Regular Article

Detecting contagious spreading of urban innovations on the global city network

Niklas H. Kitzmann[1,2,a] , Pawel Romanczuk[3,4,7] , Nico Wunderling[1,5,6] , and Jonathan F. Donges[1,5]

[1] Earth System Analysis, Potsdam Institute for Climate Impact Research, Member of the Leibniz Association, 14473 Potsdam, Germany
[2] Institute for Physics and Astronomy, University of Potsdam, 14476 Potsdam, Germany
[3] Institute for Theoretical Biology, Department of Biology, Humboldt University of Berlin, 10115 Berlin, Germany
[4] Bernstein Center for Computational Neuroscience, Humboldt University of Berlin, 10115 Berlin, Germany
[5] Stockholm Resilience Centre, Stockholm University, 10691 Stockholm, Sweden
[6] High Meadows Environmental Institute, Princeton University, Princeton, NJ 08544, USA
[7] Germany Science of Intelligence, Research Cluster of Excellence, Marchstr. 23, 10587 Berlin, Germany

Abstract Only a fast and global transformation towards decarbonization and sustainability can keep the Earth in a civilization-friendly state. As hotspots for (green) innovation and experimentation, cities could play an important role in this transition. They are also known to profit from each other's ideas, with policy and technology innovations spreading to other cities. In this way, cities can be conceptualized as nodes in a globe-spanning learning network. The dynamics of this process are important for society's response to climate change and other challenges, but remain poorly understood on a macroscopic level. In this contribution, we develop an approach to identify whether network-based complex contagion effects are a feature of sustainability policy adoption by cities, based on dose-response contagion and surrogate data models. We apply this methodology to an exemplary data set, comprising empirical data on the spreading of a public transport innovation (Bus Rapid Transit Systems) and a global inter-city connection network based on scheduled flight routes. Although our approach is not able to identify detailed mechanisms, our results point towards a contagious spreading process, and cannot be explained by either the network structure or the increase in global adoption rate alone. Further research on the role of a city's abstract "global neighborhood" regarding its policy and innovation decisions is thus both needed and promising, and may connect with research on social tipping processes. The methodology is generic, and can be used to compare the predictive power for innovation spreading of different kinds of inter-city network connections, e.g. via transport links, trade, or co-membership in political networks.

1 Introduction

Anthropogenic climate change and other human-made pressures on the environment are threatening the civilization-friendly Holocene state that the Earth system has been existing in for thousands of years [1,2]. A world-wide transformation of the economic, societal and political world towards sustainable practices and technologies is urgently needed to mitigate these effects [3]. Cities may play an important role as active agents in this global sustainability transformation, with their potential for impactful change mediated by a number of factors: Precisely because some of the largest drivers of planetary boundary transgressions are located in cities, their mitigation potential is high. For example, the urban population accounts for over 80% of global greenhouse gas emissions [4], and has large freshwa-ter [5] and chemical pollution footprints [6]. Already, over half of the global population lives in cities, a share that is projected to rise to two thirds by the middle of the 21st century [7]. These urban populations are also among the most threatened by negative consequences of environmental changes, such as flooding events [8] and heat waves [9]. Thus, if the challenge of remaining within the planetary boundaries can be solved, it must be solved in the urban context. At the same time, cities are uniquely suited to address these issues. As centers of knowledge and innovation [10,11], and possessing significant economic resources [12], innovative solutions are most likely to originate here. This is especially true for sustainability innovations, where so-called frontrunner cities have shown how experimentation with sustainability can create positive inertia for change [13–15].

To become relevant for the global sustainability transformation, locally conceived and implemented innovations must spread world-wide. As part of the

[a] e-mail: kitzmann@pik-potsdam.de (corresponding author)

human response to anthropogenic environmental changes, the spreading of sustainability innovations thus represents an important Earth system [16] process. Innovation and policy transfer has been extensively studied in the social sciences [17–20], primarily in case studies investigating individual cities or small city networks [21–23]. While this focus on individual circumstances is surely well-placed for studying such complex and diverse processes, we argue that it would be well complemented by data- and model-based studies of the macroscopic global dynamics of urban innovation spreading. It has been reported earlier that universal principles of growth, innovation and sustainability can be found by applying statistical physics to cities, leading to scaling laws and revealing the pace of urban life ("Urbanocene") [11,24–26]. Our approach is thus motivated by the complexity science of social physics and complex networks [27–29].

In this contribution, we develop a method for investigating the proliferation of urban sustainability innovations on a global scale by viewing cities as complex interacting systems [27]. Specifically, we hypothesize a complex contagion process [30,31], implying a nontrivial relationship between the probability of a city to implement an innovation and its exposure(s) to it from other cities. Throughout this work, we use terms such as "infection" and "contagion" to describe the innovation spreading process, in line with earlier literature [32]. In this context, the "infection" of a city should simply be understood as the adoption of the studied innovation by a city, without negative connotations or the inference of passivity. Likewise, "contagion" only implies that this process may be directly influenced by other cities which have already implemented the innovation, and which are connected to the original city in some way.

Networked spreading and contagion processes have been studied in many different scientific subjects, such as epidemics [33], cascading failures [34], and the formation and spreading of social norms, opinions and behaviors [35–37]. The spreading of social, political and technological innovations relevant for sustainability transitions and rapid decarbonisation have also been identified as a promising approach for understanding the emergence of social tipping points in this context [3,38–40].

We use a set of example data sets to develop and demonstrate our method. As a proxy for the inter-city links facilitating the spreading of innovations, we use the global network of scheduled flight routes. We correlate this (static) network with the spreading of Bus Rapid Transit Systems (BRTs), a public transport innovation which combines features of bus networks and light rail systems [41]. The spreading of BRTs has previously been investigated in case studies [42,43], but never on a global scale.

Detecting contagion in such low-rate spreading processes on a large, high-density network poses a statistical challenge. Because of the small number of total infections, we cannot rely on interpreting the functional shape of the infection rate as in [44]. Instead, our

approach is based on dose response functions (DRFs) [30,45,46], an analytical tool that has been used to study a variety of simple and complex spreading processes, such as the diffusion of information on social media networks [47] and the spreading of health-related behaviors among students [48]. Dose response functions encode the probability of infection of a node, as a function of the exposure (or "dose") received from connected nodes. We modify this measure for a fractional contagion paradigm [37], to describe a city's probability of adopting a new innovation dependent on the fraction of its network neighborhood that has already implemented the innovation. We then develop a hierarchy of surrogate models, successively excluding non-contagion-related mechanisms that may confound the observation of contagion processes. The surrogate model method relies on Monte-Carlo-based, data-derived hypothesis tests to analyze specific data features without prescribing concrete underlying mechanisms. Surrogate models have successfully been used as a tool in exploratory data analyses [49,50], particularly for investigating networked processes [51,52], including epidemic and social contagion [48,53,54], and time series data [55,56]. Partially randomizing the empirical data in line with specific null hypotheses, and then comparing key measurements (here, DRF functional shapes) allows us to investigate correlations found in the data, and their causal relevance.

This paper is structured as follows: In Sect. 2, a description of the data sets used to demonstrate the method is given, along with a motivation for their selection. This is followed by the description of the employed methods in Sect. 3, detailing our use of DRFs (Sect. 3.1) and surrogate models (Sect. 3.2). The results of our analysis are reported in Sect. 4; we discuss them and conclude in Sect. 5.

2 Data

To demonstrate the methodology, we use one example data set each for the city network and the spreading innovation, respectively. The choice for these illustrative data sets is driven by three considerations: the availability of an adequate data set, the plausibility of finding contagious spreading behavior in this data, and the scientific or social relevance of understanding the spreading process of the system itself and its analogues. The data sets used here are the global network of scheduled flight routes for the network component, and the adoption of Bus Rapid Transit System (BRT) public transport innovation for the spreading component.

2.1 Flight route network

The choice of flight route connections for the city network component has several advantages. First, a globally homogeneous data set is available from public sources [57]. This data also includes smaller cities, which are often not the focus of city network research
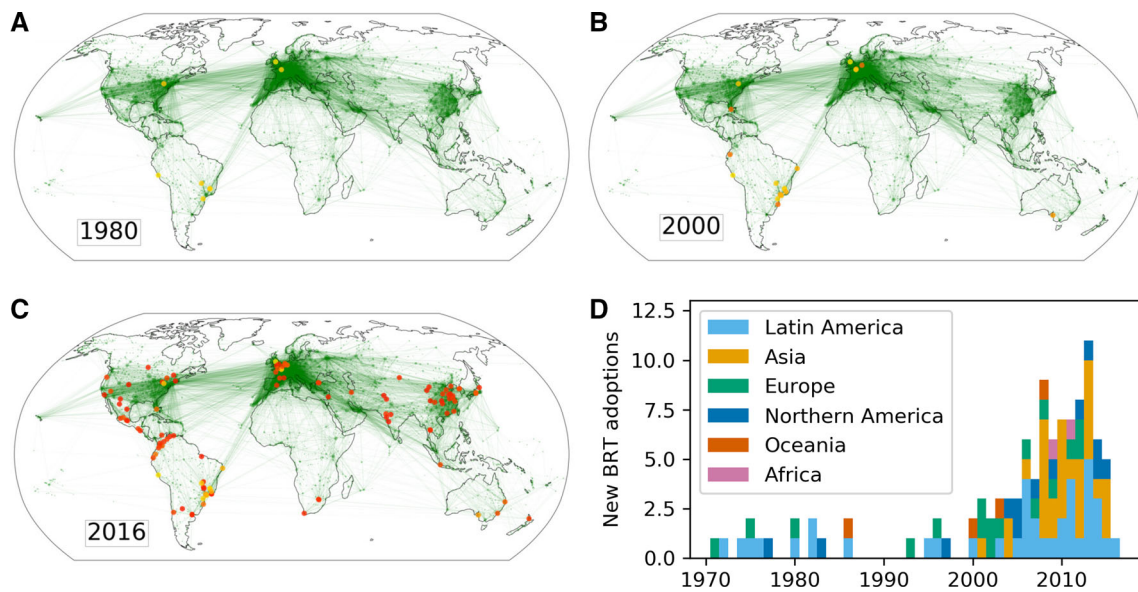
Fig. 1 Visualization of the spreading of Bus Rapid Transit Systems (**A**–**C**) and rate of implementation (**D**). Spread of Bus Rapid Transit System in 1980 (**A**), 2000 (**B**) and 2016 (**C**); the latter two represent the bounds of the time interval investigated here. The date of implementation is displayed on a color scale from yellow (1972) to red (2016). Overlaid in green is the global network of flight routes. (**D**) Implementation rate of Bus Rapid Transit Systems in a stacked histogram, color coded by continent. A marked rise of implementations is apparent after the year 2000, prompting our scrutiny of this time interval. A version of (**D**) resolved on the country scale is presented in Appendix B

[58]. Furthermore, we expect any kind of city-to-city connection, be it on an economic, political, or cultural level, to also produce some amount of flight traffic. The flight route network can thus serve as a plausible proxy for any underlying inter-city linkages.

