Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

# Cities have an effective temperature: A statistical mechanics description of user equilibrium

**Luis Eduardo Olmos Sánchez**

leolmoss@eafit.edu.co

Universidad EAFIT

**Article**

**Keywords:** user equilibrium, Beckmann potential, maximum entropy, volume-overcapacity, effective temperature, urban congestion

**Posted Date:** November 18th, 2026

**DOI:** https://doi.org/10.21203/rs.3.rs-11541700/v1

**License:** This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

**Additional Declarations:** No competing interests reported.

# Cities have an effective temperature: A statistical mechanics description of user equilibrium


## Abstract

Every morning, millions of drivers choose their routes independently, yet the traffic they produce follows a strikingly simple pattern. Physics has long tried to describe cities as it describes gases, with a few macroscopic quantities, but it has lacked the equivalent of temperature: a single number that tells how strongly a road network is loaded. We show that such a temperature emerges from user equilibrium, the standard model of route choice used by transportation engineers. In six cities as different as Boston, Lisbon, Rio de Janeiro and Beijing, how full each road is relative to its capacity follows the same exponential distribution, the maximum-entropy (least-biased) distribution compatible with the average load. Its single scale acts as an effective temperature of the city. This pattern persists when demand changes, in an independent dataset of 20 US cities, and in simulations where drivers never reach equilibrium. The temperature reveals what aggregate demand hides: cities equally hot can differ by 30% in their ratio of demand to road supply, depending on which roads carry the traffic. It predicts, with no adjustable parameters, how many roads become heavily loaded as demand grows, and a single law links it to travel times in every city. The effective temperature offers a common language for engineering models of traffic and the physics of congestion, from the percolation of traffic jams to the macroscopic fundamental diagram.



**Keywords:** *user equilibrium; Beckmann potential; maximum entropy; volume-over-capacity; effective temperature; urban congestion.*


## Introduction

Describing urban traffic with the language of physics has a long history. Millions of travelers interact indirectly through a shared road network, and their individual route choices produce collective patterns that are far simpler than the underlying microscopic dynamics might suggest. Herman and Prigogine, for example, described town traffic through moving and stopped vehicle populations in analogy with a two-fluid system, while subsequent work developed kinetic and fluid descriptions of traffic flow and network-level relations such as the macroscopic fundamental diagram (MFD)[1–7]. Percolation theory has shown how congested roads organize into clusters that grow and fragment the network as demand increases[8–10]. Other approaches have shown that aggregate demand-to-supply ratios can organize differences in congestion and recovery times across cities[11,12]. These studies show that urban traffic can be described macroscopically, even though individual roads and trips are highly heterogeneous.

What is still missing is a variable that sets the loading scale of a road network. In thermodynamics, temperature plays this role: it controls how microscopic states are populated while remaining a property of the whole system. For traffic, several quantities describe congestion at the network level[13–16], but none does this for the

distribution of traffic across individual roads. Network averages do not say how demand is spread over heterogeneous roads: two cities can carry the same demand relative to their total road capacity and still load their individual roads very differently.

Here we approach this problem through the volume-to-capacity ratio of a road segment, $v_e = x_e/C_e$ where $x_e$ is the traffic flow and $C_e$ its capacity. Unlike absolute flow, $v_e$ measures how strongly a road is used relative to what it can carry, making links of different sizes and road classes directly comparable. Across the six metropolitan networks studied here, the cross-sectional distribution of equilibrium $v$ is approximately exponential over its dominant sub-capacity range (see Fig. 1), even though the cities have very different network structures, road hierarchies and travel demand. A similar pattern had been reported for Boston and the San Francisco Bay Area[17,18], but the origin and characteristic scale of this regularity have not been explained.

The equilibrium setting provides a natural starting point for understanding this pattern. Drivers independently choose routes to reduce their own travel times. When no traveler can improve their travel time by unilaterally changing route, the network satisfies Wardrop's first principle[19] and reaches user equilibrium. For separable and increasing link travel-time functions, Beckmann, McGuire and Winsten showed that these flows can be obtained by minimizing a convex potential[20–22],

$$min_x\, z(x) \;=\; \Sigma_{e\in E}\int_0^{x_e} t_e(\omega)\, d\omega \qquad (1),$$

subject to flow conservation and non-negativity. This potential fixes the equilibrium flows, but says nothing directly about how loading is distributed across links. Our question is therefore different from the usual assignment problem: once equilibrium has selected the flows, what fixes the macroscopic scale of their distribution across the network?

For the commonly used Bureau of Public Roads travel-time function[23], the Beckmann potential can be expressed in terms of $v_e$. Under moderate congestion, its value is dominated by a contribution linear in relative loading, so the equilibrium objective is controlled primarily by a weighted mean network load. This observation motivates a coarse-grained statistical description. If link identities and their spatial arrangement are discarded while retaining only a non-negative loading variable and its mean, the maximum-entropy reference distribution is exponential[24]. Its scale is the mean load, $T = \langle v \rangle$, and we refer to $T$ as an effective traffic temperature. The name is an analogy: drivers do not maximize entropy, the equilibrium flows are still fixed uniquely by the assignment problem, and $T$ is simply the scale of the loading distribution.

We test this interpretation using user-equilibrium assignments for six metropolitan networks: Boston, the San Francisco Bay Area, Lisbon, Porto, Rio de Janeiro, and Beijing[11,12,25,26]. We first show that link loads are exponentially distributed and estimate their temperature. We then examine how $T$ relates to macroscopic demand

and road supply, showing that it differs from a distance-weighted demand-to-supply ratio in a

way that reveals how loading is spread across the road hierarchy. We further test the robustness of the exponential form under changes in travel demand and assignment procedure, and against dynamic simulations in which no equilibrium optimization is imposed. Finally, we relate $T$ to city-level travel times and to independently measured congestion levels. Together, these results show that a single number captures how a city's equilibrium traffic is spread over its roads.

