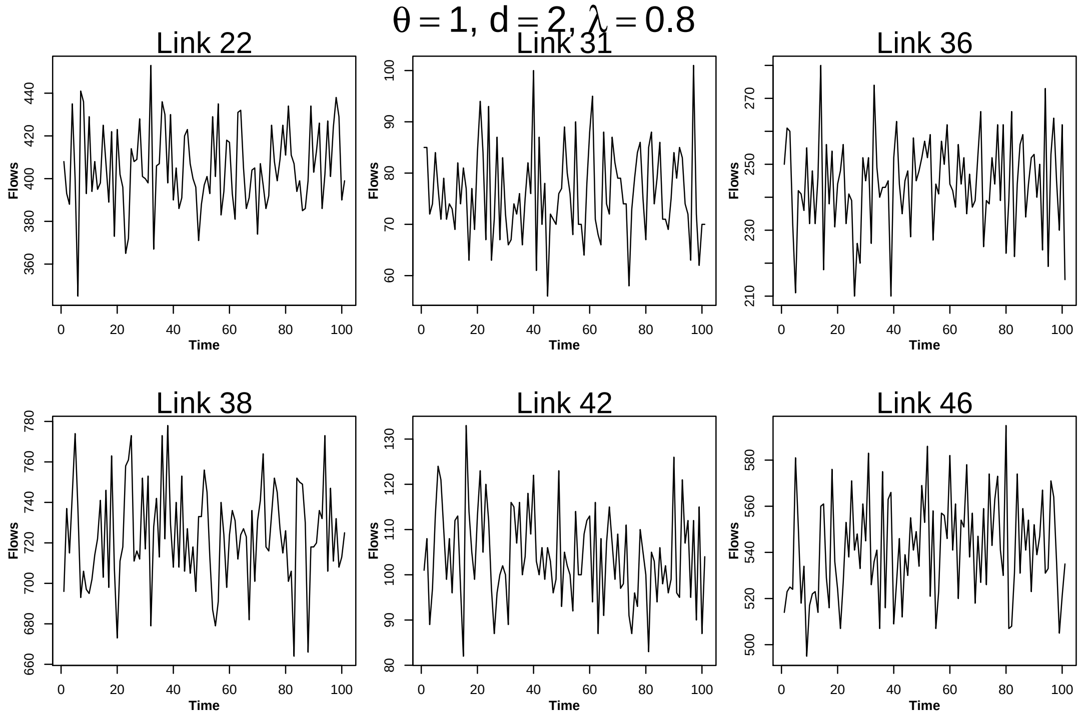


(a) Fixed $\theta$ case.