We source the network data from a publicly available data set [57], which contains information on airports and scheduled routes that are visualized in Fig. 1. We correlate this information with data on city locations and population sizes [59], to transform the airport-to-airport direct route network into an undirected, weighted city-to-city network. The exact algorithm for calculating a city's connection strength to another is described in Appendix A. Only cities with a population of greater than 60 000 are considered here, corresponding to the lowest population threshold which includes almost all cities which have implemented a Bus Rapid Transit System (see Sect. 2.2). As the flight route data source provides a snapshot of the global flight route network dated to 2014, we assume the network to be static. Limitations imposed by this choice are discussed in Sect. 5.

2.2 Bus rapid transit systems

We choose the implementation of Bus Rapid Transit Systems (BRT) as the spreading innovation component of the illustrative analysis. BRTs are a public transport innovation, first developed in the early 1970's [60]. Combining a number of measures such as dedicated bus lanes, frequent service with timed transfers, off-board fare collection, and preferential intersection treatment,

BRTs are frequently compared to light rail networks [41]. Often representing a cost-effective way of implementing a high-quality public transport network for cities, they can play an important role in shifting the modal share towards environmentally friendlier means of transport [61].

A comprehensive database of BRT implementations is jointly maintained and publicly provided by the BRT + Centre of Excellence and EMBARQ, the WRI Ross Center for Sustainable Cities signature initiative for sustainable transport [60]. Only implementations rated "Bronze", "Silver" or "Gold" by these organizations are considered in this analysis, in order to exclude systems that only share a limited amount of features with full BRT implementations. The global implementation rate of BRTs is displayed in Fig. 1; an alternative version resolved to the country level can be found in Appendix B. Following several decades of low adoption rates, a marked increase can be observed after the year 2000. To better understand this phenomenon, and exclude times of low activity that may drown out any potential contagion effects in the data, we focus on this "epidemic" phase of rapidly rising implementations between the years 2000 and 2016. At the beginning of this period, BRTs are already present on four continents (Fig. 1B). The data is considered in a time-stepped fashion, with time step length $t = 1$ year.

3 Methods

We use a dose-response-contagion approach to investigate contagion effects, described in Sect. 3.1 To differentiate between true contagion effects and confounding factors like homophily and shared environments, we use a hierarchy of surrogate data sets. This allows for a search for evidence of causal contagion effects by excluding alternative hypotheses, and is described in Sect. 3.2.

3.1 Dose–response functions

Dose-response functions (DRFs) are a useful tool in characterizing contagion effects on networks [45,46,48]. They represent the functional dependence of a node's infection probability p_{inf} on the exposure "dose" I from neighboring infected nodes. Depending on the underlying contagion process, DRFs can have different functional forms, such as smooth, sigmoidal curves, and even sharp step-like functions for threshold-based contagion processes [46]. As a measure of the "dose" received by a city, we define the infection pressure $I_i(t)$ experienced by a city i,

$$I_i(t) = \frac{\sum_{j=1}^{N_i} w_{ij} s_j(t)}{\sum_{k=1}^{N_i} w_{i,k}} . \tag{1}$$

Here, N_i represents the number of cities connected to city i, and w_{ij} holds the weight of the connection between cities i and j. The infection status $s_j(t)$ is 1 or 0 if city j is infected or uninfected at time t, respectively. This definition represents a fractional contagion paradigm, a type of complex contagion [37] that features the inhibition of infection probability by non-infected neighbors. It is motivated by the conjecture that cities with a high network degree should be less likely to adopt the innovation after an exposure of a certain strength, than cities that have few or low-weight other connections. In this way, high numbers of connections to non-infected cities may "drown out" the effects of infected neighbors. Vice versa, exposures are more effective if they are experienced by a city with a low degree of connectivity. This also solves the problem of high-degree nodes, such as the air traffic hub London, which would otherwise receive very high exposures, but cannot plausibly be expected to become infected much more readily. The system thus retains the intuitive assumption that, for an illustrative example of two connected cities with strongly differing connectivities, an infection is more likely to jump from the high-degree city to the low degree city than vice versa. This is despite the weight of the connection remaining symmetric.

From all cities' time series $I_i(t)$, we compute the total distribution of infection doses $D(I)$. Likewise, we compute the distribution of "successful" infection doses $C(I)$, which is made up of only those infection doses $I_i(\hat{t}_i)$ that were received by cities i which became

infected at $\hat{t}_i + 1$. With these two distributions, we can compute the DRF: the infection rate per time step $r_{\text{inf}}(I)$ as a function of experienced infection pressure:

$$r_{\text{inf}}(I) = \frac{C(I)}{D(I)} . \tag{2}$$

The empirically determined DRF $r_{\text{inf}}(I)$ is used as an estimator of the true probability of infection $p_{\text{inf}}(I)$ after exposure to the infection pressure I: $p_{\text{inf}}(I) \approx r_{\text{inf}}(I)$. We assume the individual exposure responses to be statistically independent, and can thus understand this process as a series of $\sum D(I)$ independent Bernoulli-experiments. As we expect the success rate to be very low in our exemplary data set, we estimate the confidence intervals using the Agresti-Coull method [62], ensuring that the two-sided confidence bounds remain within the $(0, 1)$ interval.

3.2 Surrogate models

In the presence of confounding effects such as homophily and shared local or global influences, contagion can be hard to identify [63]. In the next section (Sect. 3.2.1), we describe our use of surrogate models to address this challenge, followed by the description of the surrogate models produced for this study (Sect. 3.2.2).

3.2.1 Surrogate model method

The surrogate model approach is a statistical proof-by-contradiction method used for investigating specific features and correlations in empirical data sets. It is based on testing composite null hypotheses on data sets that are derived from the empirical data using Monte Carlo methods [51,52,55,56]. A variety of time series [48,64–66] and network data sets [48,67–69] have been analyzed using surrogate models. The method is described in the following paragraph.

First, a composite null hypotheses $\mathcal{H}_0$ is constructed, which specifies a class of processes that may be sufficient to reproduce the observed empirical data. Ideally, $\mathcal{H}_0$ excludes certain features or correlations in the data, e.g. relating to hypothetical underlying contagion processes. Based on $\mathcal{H}_0$ and the empirical data, a surrogate data set is then constructed that resembles the original data, but lacks the hypothetical features excluded by $\mathcal{H}_0$. Partially randomizing the empirical data set, in a way that is consistent with the null hypothesis, is one option of generating such data sets. This method, referred to as constrained realizations [70], forces the resulting data set to resemble the empirical data in key statistical measures as directed by $\mathcal{H}_0$. Specific correlations and data features may thus be selectively removed, without committing to a specific model. An ensemble of surrogate data sets is produced for each null hypothesis to reduce statistical uncertainties. Finally, a discriminating statistic is computed on both the empirical data and the ensemble of surrogate models. If the empirical value differs significantly from the ensemble

of surrogate values, the null hypothesis is rejected. This can be regarded as evidence that the preserved features are not sufficient to explain the observations, pointing to a more complex underlying mechanism. By carefully choosing progressively more complex null hypotheses, the nature of this underlying process can be investigated.

Using the dose-response contagion approach, we would like to compare the empirical DRF with those computed on each surrogate model realization. We must therefore find a measure that may quantify the differences of a large number of functional shapes. Comparing with the simple bin-wise average of the surrogate DRFs is not sufficient here, as individual data points within a single surrogate model realization are not statistically independent from each other. This is because of the constrained realizations method, which may preserve measures such as the total number of infected cities while randomizing other variables. In that case, a data point being raised in one bin of the surrogate data DRF (signifying a greater number of infections at this level of infection pressure) will always result in a different data point being lowered, making them statistically correlated. Since bin-wise averaging over all surrogate data sets would destroy such inter-bin correlations, we want to instead compare the individual surrogate model DRF's functional forms with the empirical DRF. We achieve this by performing weighted least-squares fits to each individual surrogate sample's DRF, and comparing the resulting fit parameter distribution to the parameters of the fit obtained from empirical data.

A number of different DRF shapes are plausible [46], and thus the fitted DRF shape in general has to be chosen with care. However, in the studied data set, only $\mathcal{O}(100)$ cities have adopted the BRT innovation. Infection doses in the network are thus generally low, and any saturation effects are unlikely to have significant effects. We therefore expect the empirically determined DRF to be close to linear, and perform the fit using a polynomial of degree one:

$$p(I) \approx m \cdot I + b. \tag{3}$$

In the fits, each bin's data point is weighted with the total number of times the corresponding infection pressure range was observed in data, displayed as the bin height in Fig. 2A. The fit parameters m (DRF slope, a measure of the sensitivity of cities' reactions to the dose I) and b (y-axis intersection point, a measure of spontaneous background infection rate at received dose $I = 0$) are chosen as the discriminating statistic for comparing empirical and surrogate model DRFs.

The fit parameter distribution of the DRFs computed on surrogate models can be visualized as a two-dimensional histogram. We non-parametrically estimate the underlying two-dimensional probability distribution $P(m, b)$ using kernel density estimation, with Gaussian kernels whose width is set by Silverman's rule [71]. We then calculate the value of the quantile function $Q(m_{\mathrm{emp}}, b_{\mathrm{emp}})$ of this distribution, integrating the probability density function over the parameter space

where it gives a lower probability than the one it gives for the empirical fit parameters ($m_{\mathrm{emp}}, b_{\mathrm{emp}}$):

$$
\begin{aligned}
&Q_{\mathcal{H}_0}(m_{\mathrm{emp}}, b_{\mathrm{emp}}) \\
&= \iint_{-\infty}^{\infty} P(m, b)\theta(m_{\mathrm{emp}}, b_{\mathrm{emp}}) \, \mathrm{d}m \, \mathrm{d}b \\
&\theta(m_{\mathrm{emp}}, b_{\mathrm{emp}}) = \begin{cases} 1 & P(m, b) \leq P(m_{\mathrm{emp}}, b_{\mathrm{emp}}) \\ 0 & \text{otherwise} \end{cases} \quad (4)
\end{aligned}
$$

This quantile represents a stochastically robust measure of the difference between surrogate DRF parameter distribution and the empirical DRF parameters. It is used to accept or reject the null hypotheses that the surrogate models are based on. We set the significance threshold for the rejection of null hypotheses at $Q_{\mathcal{H}_0} > 0.05$.