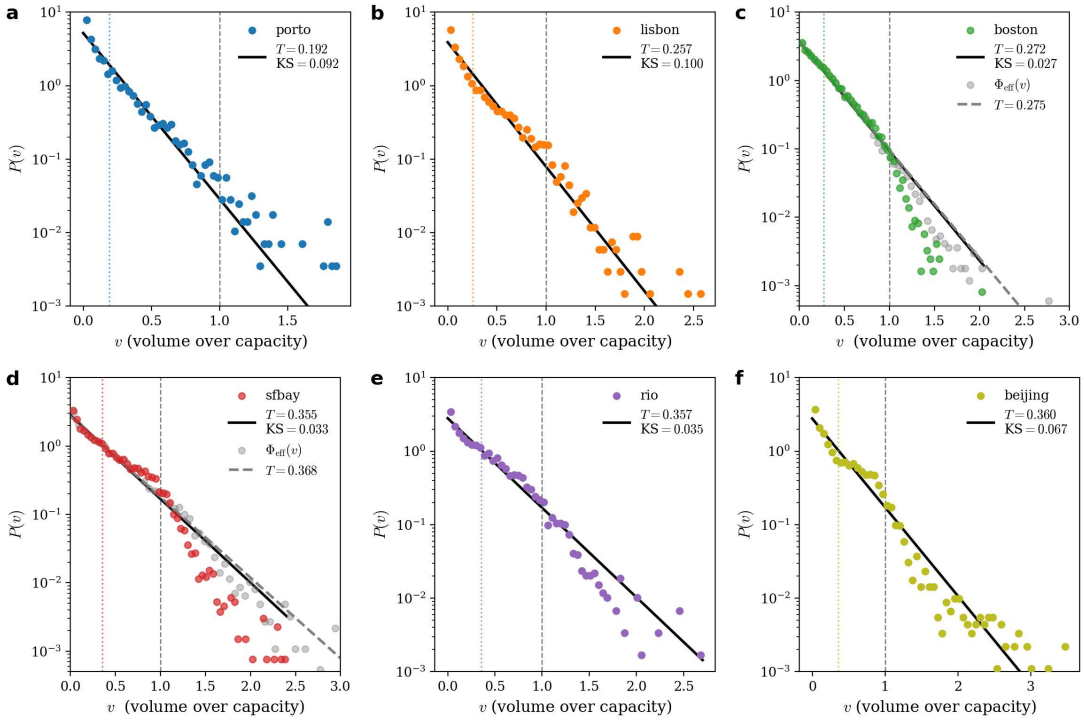



**Figure 1 | Exponential link loading at user equilibrium. (a–f)** Cross-sectional distributions of equilibrium volume-to-capacity ratios for the six metropolitan networks, evaluated over the retained higher-capacity links. Black curves show maximum-likelihood exponential fits. Vertical dashed lines indicate $v = T$, with $T = \langle v \rangle$, and nominal capacity, $v = 1$. For Boston and the San Francisco Bay Area, gray curves show the corresponding fits using the hinge-corrected effective cost of Eq. (9). Link-selection criteria are given in Methods.

## Results

### From user equilibrium to an exponential loading distribution

Minimizing the Beckmann potential (Eq. 1) determines the user-equilibrium link flows. When link travel-time functions are strictly increasing, the equilibrium is unique in link-flow space, although its decomposition into route flows need not be[21]. We represent link travel times using the standard Bureau of Public Roads (BPR) function[23],

$$t_e(x_e) = t_e^0 \left[1 + \alpha \left(\frac{x_e}{C_e}\right)^{\beta}\right], \quad (2)$$

where $t_e^0$ is the free-flow travel time and $\alpha$ and $\beta$ control how rapidly travel time increases with loading; we use the standard values α = 0.6 and β = 4 throughout (Methods). Substituting Eq. 2 into Eq. 1 and writing the equilibrium problem in terms of $v_e$ gives

$$min_{\{v_e\}} z(v_e) \; = \; \Sigma_{e \in E} \, w_e \left[v_e \; + \; \frac{\alpha}{\beta + 1} \cdot {v_e}^{\beta+1}\right] \equiv \; \Sigma_{e \in E} \, w_e \, \Phi(v_e). \; (3)$$

where $w_e \; = \; t_e^0 C_e$ and $\Phi(v_e) = \; v_e \; + \; \frac{\alpha}{\beta+1} \cdot {v_e}^{\beta+1}$.

The equilibrium problem therefore has a simple interpretation in the loading variable $v$: each link contributes an effective loading cost $\Phi(v_e)$, weighted by a link-specific factor $w_e$.

User equilibrium is solved over the complete road network $E$, but not every link is equally informative about how loading is distributed. On local streets, the estimated loading depends strongly on the spatial resolution of the origin–destination data, and many links carry essentially no assigned flow. We therefore describe the loading distribution over the links where it is well resolved: those with capacity $C_e > 1000 \frac{veh}{h}$ and equilibrium loading $v_e \geq v_{min}$, with $v_{min} = 0.01$ as the baseline threshold. We denote this subset by $\mathcal{E} \subset E$. It consists predominantly of arterials and highways (Methods; Supplementary Table S1 and Supplementary Note 1).

Most links in $\mathcal{E}$ operate below capacity. In this regime, the nonlinear contribution to $\Phi(v)$ is small relative to the linear term, so $\Phi(v) \approx v$. The nonlinear term still matters, because it penalizes heavily loaded links and so shapes the equilibrium, but over most of the range its contribution is small.

The leading-order contribution of the retained network can therefore be summarized by a weighted mean loading

$$\Gamma_w = \frac{\Sigma_{e \in \mathcal{E}} \, t_e^0 C_e \, v_e}{\Sigma_{e \in \mathcal{E}} \, t_e^0 C_e} \, . (4)$$

Because $C_e \, v_e = x_e$, the numerator is the aggregate free-flow travel time of the assigned flows, whereas the denominator is the corresponding free-flow-time-weighted capacity of the retained network. $\Gamma_w$ thus tells how loaded the retained network is, but not how that load is shared among roads.