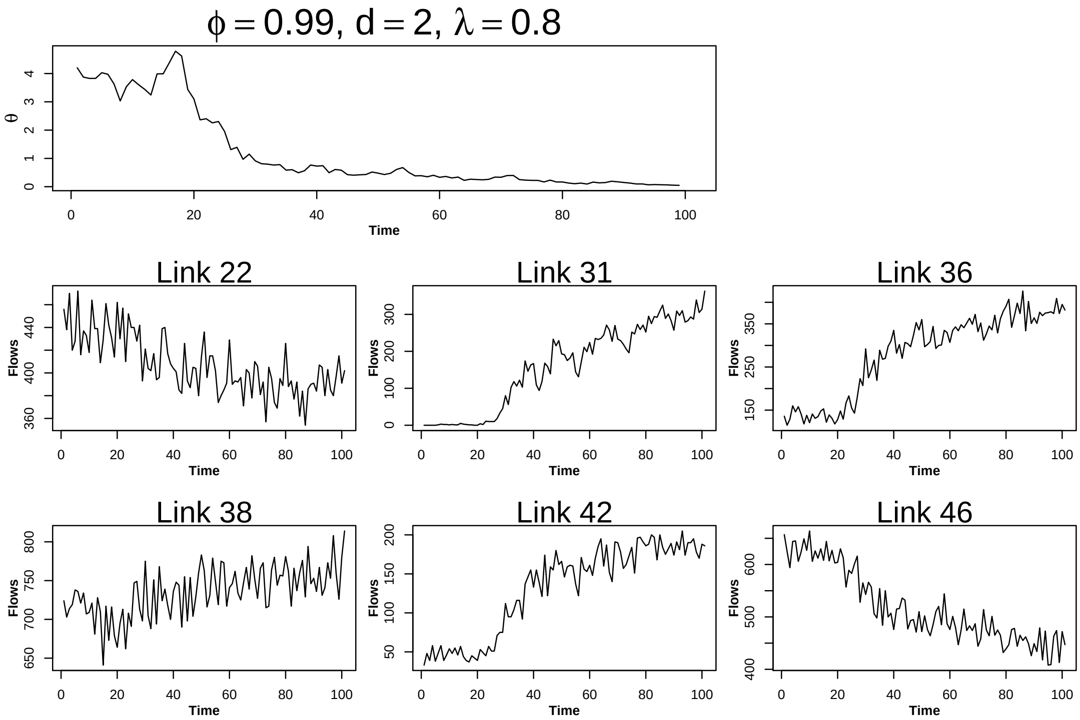


(b) Random $\theta$ case.

Figure 10: Simulated Time series for traffic counts on Leicester network.

made in the latter stages of that argument is stationarity, whereby the spatial relationships between link flows in the network are constant over time. However, if correlation patterns do happen to vary temporally then travellers might find themselves in the situation that the effects of the imagined road-widening scheme are more difficult to predict.

To investigate this aspect we calculate the correlation in a similar fashion to a moving average. That is, in Figures 11 and 12 the first value that is plotted is the correlation between the link flows from days 1 to 100 and the link flows from days 2 to 101. The next value is the correlation between the time period 2 to 101 and the time period 3 to 102, and so and so forth. This results in a sequence of points that form a line, if there is no particular correlation structure. For fixed $\theta$, as in Figure 11, none of the link pairs exhibit correlation structures that change over time, as theory predicts. However, when we allow $\theta$ to vary, as in Figure 12, some of the link flow pairs do display a noticeable change in the correlation structure over time.

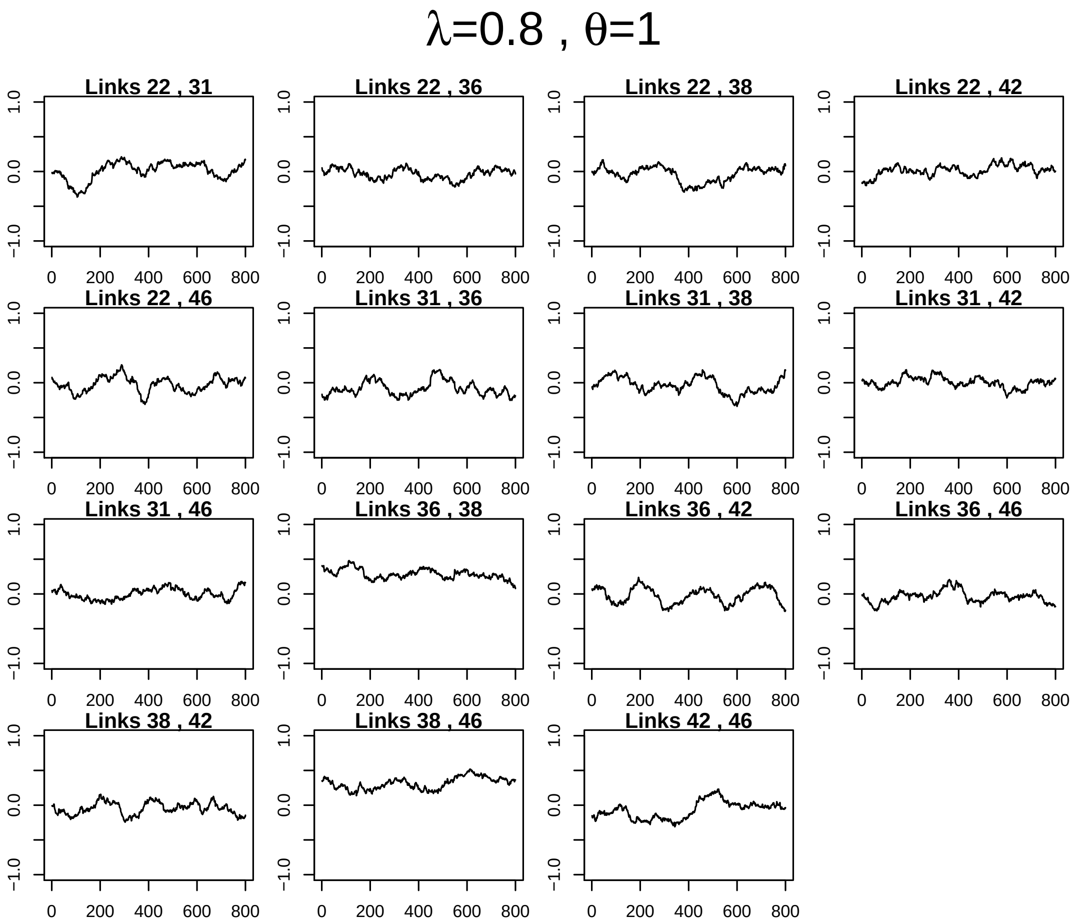


Figure 11: Spatial correlations in eight-node network for fixed $\theta$.