3.2.2 Surrogate model production

In this analysis, surrogate models for four null hypotheses are produced to probe the underlying mechanism of the spreading behavior of BRTs and their relationship with the global flight route network. A large ensemble of 5000 ($\mathcal{H}_0^1$, $\mathcal{H}_0^2$) or 10,000 ($\mathcal{H}_0^3$, $\mathcal{H}_0^4$) realizations is computed for each surrogate model, to reduce the influence of statistical fluctuations. To determine the statistical stability of the result, several of these ensembles are produced for each null hypothesis. If the value of the quantile function Q differs between the different ensembles of the same null hypothesis, five ensembles are produced, and their mean value for Q is calculated. The statistical uncertainty is then estimated conservatively as the largest difference between the mean value for Q and any of the individual ensemble's values. We use the canonical naming convention put forward in [51] to describe the surrogate models M associated with the null hypotheses $\mathcal{H}_0$. Surrogate models are thus defined by the quantities they conserve with respect to the original empirical data. To make the surrogate models generally comparable to the empirical data, the number of cities infected in the studied time period, the structure of the network, and the identity of previously infected cities are conserved in all models. The hierarchy of their null hypotheses, conserving progressively more features of the data, is described here.

1. $\mathcal{H}_0^1$: $M(w_{ij}, N_{\mathrm{inf}})$. *The empirical DRF can be reproduced with a class of models that is only based on the structure of the network w_{ij}. This most basic hypothesis is designed to check if the observed DRF is purely an artifact of the flight network structure. To produce the surrogate data set, the identities of infected cities are randomly re-assigned to other cities in the network, and their respective infection times are drawn from a uniform distribution.*
2. $\mathcal{H}_0^2$: $M(w_{ij}, p(\hat{t}))$. *The empirical DRF can be reproduced with a class of models that is only based on the network structure, and the distribution of infection timestamps $p(\hat{t})$. This hypothesis additionally*

1614

Eur. Phys. J. Spec. Top. (2022) 231:1609–1624

investigates the influence of the infection year distribution, which can be seen to have a strong upward trend in Fig. 1D. The first step to producing the surrogate model data for this null hypothesis is the randomization of the identity of newly infected cities, analogous to the previous case ($\mathcal{H}_0^1$). Here, however, the points in time that the cities become infected are drawn from a probability distribution that is derived from the empirical data. The kernel density estimation-generated probability density function used for this is displayed in Appendix B.

3. $\mathcal{H}_0^3$: $M(w_{ij}, n_i)$. *The empirical DRF can be reproduced with a class of models that is only based on the network structure, and the identity/position of the infected cities n_i in the network.* This hypothesis again builds on $\mathcal{H}_0^1$, conserving the network structure, but additionally conserves the identity of newly infected cities n_i. Thus, the state of the system at the last time step is exactly the same as in the empirical data, with the same network, and the same cities infected. Only the timing of the infections is changed. This hypothesis thus tests whether the position of infected cities in the network is sufficient to explain the observed DRF. This represents a useful test for homophilic effects: If the null hypothesis cannot be rejected, then it is apparent that the *timing* of the infections is not a relevant factor for the spreading process. Consequently, the probability of becoming infected would not depend on other infections in the network neighborhood. It could instead be dominated by, e.g., the membership in a certain closely-connected clique of cities for which a BRT is especially suitable. The surrogate model data is produced by re-drawing the infection year for each city infected in the studied time period from a uniform distribution.

4. $\mathcal{H}_0^4$: $M(w_{ij}, n_i, p(\hat{t}))$. *The empirical DRF can be reproduced with a class of models that is only based on the network structure, the network position of the infected cities, and the distribution of infection times.* Building on $\mathcal{H}_0^2$ and $\mathcal{H}_0^3$, the overall distribution of infection years is again conserved by this null hypothesis. If any artifacts are introduced by the uniform infection time distribution in $\mathcal{H}_0^3$, they should be removed by requiring the infection years to be distributed as they are in the empirical data. This is again achieved by drawing them from a distribution derived from the empirical data through kernel density estimation, as displayed in Appendix B.

4 Results

In this section, the results of the analysis are displayed and interpreted. As in Sect. 3, the empirical dose response function (DRF) is treated first (Sect. 4.1), followed by the results of the surrogate model study (Sect. 4.2).
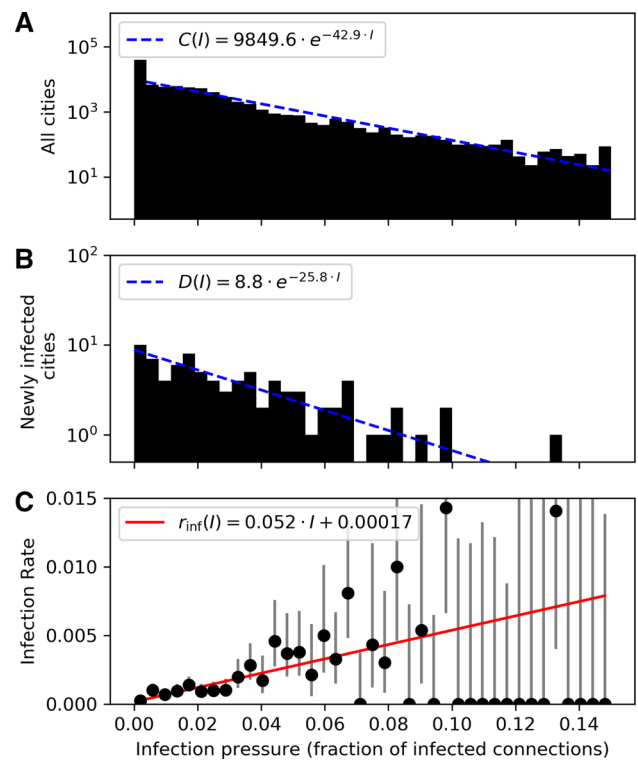

Fig. 2 Infection pressures on cities, and dose response function (DRF). Distributions of **A** the infection pressures $I_i(t)$ experienced by all cities within the studied time interval, and of **B** infection pressures experienced by those cities that implemented a Bus Rapid Transit System (BRT) the following year. An exponential function fitted to the data is given for each (dashed blue) for ease of viewing; note that the fit in **A** excludes the first bin. The infection pressure experienced by a city is defined as the fraction of weighted network connections to cities that have adopted the BRT innovation. In **C**, the empirical dose response function $r_{\text{inf}}(I)$ is displayed, obtained by the bin-wise division of **B** and **A**. The binomial error of each data point is estimated using Agresti-Coull intervals. An alternatively scaled version of this figure, with fully displayed errorbars in C and an identical y-scale in A and B, can be found in Appendix C. A clear upward trend is visible: cities whose network neighborhood already featured more BRT implementations appear to implement the BRT implementation more frequently

4.1 Empirical dose response function

The distribution of infection pressure (dose) values I, experienced by any city at any time step, is displayed in Fig. 2A. The lowest bin, containing cities that had no or very light connections to cities with BRT, is strongly elevated, followed by a brief section up to an infection pressure of about 0.03 where data points are scattered around an infection rate of 0.001. Above $I = 0.03$, a roughly exponential decay in frequency from low to high pressures is visible. "Successful" infection pressures, that is, infection pressure experienced by those cities that are about to adopt a Bus Rapid Transit System (BRT), are displayed in Fig. 2B. The limited

Eur. Phys. J. Spec. Top. (2022) 231:1609–1624

1615

number of BRT adoptions becomes very apparent here; the number of cities newly infected in the studied time period, and thus the number of data points in Fig. 2B, is 86. The downward trend can nevertheless be clearly observed to be flatter than the one in Fig. 2A. Likewise, the lowest bin is not nearly as emphasized. Dividing the two histograms to obtain the empirical DRF thus yields a graph with a marked upward trend, displayed in Fig. 2C. The shape of the DRF appears roughly linear, the linear fit yielding a slope of $m = 0.052$, and a y-axis cutoff of $b = 1.7 \times 10^{-4}$. This confirms our expectation and justifies the linear fits performed for the surrogate data study (described in Sect. 3.2). While the DRF values for many high infection pressures are zero, the very large error bars show that a lack of data introduces significant uncertainty in this regime of larger I.

The strong positive correlation between infection pressure and infection rate points to contagion effects at first glance: Cities whose network neighbors have previously adopted BRTs to a large fraction, more frequently also adopt BRTs. However, this correlation is not necessarily causal. The correlation could be an artifact of some other process, such as homophilic effects: If BRTs were especially suited for a certain clique of cities, which happens to be closely connected internally, similar correlations might arise. There may also be the possibility of other, more basic attributes of the data such as the network structure itself being sufficient in explaining the observed DRF. To investigate and exclude these confounding effects, surrogate model tests are performed, as discussed in the following section.

4.2 Surrogate models

In this section, the results of the surrogate model tests are presented. For orientation, the short descriptions of the surrogate models are repeated here; they are described in more detail in Sect. 3.2.2 As mentioned there, the number of cities infected in the studied time period, the structure of the network, and the identity of cities infected before the studied time interval are conserved in all models.

First surrogate test. $\mathcal{H}_0^1 : M(w_{ij}, N_{\text{inf}})$. *The empirical DRF can be reproduced with a class of models that is only based on the structure of the network, and the number of newly infected cities N_{inf}.* This represents the most basic assumption testing whether the observed DRF is merely a product of the structure of the network structure itself, without regard for which or when cities implement a BRT. This would imply that the spreading of the BRT innovation proceeds completely independent of the network, and thus also completely independent of any other variables correlated with a city's position in the network. If, instead, a city's network connections are in some way relevant for its BRT adoption probability, the positive correlation of the empirical DRF should not be found in this surrogate model. As shown in Fig. 3A, the latter is evidently the case: The average surrogate DRF (in blue) is nearly com-

pletely flat over the entire infection pressure range. The fit parameter comparison in Fig. 3B thus only confirms the obvious. The empirical DRF's fit parameters are strongly separated, with $Q_{\mathcal{H}_0^1} \approx 0$. This result remains the same for repeated re-creations of the ensemble. The null hypothesis is thus rejected.