To describe how this load is spread over links, we use the corresponding unweighted loading scale, $T \; \equiv \langle v \rangle_{\mathcal{E}}$. The weighted and unweighted loading scales are related exactly by

$$\Gamma_w = \frac{\langle wv \rangle_\mathcal{E}}{\langle w \rangle_\mathcal{E}} = T(1 + \chi_w); \quad \chi_w \equiv \frac{\mathrm{Cov}_\mathcal{E}(w, v)}{\langle w \rangle_\mathcal{E} \langle v \rangle_\mathcal{E}}. \quad (5)$$

The two scales therefore differ by a factor $1 + \chi_w$ (Supplementary Note 3). Across the six networks $|\chi_w| \leq 0.25$, so $\Gamma_w$ and T can differ by up to 25%. As shown below, the exponential form holds even where this coupling is largest, and the coupling itself carries information about the road hierarchy (Supplementary Table S2).

This suggests a coarse-grained description: we are interested in how loads are distributed across the network, not in which link carries which load. User equilibrium assigns to every link in $\mathcal{E}$ a load $v_e$, a weight $w_e$ and a position. If we discard link identities, treating links as exchangeable, the weights and positions are lost with them, and only the distribution of the non-negative variable $v$ remains, constrained by its normalization and its mean: $\int_0^\infty P(v)dv = 1$, $\int_0^\infty vP(v)dv = T$.

Many distributions satisfy these constraints. The one that maximizes the Shannon entropy, $S[P(v)] = -\int_0^\infty P(v) \log P(v)\, dv$ , is

$$P(v) = \frac{1}{T} e^{-v/T}. \quad (6)$$

The exponential is the least-structured distribution consistent with the information retained after coarse-graining[24]. Its only parameter is the mean load $T = \langle v \rangle$, the effective traffic temperature.

Maximum entropy does not force the equilibrium to be exponential, and it does not mean that drivers maximize entropy: the link flows remain the deterministic outcome of the assignment problem. Maximum entropy only provides a reference distribution once link identities, weights and positions are discarded. Whether real equilibria follow it is an empirical question, which we address next.

**Table 1: General properties of the six studied cities.**

| City | Links (total) | Retained $N_e$ (%) | Trips (7-8am) | $\bar{u}_f$ (km/h) | $T$ | $f_h$ (%) | $\Gamma$ | $\Gamma_w$ | $\chi_w$ |
|---|---|---|---|---|---|---|---|---|---|
| porto | 18 895 | 48.8 | 171 324 | 53 | 0.193 | 2.2 | 0.171 | 0.163 | −0.15 |
| lisbon | 32 756 | 47.8 | 323 584 | 49 | 0.257 | 6.3 | 0.205 | 0.194 | −0.24 |
| boston | 58 212 | 62.3 | 915 875 | 47 | 0.272 | 4.1 | 0.260 | 0.241 | −0.11 |
| sfbay | 73 929 | 44.9 | 1 015 120 | 55 | 0.355 | 11.2 | 0.429 | 0.382 | +0.08 |
| rio | 27 661 | 47.7 | 432 014 | 47 | 0.359 | 10.4 | 0.329 | 0.322 | −0.10 |
| beijing | 81 524 | 19.2 | 504 862 | 55 | 0.360 | 12.5 | 0.369 | 0.375 | +0.04 |

## Exponential link loading across six cities

Figure 1a–f shows that the equilibria follow Eq. (6). In every city $P(v)$ decays approximately exponentially over its dominant sub-capacity range, and the temperature $T$ spans nearly a factor of two, from $T = 0.193$ in Porto to $T = 0.360$ in Beijing. Kolmogorov–Smirnov distances between the empirical and fitted distributions are at most 0.10 in every city, reported here as a measure of discrepancy rather than as a formal test, since with samples of this size any goodness-of-fit test rejects. Fewer than 6.5% of retained links operate above nominal capacity, supporting the sub-capacity approximation used above.

A complementary view of the temperature comes from weighting each loading level by how much it contributes to the mean load. This contribution is $v \cdot P(v)$, and for an exponential distribution it peaks exactly at $v = T$: the temperature is both the mean load and the loading level that carries most of it (Fig. 2g for Boston; Supplementary Fig. S1 for the other cities). The analogy with thermal radiation is direct. Just as the colour of a hot body is set by the peak of its spectrum, the dominant loading level of a city is set by $T$. The maps of Fig. 2 use this weighting to display the equilibria: links near $T$ are drawn most prominently, so hotter cities (Beijing, Rio de Janeiro and the San Francisco Bay Area) appear bluer and colder ones redder. The weighting is a visual aid and highlights links near $T$ by construction; it does not add information beyond the distribution itself.

Lisbon shows what the temperature does not capture. Its temperature is among the lowest of the six cities, yet a compact cluster of heavily loaded roads stands out in the centre. The temperature sets the scale of the loading, not where it occurs: two cities with the same temperature can have their congested links spread across the network or concentrated in a few bottlenecks.

A two-parameter gamma distribution fits slightly better in most networks, but it has the same mean $T$, and its departure from the exponential is confined to links with almost no assigned flow (Methods, Supplementary Note 4). The absolute value of $T$ shifts with the lower cut-off $v_e \geq v_{min}$ in the networks with coarser demand data, by up to 15% in Porto, but the ordering of cities by $T$ is unchanged across the whole range (Supplementary Note 5).

Three further tests show that the exponential does not depend on how the flows are computed. First, it appears with assignment procedures different from ours: incremental assignment of mobile-phone demand in Boston and the San Francisco Bay Area, an approximation to user equilibrium, gave a mean scale of about 0.28[17], and Frank–Wolfe equilibria on independent networks of three Bay Area counties, loaded with census commuting demand, gave the same form[18]. Second, it holds on 20 independent US metropolitan networks from a published traffic dataset[27], with their own road data, capacity coding and LODES commuting demand (Supplementary Fig. S2). Third, and most stringently, it survives without equilibrium. We re-ran the mesoscopic simulations of ref. [8] for Rio de Janeiro and the San

Francisco Bay Area, in which vehicles reroute stochastically and queues form and spill back. The hourly distributions of $v$, computed with the same capacities and link selection as the static analysis, remain exponential, with temperatures between $0.107$ and $0.244$ (Methods, Supplementary Fig. S3). The exponential therefore does not depend on our solver, our networks, our demand estimates or the Beckmann argument that motivated it.