## 5 Discussion

Non-stationarity can be observed in day-to-day traffic flow data from real world transport systems. However, existing day-to-day traffic assignment models such as those proposed by Cascetta (1989) are by nature stationary, and so cannot faithfully represent such systems. In this paper, we describe how such Markovian day-to-day dynamic traffic assignment models can be extended in a straightforward manner by replacing a subset of the fixed parameters in the Markov model with random processes. The resulting models are analogous to Cox process models. They are conditionally non-stationary (given any realization of the parameter processes) and hence can reproduce features such as heteroscedasticity in traffic flows. In particular, our numerical examples were able to mimic some of the features of link flow variation

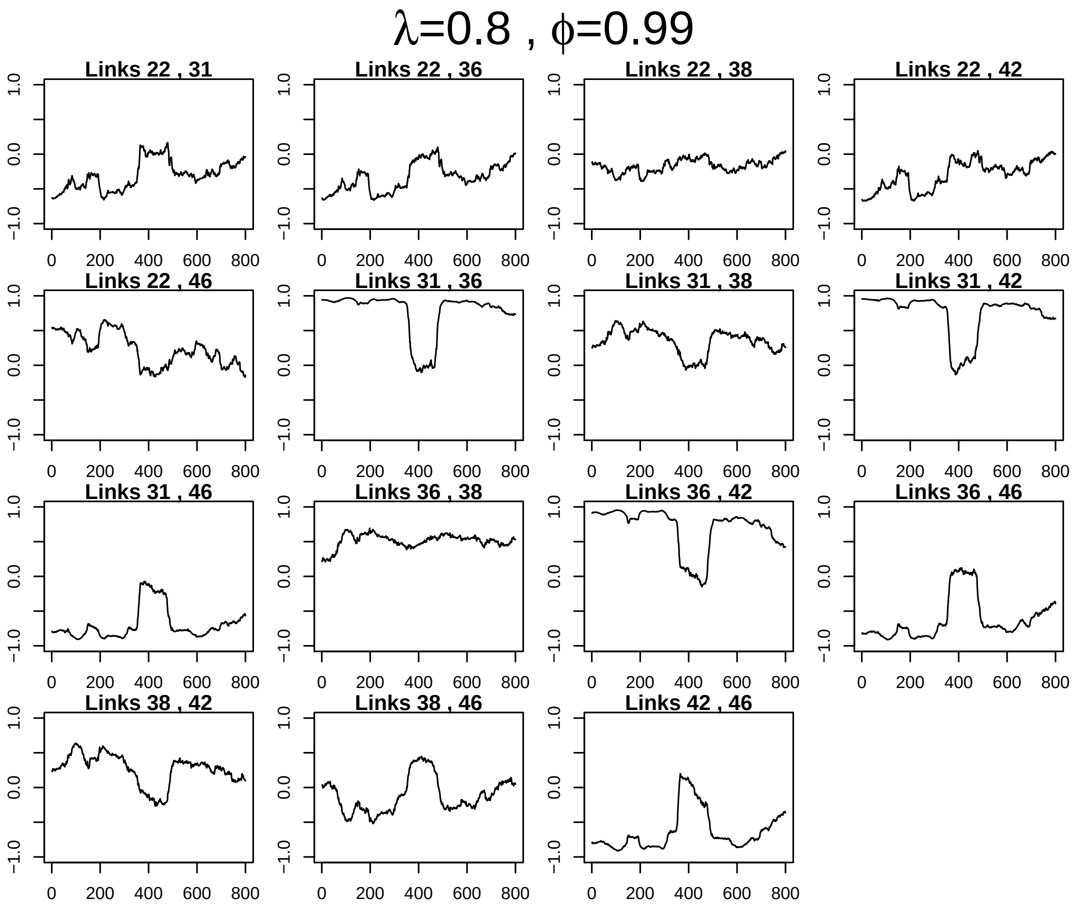


Figure 12: Spatial correlations in eight-node network for random $\theta$.

observed in Figures 3 and 4.