Second surrogate test. $\mathcal{H}_0^2 : M(w_{ij}, p(\hat{t}))$. *The empirical DRF can be reproduced with a class of models that is only based on the network structure, and the distribution of infection times $p(\hat{t})$.* The second null hypothesis builds on the first one, and additionally preserves the yearly BRT adoption rate; that is, the distribution of the years in which new BRTs are implemented. In view of the strong rejection of $\mathcal{H}_0^1$, we expect a similar result for this surrogate model, since the potential effect of the variable infection rate appears negligible compared to the randomization of which cities become infected. The result of this test is displayed in Fig. 3C, D. As expected, the average surrogate model DRF remains flat, and the distribution of surrogate DRF fit parameters is strongly separated from the empirical DRF fit parameters. A difference to $\mathcal{H}_0^1$ is visible, however: the fit parameter distribution (Fig. 3D) is shifted towards positive slopes, whereas it was centered around $m = 0$ for $\mathcal{H}_0^1$ (Fig. 3B). Combined with the correspondingly lower values for b, this reduces the separation between the distribution and the empirical for parameters. However, the value of the quantile function remains $Q_{\mathcal{H}_0^2} \approx 0$, which remains the same for repeated re-creations of the ensemble. $\mathcal{H}_0^2$ is thus rejected. The time distribution of the BRT implementations appears to hold significance for the DRF, but is not nearly sufficient to explain the empirically observed correlation.

Third surrogate test. $\mathcal{H}_0^3 : M(w_{ij}, n_i)$. *The empirical DRF can be reproduced with a class of models that is only based on the network structure, and the position of the infected cities in the network.* This hypothesis again builds on $\mathcal{H}_0^1$, this time by conserving the the identity of the infected nodes in the network, meaning which of the cities implement BRT. The resulting surrogate model is thus much more strongly constrained to resemble the empirical data than for the previous two hypotheses. The times at which these infections occur are randomly drawn from a uniform distribution. Using this null hypothesis, we can probe the system for homophilic effects. If, for example, there is a closely connected clique of cities that BRT is especially suited for, or that share a BRT-favorable political environment, then the infected cities' position in the network would be expected to be sufficient to reproduce the observed DRF. In particular, the specific times of infection / BRT adoption would be irrelevant, as the null hypothesis posits.

The comparison of the empirical DRF with the ensemble of surrogate model DRFs is displayed in Fig. 4A, B. Here, the importance of not relying on the bin-wise average of the surrogate DRF ensemble becomes apparent: while the difference between
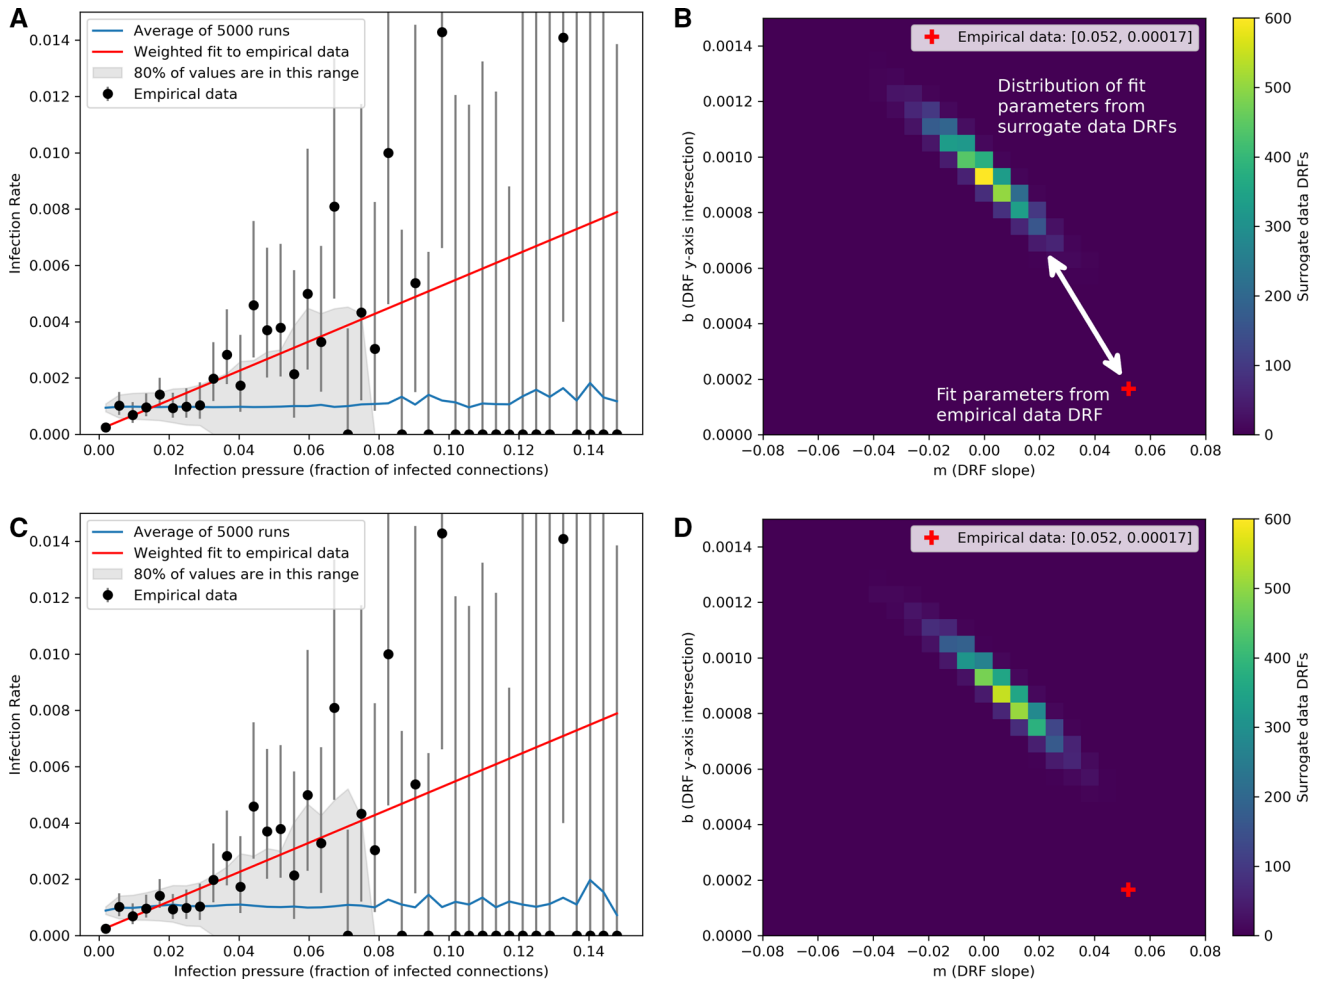
Fig. 3 Surrogate data sets with randomly reassigned infected cities. **A**, **C** Comparison of the empirical dose response function (DRF, black circles) and the bin-wise average of the DRFs of 5000 surrogate model runs (blue), for the surrogate models where the identity of cities implementing BRT ("infected cities") are randomly reassigned to cities in the network. In **A** ($\mathcal{H}_0^1$), the infection times (BRT implementation timestamps) of the infected cities are randomized uniformly, while in **C** ($\mathcal{H}_0^2$), the distribution of infection times is conserved. The error bars of the empirical DRF are calculated as described in Sect. 3.1. Weighted linear least-squares fits are computed for each surrogate model

realization; their fit parameters are displayed in **B** and **D**. The weighted linear least-squares fit to the empirical data is displayed as a red line in **A** and **C**, whose parameters are marked as a red cross in **B** and **D**, respectively. As opposed to the empirical DRFs, the average surrogate DRFs appear flat throughout most of the infection pressure range (**A**, **C**) and the corresponding fit parameter distributions differ strongly from the empirical fit parameters. The null hypotheses $\mathcal{H}_0^1$ and $\mathcal{H}_0^2$ can thus be rejected. These figures show that the network position of infected cities is relevant for the spreading of the innovation, and cannot be ignored in a model describing this process

empirical and surrogate data is not readily apparent in Fig. 4A, it can more clearly be discerned in B. While the value of the quantile function $Q_{\mathcal{H}_0^3} = 0.0296$ comes closer to the chosen significance threshold than in the previous two tests, the null hypothesis is nonetheless rejected. This value for $Q_{\mathcal{H}_0^3}$ remains the same for repeated re-creations of the ensemble. Evidently, this surrogate model can reproduce the empirical DRF much better than the two previous ones. Which nodes become infected is thus apparently correlated with their network position, and homophilic effects may be at play. However, they do not appear sufficient on their own to explain the observed empirical DRF.