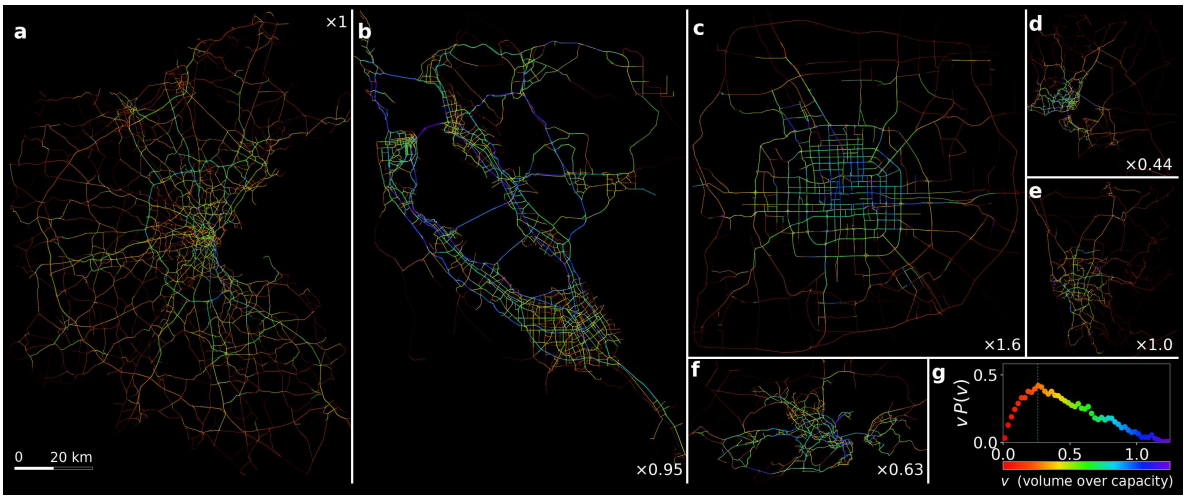



**Figure 2 | Spatial distribution of loading at user equilibrium.** (**a**–**f**) Volume-to-capacity ratios $v$ at user equilibrium for (**a**) Boston, USA; (**b**) the San Francisco Bay Area, USA; (**c**) Beijing, China; (**d**) Lisbon, Portugal; (**e**) Porto, Portugal; and (**f**) Rio de Janeiro, Brazil. Road segments are colored by $v$ on a visible-spectrum scale, from red at low loading through green to blue and violet at high loading. Their visual prominence is weighted by $v\,\mathrm{P}(v)$, the contribution of each loading level to the mean load, so that links near the effective temperature $T = \langle v \rangle$ stand out while lightly and heavily loaded links recede. Higher-temperature cities therefore appear predominantly bluer, and lower-temperature cities redder. Numbers in each panel give the map scale relative to Boston; the scale bar in **(a)** refers to Boston. **(g)** $v\,\mathrm{P}(v)$ for Boston; for an exponential distribution it peaks at $v = T$, which sets the prominence weighting of the maps. The other cities are shown in Supplementary Fig. S1. Maps © OpenStreetMap contributors, ODbL.

## Effective temperature as a macroscopic congestion scale

We now compare the temperature with two aggregate indices computed from the same equilibrium flows. Alongside $\Gamma_w$ of Eq. (4), Çolak et al.[11] introduced on heuristic grounds a distance-weighted demand-to-supply ratio,

$$\Gamma = \frac{\Sigma_{e\in\mathcal{E}}\, l_e x_e}{\Sigma_{e\in\mathcal{E}}\, l_e C_e}, \quad (7)$$

with $l_e$ the length of link *e*, and showed that it organizes characteristic congestion and recovery times across cities. Our framework explains why $\Gamma$ *works: it is* the length-weighted analogue of Eq. (4), and the two coincide when free-flow speeds are approximately uniform across the retained network. Its numerator is the total vehicle-distance traveled on the retained network, so $\Gamma$ compares aggregate travel

against aggregate capacity without reference to how that travel is distributed among links.

Figure 3a compares the two measures against $\Gamma_w$, so the vertical distance to the diagonal measures how far each departs from the weighted mean load.The distance-weighted ratio stays close to it ($|\Gamma - \Gamma_w| \leq 0.019$ in five cities), because free-flow speeds are fairly uniform within each network; the largest gap, $0.047$, occurs in the freeway-dominated San Francisco Bay Area, where speeds are most heterogeneous.

The temperature departs more, and in different directions in different cities. Rio de Janeiro and the San Francisco Bay Area have indistinguishable temperatures ($T$ = 0.359 and 0.355), yet their demand-to-supply ratios differ by nearly 30% ($\Gamma$ = 0.329 and 0.429). By Eq. (5), this gap is the coupling, $\chi_w = \frac{\Gamma_w - T}{T}$ between link weight and relative load. Positive $\chi_w$ means that load concentrates on high-capacity roads; negative $\chi_w$, that smaller roads carry a disproportionate share. The coupling is positive in the San Francisco Bay Area ($\chi_w = +0.08$), whose freeways absorb more load than their share of links, and weakly so in Beijing $\chi_w = +0.04$); it is negative in the other four cities, most strongly in Lisbon ($-0.24$). The distribution remains exponential even at this level of coupling, so link weights shape the aggregate but not the form of $P(v)$. $T$ and $\chi_w$ thus capture, respectively, how loaded a network is and which part of the road hierarchy carries that load.

The temperature also sets city-level travel times. We describe the mean travel time of trips of route distance $d$ by

$$t(d) \simeq \frac{d}{\bar{u}_f}\left[1 + \alpha^* T^{\beta^*}\right], \quad (8)$$

where $\bar{u}_f$ is a characteristic free-flow speed and $\alpha^*$and $\beta^*$are effective route-level parameters. Equation (8) is not an average of the BPR function. We first treat it as an empirical law and ask whether one parameter pair fits all six cities; the next section explains its form and the relation between its two parameters.