Our illustrative model and numerical examples focused on cases where model parameters were random processes in time, but were constant across space. In part this is a response to the nature of our motivating example. Moreover, focusing on models with just time variation in the parameters simplified our exposition of doubly stochastic models. As an extension to our proposed model, we may also consider adding in spatial random variation in model parameters, and in such a case we need to decide how the value of a parameter in space will relate to the route choice model. For example, suppose that we wish the logit parameter $\theta$ from our illustrative example to vary across the network. Which specific value applies when computing probabilities according to equation (7)? There are many possibilities, including a stochastic integral of the random process across the spatial path. Rather more simply, we might define the value of the random process only at the network nodes, and then allocate the value from the appropriate origin node to each traveller. In that case a Gaussian spatio-temporal model for the model parameters would be

$$\begin{bmatrix} \boldsymbol{\zeta} \\ \boldsymbol{\psi} \end{bmatrix} \sim \mathcal{N}(\mathbf{0}, \Sigma) \tag{12}$$

where $\boldsymbol{\zeta}$ and $\boldsymbol{\psi}$ now denote the values of the process at all required times and nodes. In other words, the former vector is the concatenation of vectors $\{\boldsymbol{\zeta}(\ell, t)\}$ where $\boldsymbol{\zeta}(\ell, t)$ gives the values at node $\ell$ and time point $t$. The covariance matrix $\Sigma$ in equation (12) in principle allows for huge flexibility in the form of spatio-temporal correlations. In practice it will be necessary to impose some highly parsimonious parameterization.

Finally, we note that our doubly stochastic models can be regarded as describing the effects of unobserved covariates. However, cyclical variation due to day-of-the-week effects may be concerned with factors that are *observable*. Such properties of the system can in principle be handled by incorporating fixed covariates into the model, although we are not aware of any work in this direction in the literature for dynamic day-to-day models. Thus an interesting future direction would be to consider extended models incorporating not only the random (unobserved) factors that we have considered, but also some fixed (observed) covariates where that information is available.

## References


T. Bollerslev. Generalised autoregressive conditional heteroskedasticity. *Journal of Econometrics*, 31:307–327, 1986.

A. Brix and P. J. Diggle. Spatiotemporal prediction for log-Gaussian Cox processes. *Journal of the Royal Statistical Society. Series B (Statistical Methodology)*, 63(4):823–841, 2001.

G.E. Cantarella and E. Cascetta. Dynamic processes and equilibrium in transportation networks: Towards a unifying theory. *Transportation Science*, 29:305–329, 1995.

E. Cascetta. A stochastic process approach to the analysis of temporal dynamics in transportation networks. *Transportation Research Part B*, 23:1–17, 1989.

E.A. Cascetta, F. Nuzzolo, A. Russo, and A. Vitetta A modified logit route choice model overcoming path overlapping problems: specification and some calibration results for interurban networks In B. Lesort (ed.) *Transportation and Traffic Theory. Proceedings from the Thirteenth International Symposium on Transportation and Traffic Theory*, Lyon, France, Pergamon, pp. 697–711, 1996.

C.F. Daganzo and Y. Sheffi. On stochastic models of traffic assignment. *Transportation Science*, 11:253–274, 1977.

G.A. Davis and N.L. Nihan. Large population approximations of a general stochastic traffic assignment model. *Operations Research*, 41(1):169–178, 1993.

M.L. Hazelton and D.P. Watling. Computation of equilibrium distributions of Markov traffic assignment models. *Transportation Science*, 38:331–342, 2004.

J.L. Horowitz. The stability of stochastic equilibrium in a two-link transportation network. *Transportation Research Part B* 18(1):13-28, 1984.

J. Møller and R.P. Waagepetersen. Modern statistics for spatial point processes. *Scandinavian Journal of Statistics*, 34:643–684, 2007.

R Core Team. R: A Language and Environment for Statistical Computing. R Foundation for Statistical Computing, Vienna, Austria. ISBN 3-900051-07-0, URL `http://www.R-project.org/`, 2013

A. Thavaneswaran, S. S. Appadoo, and C. R. Bector. Recent developments in volatility modeling and applications. *Journal of Applied Mathematics and Decision Sciences*, 2006:1–23, 2006.

J.G. Wardrop. Some theorectical aspects of road traffic research. *Proceedings of the Institute of Civil Engineers*, Part II(1):325–378, 1952.

D.P. Watling and M.L. Hazelton. The dynamics and equilibria of day-to-day assignment models. *Networks and Spatial Economics.*, 3:349–370, 2003.