We observe that the *individual* values of m_{emp}, and to a lesser extent b_{emp}, are not uncommon in the fit parameter distribution of the surrogate DRFs. Only in the two-dimensional graph does it become apparent that their *combination* (m_{emp}, b_{emp}) is very rarely observed. It appears that the empirical DRF's data point at $I \approx 0$, being very close to $r_{inf}(I) = 0$ and having very small error bars (Fig. 4A), more strongly constrains b_{emp} to small values than is the case for the surrogate model data. We interpret this as the observation that, in the absence of an infected, or "seed" city in a city's neighborhood, the infection rate is inhibited. Thus, in the empirical data, cities are unlikely to adopt
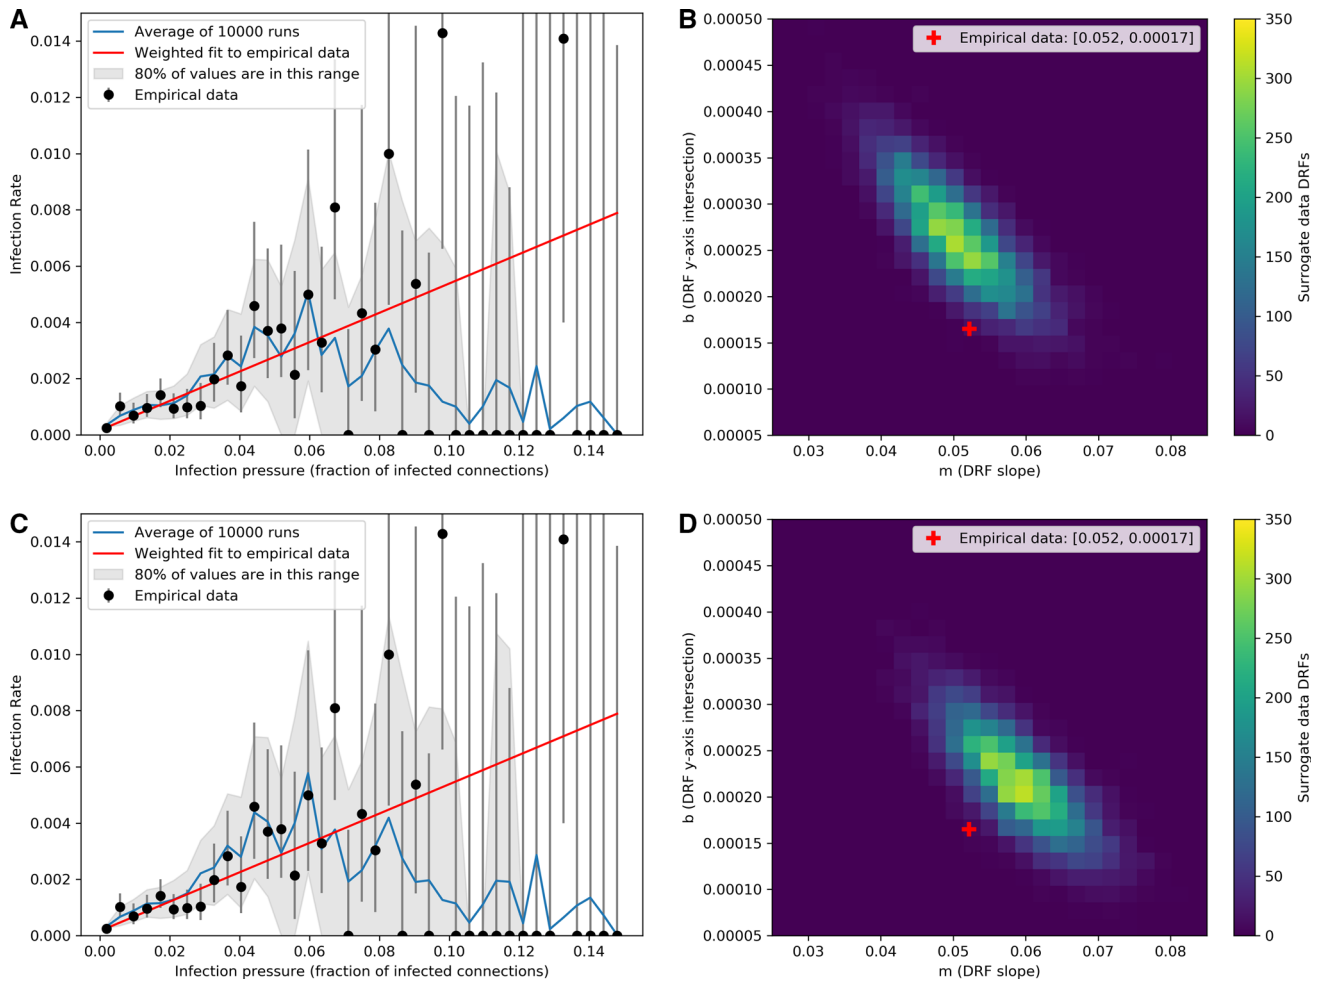
Fig. 4 Surrogate data sets with identities of infected cities conserved. **A**, **C** Comparison of the empirical dose response function (DRF, black circles) and the bin-wise average of the DRFs of 10 000 surrogate model runs (blue), for the surrogate models where the identity of cities implementing BRT ("infected cities") are conserved. In **A** ($\mathcal{H}_0^3$), the infection times of the infected cities are randomly drawn from a uniform distribution covering the investigated time interval. In **C**, ($\mathcal{H}_0^4$) the infection times are drawn from a distribution derived from a kernel density estimation of the distribution found in empirical data. The error bars of the empirical DRF are calculated as described in Sect. 3.1. Weighted linear least-squares fits are computed for each surrogate model realization; their fit parameters are displayed in **B** ($\mathcal{H}_0^3$) and **D** ($\mathcal{H}_0^4$). The fit to the empirical data is displayed as a red line in **A** and **C**, whose parameters are marked as a red

cross in **B** and **D**. In both **A** and **C**, the average surrogate DRFs appear to match the empirical DRF. However, individual data points are correlated within each surrogate model realization's DRF. The fit parameter comparison in **B** and **D** takes this into account, revealing a significant difference between the distribution of surrogate model fit parameters and those of the empirical DRF. The null hypotheses $\mathcal{H}_0^3$ and $\mathcal{H}_0^4$ can thus be rejected as well. These figures demonstrate that the timing of cities' innovation adoptions relative to one another is not random, and important for the innovation spreading process. An explanation solely based on homophilic effects, implying a spreading process that is only based on preferential attachment between cities that are likely to adopt BRT, is thus not sufficient to explain the data

BRT if *none* of their neighbors have done so yet, suggesting a contagion process to be possibly at play.

The analysis of $\mathcal{H}_0^2$ demonstrated that the non-constant rate of infections may play a role for the underlying process. Building on $\mathcal{H}_0^3$, but conserving this distribution as well, thus remains as the final test that may distinguish between homophily and potential contagion effects. This most restrictive surrogate analysis

performed in this study is described in the following paragraph.

Fourth surrogate test. $\mathcal{H}_0^4$: $M(w_{ij}, n_i, p(\hat{t}))$. *The empirical DRF can be reproduced by a class of models that is only based on the network structure, the network position of the infected cities, and the distribution of infection times.* Combining null hypotheses $\mathcal{H}_0^2$ and $\mathcal{H}_0^3$, this hypothesis is designed as the most stringent test for homophily. The infection timestamps of newly

1618

Eur. Phys. J. Spec. Top. (2022) 231:1609–1624

infected cities are, again, sampled from a probability distribution derived from empirical data via kernel density estimation.

In this test, the value for the quantile function $Q_{\mathcal{H}_0^4}$ fluctuated slightly for subsequent re-creations of the ensemble. The ensemble was thus produced five times, and the mean value for $Q_{\mathcal{H}_0^4}$ is given in the following. The uncertainty interval is estimated conservatively as the largest difference between the mean and any of the constituent values. We explain the greater fluctuation with the reduction of the randomization phase-space, by the additional constrains placed on the surrogate data. The result most closely matching the mean value for $Q_{\mathcal{H}_0^4}$ is displayed in Fig. 4C, D.

Similar to $\mathcal{H}_0^3$, no significant difference between the empirical and average surrogate DRFs can be observed (Fig. 4C). However, the distribution of surrogate DRF fit parameters is clearly removed from the empirical DRF's fit parameter. The additional conservation of the infection time distribution has shifted the distribution towards higher m and lower b, similar to the difference between $\mathcal{H}_0^1$ and $\mathcal{H}_0^2$. With a quantile function value of $Q_{\mathcal{H}_0^4} = 0.0138 \pm 0.0037$, null hypothesis $\mathcal{H}_0^4$ is actually somewhat more strongly rejected than $\mathcal{H}_0^3$. The uncertainty is too small to have an effect on the hypothesis rejection decision.

Similarly to $\mathcal{H}_0^3$, the individual values for m_{emp} and b_{emp} are common in the distribution of surrogate DRF fit parameters, and steeper DRF slopes $m > m_{\mathrm{emp}}$ are actually more likely to occur. However, the combined analysis of $(m_{\mathrm{emp}}, b_{\mathrm{emp}})$ again shows the significant difference between the empirical and surrogate data. This demonstrates the strength of the method in distinguishing differences between empirical and surrogate data, beyond what would have been visually perceivable in Fig. 4C.

Both contagion and homophilic mechanisms are expected to create a strong correlation between the nodes' network position and their probability to become infected. However, only in the case of contagion dynamics does the relative timing of infection events matter: For a kind of "infection wave" traveling from node to node across the network, the order of infections is not random. Thus, with the rejection of both $\mathcal{H}_0^3$ and $\mathcal{H}_0^4$, we find homophilic effects to be an insufficient explanation of the data. We interpret this as evidence pointing towards underlying contagion mechanisms controlling the spreading behavior of BRT innovations between cities, where the linkages provided by the global flight route network appear to be an adequate proxy (Table 1).

5 Discussion and conclusion

Cities may learn from each other which policies and technologies to adopt. As part of the human response to anthropogenic environmental changes, this represents an Earth System process that has attracted lit-

Table 1 Values of the quantile function $Q(m_{\mathrm{emp}}, b_{\mathrm{emp}})$, as defined in Eq. 4 for the four investigated null hypotheses. All four are rejected, though much more narrowly for $\mathcal{H}_0^3$ and $\mathcal{H}_0^4$. Only $Q_{\mathcal{H}_0^4}$ showed any significant fluctuation between identically produced ensembles. For details on the kernel density estimations which form the basis for the calculation of Q, see Appendix D

Null hypothesis	Surrogate model	$Q_{\mathcal{H}_0^i}$
$\mathcal{H}_0^1$	$M(w_{ij}, N_{\mathrm{inf}})$	0.0000 ± 0.0000
$\mathcal{H}_0^2$	$M(w_{ij}, p(\hat{t}))$	0.0000 ± 0.0000
$\mathcal{H}_0^3$	$M(w_{ij}, n_i)$	0.0296 ± 0.0000
$\mathcal{H}_0^4$	$M(w_{ij}, n_i, p(\hat{t}))$	0.0138 ± 0.0037

tle research attention at the macroscopic scale so far. In this contribution, we proposed a method for investigating the spreading of urban innovations related to sustainability on the global network of cities. The method is made up of two steps: first, we estimate dose-response functions (DRFs) from empirical data on the network proxy and spreading process. Second, we perform hypothesis tests on surrogate data models generated from the empirical data, to probe and exclude specific effects that may confound the detection of contagion effects.