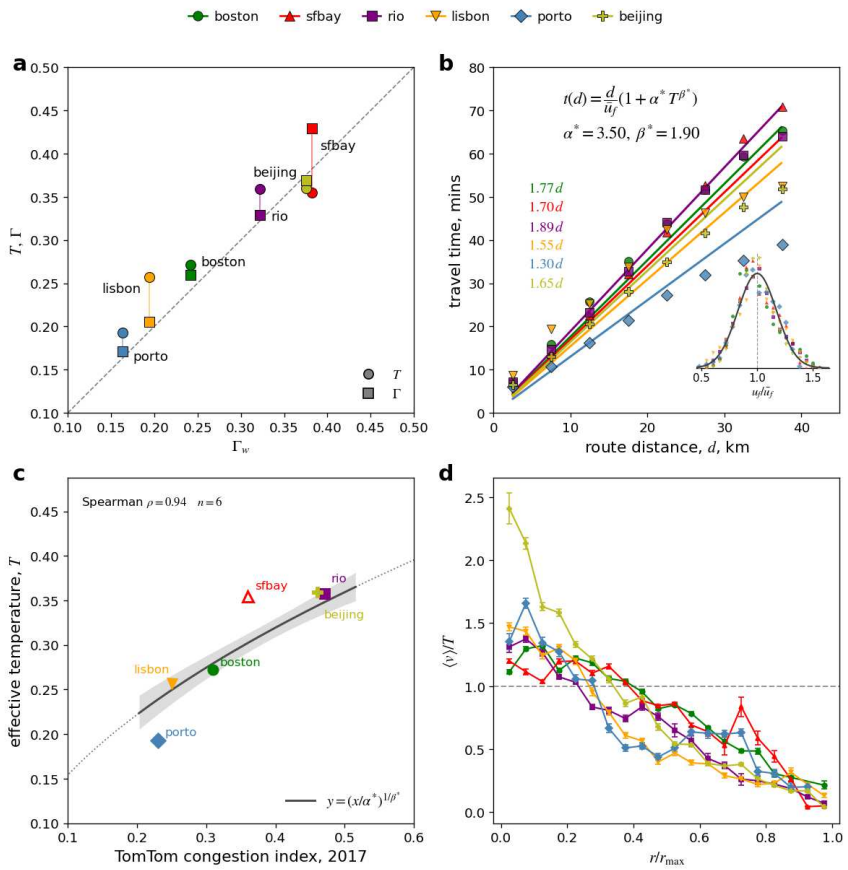



**Figure 3 | Loading indices, travel times, and spatial structure of loading. (a)** Effective temperature $T$ (circles) and demand-to-supply ratio $\Gamma$ (squares) versus weighted loading $\Gamma_w$; the dashed line marks equality and vertical segments join values for the same city. The distance between $T$ and $\Gamma_w$ is the hierarchy coupling $\chi_w$ of Eq. (5). **(b)** Mean equilibrium travel time against route distance $d$. Lines show Eq. (8) with a single global pair $(\alpha^*, \beta^*)$; numbers give the resulting city-specific slopes in min $km^{-1}$. Inset: route free-flow speeds, normalized by the city mean and weighted by route flow. **(c)** Effective temperature against the 2017 TomTom congestion index [28], the fractional increase in travel time over free-flow conditions. The solid curve inverts Eq. (8) with the parameters of panel **b**, without further fitting; the band is the 1σ uncertainty propagated from their covariance, and the curve is dotted outside the range of the data. The open symbol marks the San Francisco Bay Area, for which the index covers only the city of San Francisco. (d) Local mean load normalized by the city temperature, $\langle v \rangle(r)/T$, against normalized distance from the city centre, $r/r_{max}$; error bars are standard errors within each radial bin, and the dashed line marks $\frac{\langle v \rangle(r)}{T} = 1$. Colours and markers are consistent across panels.

Figure 3b shows that this form describes all six cities with a single global pair of parameters, $\alpha^* = 3.50$ and $\beta^* = 1.90$. No city-specific parameter is needed once $\bar{u}_f$ and $T$ are known, even though both the free-flow speed of a route and the congestion met along it vary with route length. The two parameters are not independently constrained: over the range of temperatures spanned by the six cities, $\alpha^*$ and $\beta^*$ trade off along a narrow ridge (correlation 0.99), which we mapped by sweeping $\beta^*$over [0.5, 6.0] and refitting α* (Supplementary Note 6 and Supplementary Fig. S5). The curve is well determined, but the two values are not; in particular, $\beta^*$ should not

be read as a route-level analogue of the BPR exponent. The next section explains the direction of this degeneracy.

Re-arranging Eq. (8) gives a direct mapping between the effective temperature and fractional travel-time delay, $(t(d) - t_0)/t_0 \simeq \alpha^* T^{\beta^*}$ (with $t_0 = \frac{d}{\bar{u}_f}$), which is independent of trip distance. Across the observed range of temperatures this implies delays from about 15% at T = 0.193 to about 50% at T = 0.360. Figure 3c shows that this mapping orders cities consistently with the TomTom congestion index[28] (Spearman $\rho = 0.94$, $n = 6$), despite differences in spatial coverage and temporal averaging between the two quantities; with six cities this is a consistency check rather than a validation.

The temperature is a network-wide number, not a map of congestion. Figure 3d shows that local mean loading decreases systematically with distance from the urban core, while Eq. (5) shows that loading may also be distributed differently across the road hierarchy. T sets the overall scale of network use, while the radial profile and $\chi_w$ say where that load sits. $T$ alone also does not tell how many links are heavily loaded; for that we need the nonlinear part of the link cost.

**Beyond the linear loading regime**

The sub-capacity exponential follows from the approximation $\Phi(v) \approx v$. As loading approaches capacity, however, the nonlinear BPR contribution can no longer be neglected. The exact effective cost is $\Phi(v) = v + \frac{\alpha}{\beta+1} \cdot v^{\beta+1}$, which grows much faster than $v$ near capacity.

A useful first-order description is obtained by linearizing the nonlinear contribution about $v = 1$. This tangent becomes negative at low loading, where the exact correction is negligible; keeping it only where it is positive gives the hinged effective cost (Supplementary Note 2).

$$\Phi_{\text{eff}}(v) = v + \alpha \, max\,(0, v - v_h);\; v_h = \frac{\beta}{\beta + 1} \quad (9)$$

For $\beta = 4$, $v_h = 0.8$: links begin to feel the nonlinear cost at about 80% of nominal capacity, where the exact correction already amounts to about 5% of the linear term for our baseline $\alpha = 0.6$. Because $\Phi_{\text{eff}}(v_h) = v_h$, the fraction of links beyond the breakpoint is the same whether measured in $v$ or in $\Phi_{\text{eff}}$, a property we use below. In this section the temperature is evaluated on the effective cost, $T(\lambda) \equiv \langle \Phi_{\text{eff}}(v) \rangle$, which differs from $\langle v \rangle$ by at most 5% at baseline demand.

Varying $\alpha$ provides a direct test of this approximation. Figure 4a shows equilibrium loading distributions for Boston computed with BPR coefficients from $\alpha = 0.15$ to $\alpha = 0.8$. Although the distributions in $v$ change as larger $\alpha$ increasingly penalizes highly loaded links, expressing them through $\Phi_{\text{eff}}(v)/T$, with the breakpoint and

slope change determined by each value of $\alpha$, produces an approximate collapse onto the unit exponential. No parameter is fitted: the tail follows from the link cost that generated the equilibrium. The effective temperature itself varies by less than 1% across this range, so $\alpha$ mainly redistributes a given load, while T sets its overall scale.

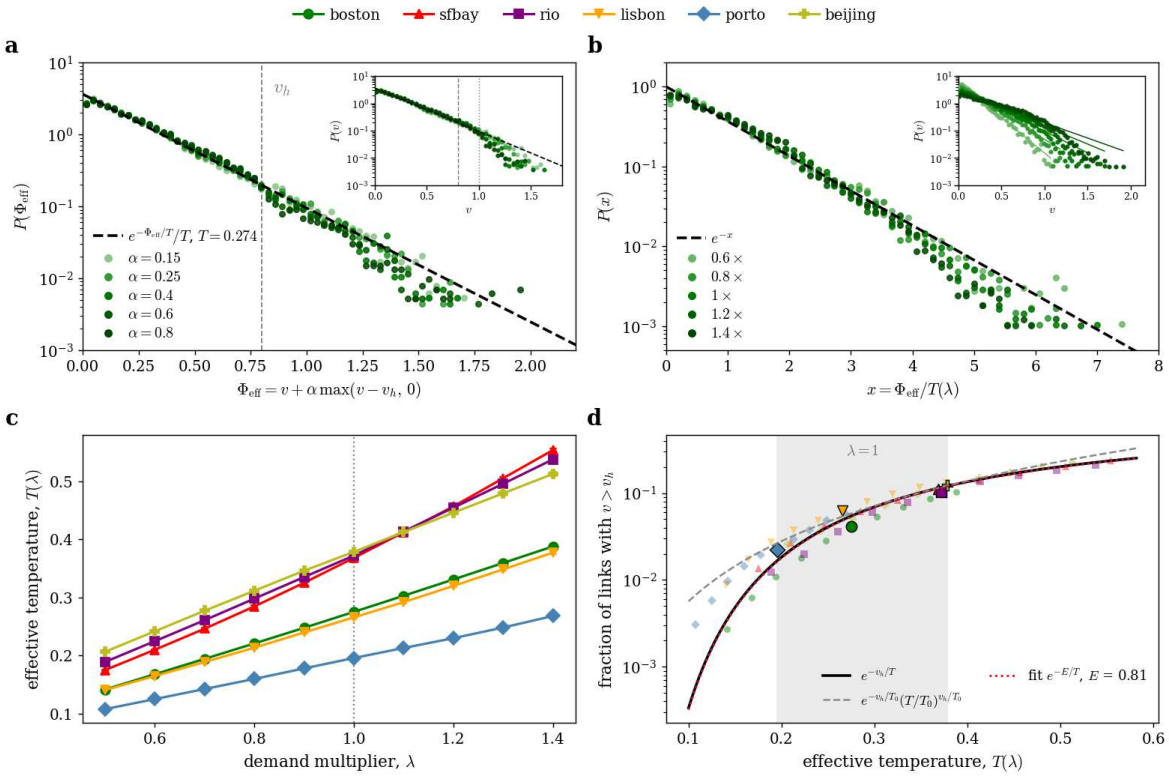



**Figure 4 | Beyond the linear loading regime.** (a) Equilibrium loading distributions for Boston computed with BPR coefficients from $\alpha = 0.15$ to $0.80$, expressed through the hinge-corrected cost $\Phi_{\text{eff}}(v) = v + \alpha\, max\,(0,\ v - v_h)$; the dashed line is the exponential with $T = 0.274$ and the vertical line marks the breakpoint $v_h = \beta/(\beta + 1) = 0.8$. Inset: the same equilibria in the raw variable $v$, where the distributions differ. (b) The same construction across demand multipliers $\lambda = 0.6$–$1.4$, expressed as $x = \Phi_{\text{eff}}(v)/T(\lambda)$, with the unit exponential $e^{-x}$ dashed. (c) Effective temperature versus demand multiplier for the six networks. (d) Fraction of retained links beyond the breakpoint, $f_h = P(v > v_h)$, against effective temperature. Small symbols are all six cities across all demand levels; large symbols mark the baseline states (λ = 1), whose temperature range is shaded. The solid and dashed curves carry no adjustable parameters: Eq. (10), $f_h = exp(-v_h/T)$, and its local power-law form about $T_0 = 0.347$, with exponent $\frac{v_h}{T_0} = 2.3$. The two nearly coincide within the shaded range and separate outside it. The dotted curve is the only fit in the panel, a single exponential $exp(-E/T)$ across all cities and demand levels, giving $E = 0.81 \pm 0.08$, consistent with the predicted $v_h = 0.8$. In all panels, T is the effective cost scale $\langle \Phi_{\text{eff}}(v) \rangle$, not $\langle v \rangle$. Colors and markers are consistent across panels.