This method was demonstrated on a pair of example data sets: We correlate the spreading of Bus Rapid Transit Systems, a public transport innovation, with the global network of flight routes as a proxy for inter-city learning connections. We find significant evidence towards contagion processes in this data, which cannot be sufficiently explained by homophilic effects such as a shared environment or clustering. Cities whose neighborhood is free of BRT implementations appear especially unlikely to adopt a BRT themselves. This underlines the fact that cities often base their policy decisions on each others' experiences.

Our results indicate that it is possible to find proxies for the global city learning network. This forms the basis for future work, for example on comparing which inter-city linkages offer the largest predictive power for the spreading of innovations, and what kind of (complex) contagion processes are at work. The investigation of these processes has the potential to improve our understanding of the human components of the Earth System. Identifying "gaps" in this network, such as cluster boundaries that innovations rarely cross, may even open the doors for targeted interventions in the future. Such actions could ensure that vital innovations are adopted quickly around the globe, and reduce the time lag between the inception of a good idea and its environmental impact on a global scale. Even social tipping processes might be facilitated, with innovation uptake starting locally and then spreading rapidly via self-enforcing positive feedbacks. Similar to this work, social tipping processes are inherently based on network processes [72], which include changes of network nodes internally, or of the network structure itself. Sen-

Eur. Phys. J. Spec. Top. (2022) 231:1609–1624

1619

sitive interventions points, as mentioned above, have been linked to social tipping processes relevant for transitions in and around urban systems [73], decarbonization [39], and other positive tipping points [38,74].

Our analysis has a range of limitations, which we would like to address here. Firstly, those relating to the datasets used to illustrate the method. We base the network component of our analysis on a publicly available dataset of scheduled flight routes, which provides only a static snapshot. Temporally resolved data, especially on actual origin-to-destination passenger flows instead of scheduled airport-to-airport routes, could be expected to contain a clearer signal. Using a different weighting algorithm [75] to create a different network using a rescaled network distance from the same underlying data could possibly also result in a better predictor. On the matter of the spreading innovation, the low ($\mathcal{O}(100)$) number of total BRT adoptions likely significantly reduces our ability to discern actual trends from noise and fluctuations in the data. On the methodological side, our analysis shares the drawbacks of all statistical methods, in that hypotheses about the data can only be rejected, and a confirmative result is impossible. The proof-by-contradiction nature of our method can thus only indirectly infer contagion processes. Furthermore, DRFs are a highly aggregate statistical measure, which may not be specific enough for accurately characterizing subtle spreading processes under the data limitations. A systematic analysis of the statistical power of the method under varying conditions of noise, system size, different underlying spreading processes, and other parameters would be desirable to more accurately interpret the results. Such an analysis, which could be achieved using large ensembles of simulations, lies beyond the scope of this contribution.

Furthermore, there are a number of potential external factors that may drive or facilitate the innovation's spreading, confounding the results extracted from the example data sets. Processes such as increasing urbanization (growing populations and population densities, demographic changes), large-scale economic shifts (arrival of new industry sectors, economic and cultural effects of globalization), or changes in the political landscape (regulatory changes on national and supernational level, e.g. in the EU), and countless others surely play a role in cities' individual decisions to adopt Bus Rapid Transit Systems. These may occur simultaneously for entire regions, and, for example, greatly raise the "susceptibility" of affected cities to implement the innovation. Such factors could possibly either mimic, or drown out, contagion processes in the data. Unfortunately, quantifying these factors, and controlling for them in a global data set, poses great methodological and data availability challenges, and is beyond the scope of this study. However, it should be noted that even under conditions of greater "susceptibility," the decision to implement a particular public transport innovation (as opposed to its competitors) may still depend on the experience of other, closely connected cities. Analyzing the spreading of a set of innovations with very different characteristics, spanning both tech-

nological and governance innovations and addressed at different kinds of urban challenges, may allow future research to mitigate the masking effects of these drivers without requiring a detailed understanding of each factor.

While our method can exclude specific correlations and causal connections in the data, it cannot provide a detailed process-based understanding of urban innovation transmission. In future work, hypotheses of concrete mechanisms should be tested, for example using combinations of Monte Carlo and maximum likelihood methods. A similarly promising approach would be to use higher-order statistics, such as multi-node correlations for identifying longer contagion chains and network motifs related to spreading dynamics.

The applications of our method are limited by data availability. Comprehensive databases on intercity linkages are rare, especially on the global scale. While abundant data often exists on the national and inter-national level, homogeneous city-scale data that covers more than just the largest "world cities" [12,76] is much harder to find. If cities' role as active agents in the Anthropocene is to be taken seriously, then the collection of high-quality data on this scale is paramount. The comparative analysis of multiple spreading processes, correlated with a number of different inter-city network proxies, promises valuable insights on the real-world mechanisms underlying the transfer of urban innovations.

Our proposed method is generic, and can be applied to probe for contagion effects in any combination of city network and spreading innovation. It can also easily be generalized to temporally dynamic networks. Furthermore, it may be useful for the analysis of other complex systems wherever contagion processes are hypothesized in low-rate spreading phenomena on densely connected networks. Potential applications include the spreading of opinions, social norms and behaviors among individuals and more abstract agents, wherever such data sets are available. As highlighted, this proposes a potential connection to research on social tipping processes [3], which are one of the most promising pathways for mitigating dangerous anthropogenic global warming and the long-term crossing of other planetary boundaries.

1620

Eur. Phys. J. Spec. Top. (2022) 231:1609–1624

Bus Rapid Transit Systems was obtained from BRTData [60] under a *Creative Commons BY-NC-ND 3.0* license [77]. The data on scheduled flight routes was obtained from OpenFlights [57], made available under Open Database [78] and Database Contents Licenses [79]. City data sourced from wikidata [59] is in the public domain [80].

Author contribution statement

NHK, PR and JFD conceived this study. NHK prepared the data, implemented the data analysis methodology, and visualized the results, with supervision by JFD and PR. NW contributed the link weight algorithm for the flight route network. All authors interpreted and discussed the results. NHK wrote the manuscript, with support from PR and JFD and NW. All the authors give their final approval of the article version to be published.

Code and data availability The data analysis scripts and (to the extent that licenses allow) curated data sets used in this study will be made available as open source code upon publication of this article. This manuscript has associated data in a data repository. [Authors' comment: They are available under the https://doi.org/10.5281/zenodo. 6105486.]

A City connection weight calculation

The connection weights of the cities, interpretable as abstract proximity measures, are derived from the flight route database [57] and the geographic distance between airports and city centers [57,59]. The network is constructed as follows:

1. All cities with a population above 60 000 are considered as nodes in the network.
2. All airports within 60 km of a city's center are considered to belong to this city. Cities sharing airports and vice versa are possible.
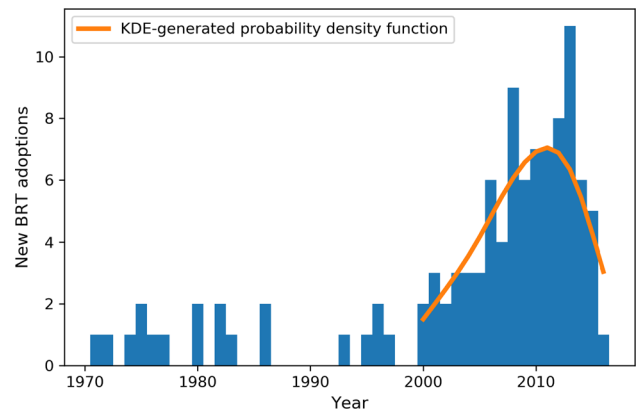
Fig. 5 KDE-derived time-dependent infection probability. Displayed are the yearly Bus Rapid Transit System (BRT) adoption rate (blue), and the kernel density estimation-derived approximation of the underlying probability distribution (orange). This distribution is used in the surrogate data production of $\mathcal{H}_0^2$ and $\mathcal{H}_0^4$
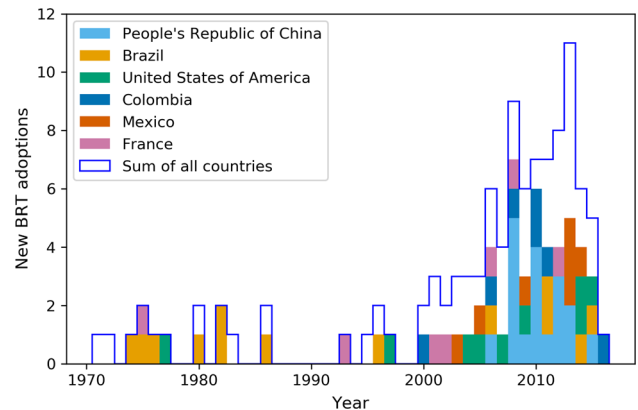

Fig. 6 Time-dependent BRT implementation rate for the countries with the most adoptions. Superimposed with the global BRT adoption rate (blue outline), the yearly BRT adoptions for the six countries that have the most cities with BRT adoptions are displayed according to the color key in the figure

3. Two cities are connected to each other if one can be reached from the other with ≤ 2 flights. The weight w_{ij} of this connection functionally depends on the multiplicity of the flight routes $r_{i \to j}$ and $r_{j \to i}$ that directly connect i and j (if any), as well as the number and multiplicity of viable 2-flight paths between the cities. The functional dependence is defined here:

$$w_{ij} = f_{i \to j} + f_{j \to i} \tag{5}$$

$$f_{i \to j} = r_{i \to j} + \frac{\sum_{k=1}^{M_i} r_{i \to k} r_{k \to j}}{\sum_{l=1}^{M_k} r_{k \to l}} \tag{6}$$