The same rescaling works when demand changes. We scale Boston's origin–destination matrix by a factor $\lambda$, from $\lambda = 0.5$ to $1.4$ of its baseline morning value, holding the network, capacities, and relative spatial distribution of trips fixed. Figure 4b shows that the resulting distributions again collapse approximately when

expressed as $\Phi_{\mathrm{eff}}(v)/T(\lambda)$. Increasing demand therefore acts mainly by changing the characteristic scale $T$, while the nonlinear correction accounts for the steepening of the high-loading tail. The collapse is not exact, since origin–destination constraints and changes in the set of used links continue to shape each equilibrium.

The number of links in the nonlinear regime behaves very differently. Equation (9) provides a natural threshold, so we define $f_h \equiv P(v > v_h)$, the fraction of links loaded beyond the breakpoint. In the exponential approximation,

$$f_h(T) = exp(-v_h/T). \quad (10)$$

This relation contains no adjustable quantity: $v_h$ is fixed by the BPR exponent alone, so Eq. (10) predicts $f_h$ from the effective temperature and nothing else. Figure 4d tests it directly. Measured fractions for all six cities and all demand levels, spanning temperatures from $0.11$ to $0.55$, collapse onto a single exponential in $1/T$ whose fitted decay constant, $E = 0.81 \pm 0.08$, is consistent with the predicted $v_h = 0.8$. Six very different cities, across all demand levels, thus follow one relation between the temperature and the share of heavily loaded links. The figure also shows the local power-law form of Eq. (10) about $T_o = 0.347$, the temperature at which the travel-time fit of Fig. 3b is anchored (Supplementary Note 6; evaluated on $\langle v \rangle$), drawn with the same predicted $v_h$. The two forms nearly coincide within the narrow band of temperatures spanned by the six cities at $\lambda = 1$ and diverge outside it: varying demand is what makes the exponential form testable.

Combined with the approximately linear $T(\lambda)$ of Fig. 4c, Eq. (10) implies that a nearly linear increase in the macroscopic loading scale translates into a rapidly growing population of highly loaded links: over the demand range explored, $T(\lambda)$ roughly doubles in all six networks while $f_h$ grows by an order of magnitude in the hotter cities and by substantially more in the colder ones.

This result explains the form of the route-level travel-time law of Eq. (8), and suggests a connection with the two-fluid model of Herman and Prigogine [1]. In that description, travel-time delay is controlled by the fraction of stopped vehicles, $f_s$, through $t/t^0 = (1 - f_s)^{-(n+1)}$, where $n$ characterizes network quality. If, at the network level, the stopped fraction scales with the population of links beyond the breakpoint, $f_s \propto f_h$, then for small $f_h$ one obtains $\frac{t}{t^0} - 1 \propto f_h = exp(-v_h/T)$. The route-level delay would then be governed by the same temperature-dependent abundance of highly loaded links, with no reference to the microscopic BPR function. This connection remains to be tested directly: $f_s$ is a kinematic property of vehicles, measured dynamically, whereas $f_h$ is a static property of the loading distribution, and the linearization holds only while $f_h$ remains small.

This also accounts for the behavior of the fitted parameters. Over a narrow temperature range, $exp(-v_h/T)$ looks like a power law with exponent $v_h/T_o$: $f_h(T) \simeq f_h(T_o)\,(T/T_o)^{v_h/T_o}$. A power law is therefore the local representation of an underlying exponential, which is why Eq. (8) describes the data well while its two coefficients

are not separately determined. Expanding about the same $T_o$ predicts $\beta^* = \frac{v_h}{T_0} \approx 2.3$. The free fit gives $\beta^* = 1.90$, but its profile is shallow: any exponent in [1.27, 2.73] describes the six cities within the $95\%$ criterion, and the predicted value lies inside that range (Supplementary Fig. S5). The comparison is approximate, since the travel-time fit uses $T = \langle v \rangle$ whereas Eq. (10) uses $\langle \Phi_{eff} \rangle$, which differ by at most 5% at baseline; evaluating both on the same variable shifts $\beta^*$ and $v_h/T_o$ by similar amounts and does not change this conclusion. $\beta^*$ thus measures how fast the share of heavily loaded links grows with temperature

The framework has a clear limit. The BPR function is static and separable: it assigns each link a cost regardless of the state of its neighbors. Once persistent queues form, flow can decrease as density increases, congestion propagates through spillback, and neighboring links become dynamically coupled; a severely congested link may even carry a realized flow below nominal capacity. This is the congested branch of the macroscopic fundamental diagram, and there the volume-to-capacity ratio alone is no longer a sufficient state variable: density, queue length and unmet demand are also needed. The dynamic simulations suggest that the exponential form persists while queues remain limited, but describing the fully congested regime requires extending the framework beyond static equilibrium.

## Discussion

Physical descriptions of urban traffic have produced several macroscopic regularities: the two-fluid model, which characterizes a network through the fraction of stopped vehicles, the macroscopic fundamental diagram, which relates accumulation to production[4–6], and aggregate demand-to-supply ratios, which organize congestion and recovery times across cities [11,12]. Each summarizes how much traffic a network carries; none sets the scale on which that traffic is distributed over individual roads. The effective temperature provides that scale, and its origin is what distinguishes it from earlier proposals. Recent scaling approaches to urban congestion postulate their macroscopic variables[16]. Statistical physics has treated optimization problems as disordered systems[29–31] and has been used to optimize routing on transport networks[32], but always to solve the routing problem, not to describe the loading pattern that equilibrium produces. Here the temperature is identified by the equilibrium problem that transportation engineers already solve: the linear part of the Beckmann potential makes the mean load the relevant constraint, and maximum entropy, read as statistical inference[24], turns that constraint into a temperature. Assignment outputs computed with BPR-type costs can therefore be read in thermodynamic-like terms. This also explains the exponential-like load distributions reported earlier for Boston and the San Francisco Bay Area[17,18], whose mean scale is close to the temperature we find for Boston.