Here, $\sum_{k=1}^{M_i}$ refers to the sum over all M_i cities that can be reached with a single flight from city i. Thus, the contribution to w_{ij} of any 2-flight path $i \to k \to j$ is the product of the route multiplicities $r_{i \to k}$ and $r_{k \to j}$, divided by the out-degree of node k.
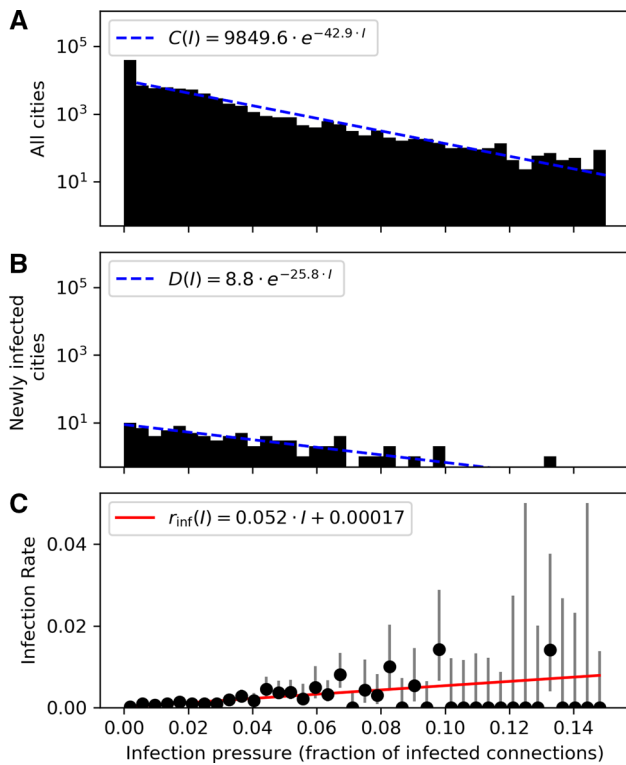
Fig. 7 Alternative version of Fig.2: infection pressures on cities, and dose response function (DRF). This figure displays the same contents as Fig. 2, but with rescaled y-axis in all three subfigures. **A** and **B** now have the same y-scale, facilitating a comparison of their downward trend. In **C**, the displayed y-axis is increased, so that all errorbars are entirely visible

Potential connections using more than two flights are not considered, as this would significantly increase the computation time and network density, while not adding significantly weighted connections. Note that the scheduled flights themselves are directed; $i \to k \to j$ represents a valid connection between i and j, while $i \to k \leftarrow j$ does not. However, the resulting city connection weight $w_{ij} = w_{ji}$ is undirected, or symmetric, by construction.

B Temporal and spatial aspects of the spreading of Bus Rapid Transit System

In this appendix, the temporal and spatial characteristics of the Bus Rapid Transit innovation are explored in more detail. Fig. 5 shows the global Bus Rapid Transit System (BRT) adoption rate, overlaid with a kernel density estimate of the underlying probability distribution. This probability distribution is used in the surrogate data production for null hypotheses $\mathcal{H}_0^2$ and $\mathcal{H}_0^4$. It is only calculated on the studied year range of 2000 to 2016.

In Fig. 6, the BRT spreading in the six countries with the most implementations is shown: China (18 adoptions), Brazil (13), USA (9), Colombia (9), Mexico (9) and France (9). All other countries have 4 adoptions or less. It can be observed that, while Brazil dominates the pre 2000-period and China does the same for the time after that year, no single country can be identified as the sole main driver of the spreading dynamics.

C Alternative version of the empirical dose response function figure

In Fig. 7, an alternative version of Fig. 2 is displayed. In this version, the y-axes of subfigures A and B cover the same range. It is thus much easier to see that the downward trend of the distribution of all infection pressures $I_i(t)$ (Fig. 7A) is stronger than the one in Fig. 7B, which displays the infection pressures experienced by those cities that implemented a Bus Rapid Transit System (BRT) the following year. This makes the upward trend of the DRF, shown in Fig. 2C and calculated by dividing the histograms in B and A, more intuitive.

Furthermore, Fig. 7C is shown with an increased y-range compared to Fig. 2C. The data points at high infection pressure and zero infection rate can be seen to have very large error bars, demonstrating their small contribution to the weighted least-squares fit.

D Quantile function calculation for fit parameter distributions

The underlying probability distributions from the fit parameter distributions displayed in Figs. 3 and 4, are estimated using kernel density estimations. Gaussian kernels were used, and their bandwidths were determined using Silverman's rule [71]. In Fig. 8, the estimated probability distributions are displayed. The fit parameters of the empirical DRF is shown as a red cross. The value of the quantile function Q, measuring the likelihood of the empirical data point being produced by the distribution, is noted in the figures' legends.

1622

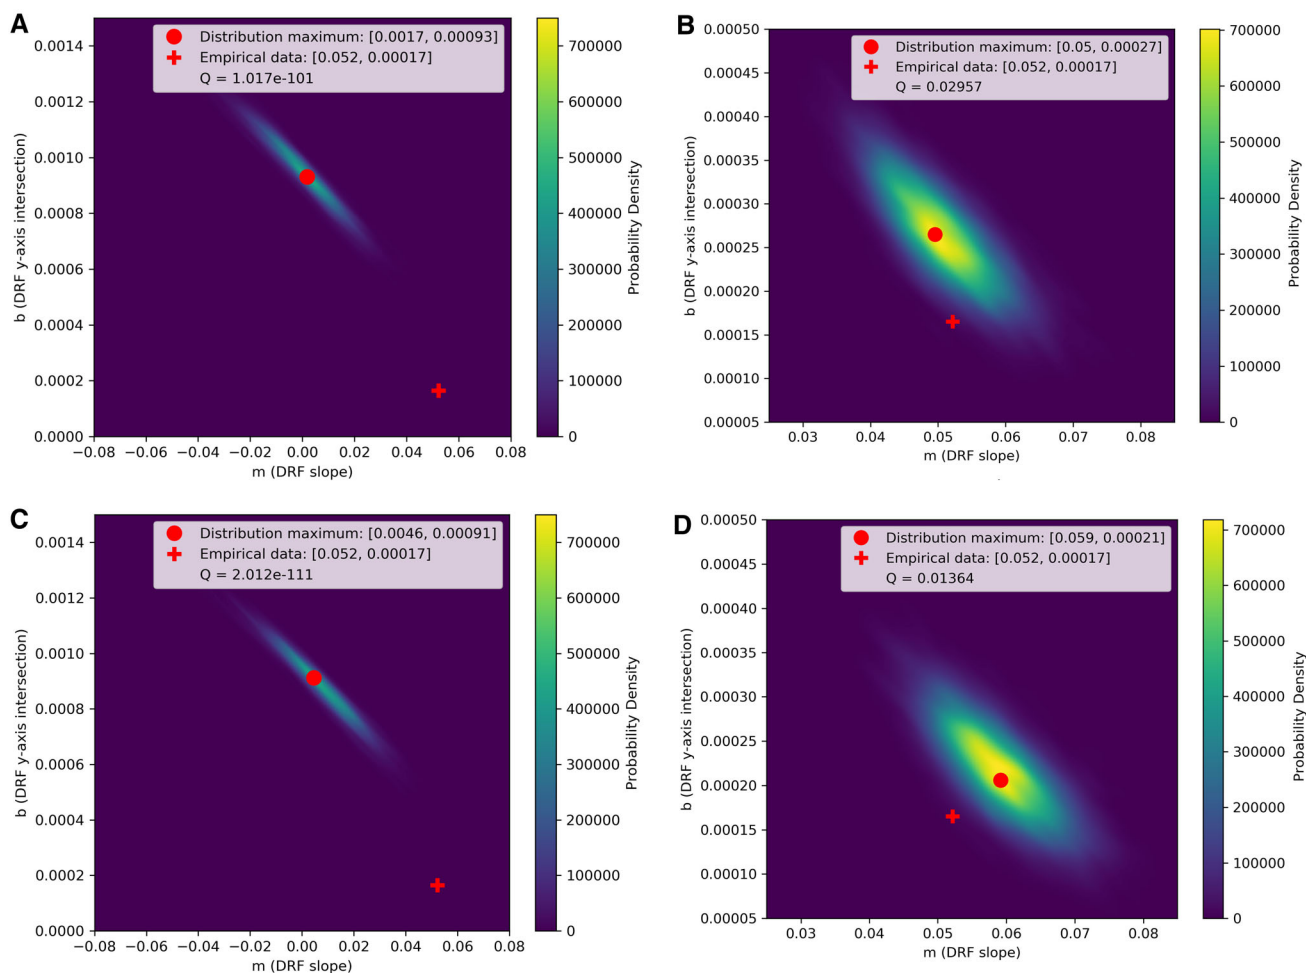Eur. Phys. J. Spec. Top. (2022) 231:1609–1624

Fig. 8 Parameter probability distributions of the linear fits on every surrogate data realization. The probability distributions are determined via kernel density estimation, and correspond to the surrogate ensembles for **A** $\mathcal{H}_0^1$, **B** $\mathcal{H}_0^2$, **C** $\mathcal{H}_0^3$ and **D** $\mathcal{H}_0^4$

References

1. J. Rockström, et al., Planetary boundaries: exploring the safe operating space for humanity. Ecol. Soc. **14**(2), 32 (2009)
2. W. Steffen et al., Planetary boundaries: guiding human development on a changing planet. Science **347**, 1259855 (2015)
3. I.M. Otto et al., Social tipping dynamics for stabilizing Earth's climate by 2050. PNAS **117**, 2354–2365 (2020)
4. D. Hoornweg, L. Sugar, C.L. Trejos Gómez, Cities and greenhouse gas emissions: moving forward. Environ. Urban. **23**, 207–227 (2011)
5. W. Paterson et al., Water footprint of cities: a review and suggestions for future research. Sustainability **7**, 8461–8490 (2015)
6. B.W. Brooks et al., Toxicology advances for 21st century chemical pollution. One Earth **2**, 312–316 (2020)
7. United Nations, *World Urbanization Prospects: The 2018 Revision* (2019)
8. P. Skougaard Kaspersen, N. Høegh Ravn, K. Arnbjerg-Nielsen, H. Madsen, M. Drews, Comparison of the impacts of urban development and climate change on exposing European cities to pluvial flooding. Hydrol. Earth Syst. Sci. **21**, 4131–4147 (2017)
9. D. Li, E. Bou-Zeid, Synergistic interactions between urban heat islands and heat waves: the impact in cities is larger than the sum of its parts. J. Appl. Meteorol. Climatol. **52**, 2051–2064 (2013)
10. R.V. Knight, Knowledge-based development: policy and planning implications for cities. Urban Stud. **32**, 225–260 (1995)
11. L.M.A. Bettencourt, J. Lobo, D. Helbing, C. Kühnert, G.B. West, Growth, innovation, scaling, and the pace of life in cities. PNAS **104**, 7301–7306 (2007)
12. P.J. Taylor, Specification of the world city network. Geogr. Anal. **33**, 181–194 (2001)
13. F. Sengers, A.J. Wieczorek, R. Raven, Experimenting for sustainability transitions: a systematic literature review. Technol. Forecast Soc. Change **145**, 153–164 (2019)
14. S. Irvine, X. Bai, Positive inertia and proactive influencing towards sustainability: systems analysis of a front-runner city. Urban Transform **1**, 1–27 (2019)
15. V. Castán Broto, H. Bulkeley, A survey of urban climate change experiments in 100 cities. Glob. Environ. Change **23**, 92–102 (2013)
16. W.V. Reid et al., Earth system science for global sustainability: grand challenges. Science **330**, 916–917 (2010)

17. D. Benson, A. Jordan, What have we learned from policy transfer research? Dolowitz and Marsh revisited. Polit. Stud. Rev. **9**, 366–378 (2011)

18. D. Dolowitz, D. Marsh, Who learns what from whom: a review of the policy transfer literature. Polit. Stud. **44**, 343–357 (1996)

19. D.P. Dolowitz, D. Marsh, Learning from abroad: the role of policy transfer in contemporary policy-making. Governance **13**, 5–23 (2000)

20. G. Marsden, K. Frick, A. May, E. Deakin, How do cities approach policy innovation and policy learning? A study of 30 policies in Northern Europe and North America. Transp. Policy **18**, 501–512 (2011)

21. G. Marsden, K.T. Frick, A.D. May, E. Deakin, Transfer of innovative policies between cities to promote sustainability: case study evidence. Transp. Res. Rec. J. Transp. Res. Board **2163**, 89–96 (2010)

22. D. Pojani, D. Stead, Going Dutch? The export of sustainable land-use and transport planning concepts from the Netherlands. Urban Stud. **52**, 1558–1576 (2015)

23. T. Lee, S. van de Meene, Who teaches and who learns? Policy learning through the C40 cities climate network. Policy Sci. **45**, 199–220 (2012)

24. G. West, *Scale: The universal laws of life, growth, and death in organisms, cities, and companies* Penguin Books, New York, 2017)

25. M. Barthelemy, The statistical physics of cities. Nat. Rev. Phys. **1**, 406–415 (2019)

26. L. Bettencourt, G. West, A unified theory of urban living. Nature **467**, 912–913 (2010)

27. M. Jusup, et al., Social physics. Phys. Rep. **948**, 1–148 (2022)

28. M.E. Newman, The structure and function of complex networks. SIAM Rev. **45**, 167–256 (2003)

29. R. Albert, A.-L. Barabási, Statistical mechanics of complex networks. Rev. Mod. Phys. **74**, 47 (2002)

30. P.S. Dodds, D.J. Watts, A generalized model of social and biological contagion. J. Theor. Biol. **232**, 587–604 (2005)

31. D. Centola, M. Macy, Complex contagions and the weakness of long ties. Am. J. Sociol. **113**, 702–734 (2007)

32. D. Helbing et al., Saving human lives: what complexity science and information systems can contribute. J. Stat. Phys. **158**, 735–781 (2015)

33. D.J. Daley, J. Gani, *Epidemic Modelling: An Introduction. Cambridge Studies in Mathematical Biology* (Cambridge University Press, Cambridge, 1999)

34. S.V. Buldyrev, R. Parshani, G. Paul, H.E. Stanley, S. Havlin, Catastrophic cascade of failures in interdependent networks. Nature **464**, 1025–1028 (2010)

35. K. Nyborg et al., Social norms as solutions. Science **354**, 42–43 (2016)

36. M. Tsvetkova, M.W. Macy, The social contagion of generosity. PLoS ONE **9**, e87275 (2014)

37. S.B. Rosenthal, C.R. Twomey, A.T. Hartnett, H.S. Wu, I.D. Couzin, Revealing the hidden networks of interaction in mobile animal groups allows prediction of complex behavioral contagion. Proc. Natl. Acad. Sci. USA **112**, 4690–4695 (2015)

38. J.D. Tàbara et al., Positive tipping points in a rapidly warming world. Curr. Opin. Environ. Sustain. **31**, 120–129 (2018)

39. J.D. Farmer et al., Sensitive intervention points in the post-carbon transition. Science **364**, 132–134 (2019)

40. S. Sharpe, T.M. Lenton, Upward-scaling tipping cascades to meet climate goals: plausible grounds for hope. Clim. Policy **21**, 421–433 (2021)

41. H. Levinson, S. Zimmerman, J. Clinger, G. Rutherford, Bus rapid transit: an overview. J. Public Transp. **5**(2), 1–30 (2002)

42. A. Wood, The politics of policy circulation: unpacking the relationship between South African and South American Cities in the adoption of bus rapid transit. Antipode **47**, 1062–1079 (2015)

43. N. Matsumoto, Analysis of policy processes to introduce Bus Rapid Transit systems in Asian cities from the perspective of lesson-drawing. In *Air Pollution Control in the Transportation Sector: Third Phase Research Report of the Urban Environmental Management Project*, 351–376 (Institute for Global Environmental Strategies, Hayama, 2007)

44. D.P. Rosati, M.H. Woolhouse, B.M. Bolker, D.J.D. Earn, Modelling song popularity as a contagious process. Proc. R. Soc. A **477**, 20210457 (2021)

45. C.N. Haas, Conditional dose–response relationships for microorganisms: development and application. Risk Anal. **22**, 455–463 (2002)

46. P.S. Dodds, D.J. Watts, Universal behavior in a generalized model of contagion. Phys. Rev. Lett. **92**, 218701 (2004)

47. N.O. Hodas, K. Lerman, The simple rules of social contagion. Sci. Rep. **4**, 4343 (2014)

48. J.F. Donges et al., Dose–response functions and surrogate models for exploring social contagion in the Copenhagen networks study. Eur. Phys J. Spec. Top. **230**, 3311–3334 (2021)

49. R. Vicente, M. Wibral, M. Lindner, G. Pipa, Transfer entropy—a model-free measure of effective connectivity for the neurosciences. J. Comput. Neurosci. **30**, 45–67 (2011)

50. M. Casdagli, Chaos and deterministic versus stochastic non-linear modelling. J. Roy. Stat. Soc. Ser. B (Methodol.) **54**, 303–328 (1992)

51. L. Gauvin, et al., Randomized reference models for temporal networks. arXiv:1806.04032 [physics, q-bio] (2020)

52. P. Holme, J. Saramäki, Temporal networks. Phys. Rep. **519**, 97–125 (2012)

53. M. Génois, C.L. Vestergaard, C. Cattuto, A. Barrat, Compensating for population sampling in simulations of epidemic spread on temporal contact networks. Nat. Commun. **6**, 8860 (2015)

54. F. Karimi, P. Holme, Threshold model of cascades in empirical temporal networks. Physica A **392**, 3476–3483 (2013)

55. J. Theiler, S. Eubank, A. Longtin, B. Galdrikian, J. Doyne Farmer, Testing for nonlinearity in time series: the method of surrogate data. Physica D **58**, 77–94 (1992)

56. T. Schreiber, A. Schmitz, Surrogate time series. Physica D **142**, 346–382 (2000)

57. Openflights, Air route and airport databases. Available at: https://openflights.org

58. P.J. Taylor, Specification of the world city network. Geogr. Anal. **33**, 181–194 (2010)

59. The Wikimedia Foundation, Wikidata query service. https://query.wikidata.org/

60. BRT + Centre of Excellence and EMBARQ, Global BRTData. Version 3.47. Last modified: July 07, 2019. Available at: http://www.brtdata.org

61. T. Deng, J.D. Nelson, Recent developments in bus rapid transit: a review of the literature. Transp. Rev. **31**, 69–96 (2011)

62. A. Agresti, B.A. Coull, Approximate is better than "Exact" for interval estimation of binomial proportions. Am. Stat. **52**, 119–126 (1998)

63. C.R. Shalizi, A.C. Thomas, Homophily and contagion are generically confounded in observational social network studies. Sociol. Methods Res. **40**, 211–239 (2011)

64. V. Venema, S. Bachner, H.W. Rust, C. Simmer, Statistical characteristics of surrogate data based on geophysical measurements. Nonlinear Process. Geophys. **13**, 449–466 (2006)

65. J.A. Scheinkman, B. LeBaron, Nonlinear dynamics and stock returns. J. Bus. **62**, 311–337 (1989)

66. W.S. Pritchard, D.W. Duke, K.K. Krieble, Dimensional analysis of resting human EEG II: surrogate-data testing indicates nonlinearity but not low-dimensional chaos. Psychophysiology **32**, 486–491 (1995)

67. M. Wiedermann, J.F. Donges, J. Kurths, R.V. Donner, Spatial network surrogates for disentangling complex system structure from spatial embedding of nodes. Phys. Rev. E **93**, 042308 (2016)

68. S. Maslov, K. Sneppen, A. Zaliznyak, Detection of topological patterns in complex networks: correlation profile of the internet. Physica A **333**, 529–540 (2004)

69. S. Maslov, K. Sneppen, Specificity and stability in topology of protein networks. Science **296**, 910–913 (2002)

70. J. Theiler, D. Prichard, Constrained-realization Monte-Carlo method for hypothesis testing. Physica D **94**, 221–235 (1996)

71. B.W. Silverman, Density estimation for statistics and data analysis, in *Monographs on Statistics and Applied Probability*, vol. 26, (Chapman and Hall, London, 1986)

72. R. Winkelmann et al., Social tipping processes towards climate action: a conceptual framework. Ecol. Econ. **192**, 107242 (2022)

73. F. Kraas, et al., WBGU–German Advisory Council on Global Change, *Humanity on the move: Unlocking the transformative power of cities. Summary* (WBGU, Berlin, 2016)

74. T.M. Lenton, Tipping positive change. Philos. Trans. R. Soc. B **375**, 20190123 (2020)

75. D. Brockmann, D. Helbing, The hidden geometry of complex, network-driven contagion phenomena. Science **342**, 1337–1342 (2013)

76. P. Taylor, G. Catalano, D. Walker, Measurement of the World City Network. Urban Stud. **39**, 2367–2376 (2002)

77. Creative Commons, Attribution-noncommercial-NoDerivs 3.0 unported (CC BY-NC-ND 3.0). Available at: https://creativecommons.org/licenses/by-nc-nd/3.0/

78. Open Knowledge Foundation, Open Database License. Available at: https://opendatacommons.org/licenses/odbl/1-0/

79. Open Knowledge Foundation, Database Contents License. Available at: https://opendatacommons.org/licenses/dbcl/1-0/

80. Creative Commons, CC0 1.0 universal public domain dedication. Available at: https://creativecommons.org/publicdomain/zero/1.0/