The temperature also gives a common variable to descriptions that have developed separately. The fraction of heavily loaded links fixed by $T$ is formally the occupation probability of percolation theory, which describes how congested roads organize into clusters that fragment urban networks[9,10] and, combined with the MFD, how

congestion propagates[8]. Percolation captures what the temperature deliberately leaves out, namely topology and spatial correlations, while the temperature supplies its input. If heavily loaded links play the role of stopped vehicles, the travel-time law implied by $T$ also connects with the two-fluid model and, through it, with the MFD. In this picture $T$ acts as a control parameter, and congested links are the natural place to look for an order parameter. Their abundance varies smoothly with temperature, but their connectivity need not: as $T$ increases, isolated jams can merge into network-spanning clusters, and the size of the largest one is a natural candidate order parameter for the transition from free-flowing to globally congested states. Locating this transition in each city would turn the temperature from a descriptive scale into a predictor of network-wide breakdown.

Three questions remain open. The first is whether $T$ is a genuine state variable: it should be robust to the choice of subnetwork, under capacity-based, geographical or random partitions, and should be recovered when the same demand is redistributed over the same network. The second concerns entropy, which is inferential here because user equilibrium selects a deterministic configuration; day-to-day fluctuations and route adaptation[33,34] provide the natural ensemble in which it could acquire a physical meaning. The third is generality: the persistence of the exponential on independent networks and without equilibrium suggests a property of how urban networks spread demand over capacity that goes beyond the Beckmann program, and the roles of route choice, topology, capacity coding and the spatial structure of demand remain to be separated.

Our results have limitations. Temperatures depend on the resolution of the demand data. The travel-time law was fitted in sample on six cities, with a free-flow speed held constant per city, and its comparison with independent congestion indices rests on six points. The framework is static and does not reach the fully congested regime, where queues and spillback dominate.

In summary, an effective temperature emerges from user equilibrium in six very different cities and links the equilibrium models of transportation engineering with the statistical physics of networks. It does not yet amount to a thermodynamics of cities, but it supplies the variable such a theory needs.

## Methods

**Data and networks.** Morning peak-hour origin–destination matrices were inferred from mobile-phone call-detail records for six metropolitan regions. For Boston, the San Francisco Bay Area, Lisbon, Porto, and Rio de Janeiro we follow the demand-estimation and network-construction procedures of [11]; the Beijing origin–destination matrix and network were taken as constructed in [25,26]. Road networks were extracted from OpenStreetMap (OSM) [35] in all cases, retaining road classes relevant to vehicular traffic. Link capacities $C_e$ were assigned according to road class; the complete capacity assignment is reported in Supplementary Note 1. Origin–destination matrices are defined between traffic-analysis zones. Further details on demand inference and network construction are given in[11,12,26].

**Traffic assignment.** Deterministic user-equilibrium flows were computed with an origin-based assignment algorithm[36], using the Bureau of Public Roads travel-time function $t_e(x_e) = t_e^0\left[1+\alpha\left(\frac{x_e}{c_e}\right)^{\beta}\right]$ with $\alpha = 0.6$ and $\beta = 4$ and a relative-gap convergence criterion below $0.1\%$. Assignment is performed over the complete road network; the link selection described below applies only to the cross-sectional analysis. Free-flow travel times $t_e^0$ carry the city-specific factor described in Supplementary Note 1; no further adjustment is applied to the resulting equilibrium travel times.

**Link selection.** Cross-sectional statistics were computed over links with $C_e > 1000\ veh/h$ and $v_e > v_{min}$, with $v_{min} = 0.01$. The resulting subnetworks contain $N_e$ links per city, with retained fractions from 19.2% in Beijing to 62.3% in Boston (Table 1). All cross-sectional quantities reported in the main text, including $T, \Gamma, \Gamma_w$, and $\chi_w$, are evaluated over this same subset. Sensitivity to the capacity threshold is reported in Supplementary Fig. S4 and Supplementary Note 5.

**Temperature estimation and exponential fits.** The effective temperature is the mean volume-to-capacity ratio over the retained links, $T = \langle v \rangle$, the maximum-likelihood scale of $P(v) = \frac{1}{T} exp\ (-v/T)$. Goodness of fit is quantified by the Kolmogorov–Smirnov distance, reported as a discrepancy measure rather than as a test. Because links with $v_e \leq v_{min}$ are excluded, the sample is left-truncated and the maximum-likelihood scale of the truncated exponential is $\langle v \rangle - v_{min}$; we report the ordinary mean so that the fitted scale coincides with the mean loading entering the macroscopic relations, which shifts T upward by exactly v_min in every city. The nominal standard error $\sigma_T = {T}/{\sqrt{N}}$ reflects sampling uncertainty under the exponential model only, and understates the true uncertainty because link loads are not independent. The hierarchy coupling is evaluated with $w_e = t_e^o C_e$ so that $\Gamma_w = T(1 + \chi_w)$ holds identically; the correlation $\rho_{wv}$ is reported in Supplementary Table S2 and Supplementary Note 3.

**Travel-time scaling.** Routes were retained if their equilibrium travel time was below 150 min. Mean equilibrium travel time was computed in 5 km bins of equilibrium route distance $d$, weighted by route flow, with the fit restricted to d ≤ 40 km, giving 8 bins per city and 48 observations. We fitted $t(d) \simeq \frac{d}{\bar{u}_f}\left[1 + \alpha^* T^{\beta^*}\right]$,where $\bar{u}_f$ is each city's characteristic free-flow speed, computed as route length divided by free-flow travel time and averaged across routes with route flow as weight, and $T$ its effective temperature. Both $\bar{u}_f$ and $T$ were fixed from the network data; the two remaining coefficients were estimated jointly across all six cities by weighted nonlinear least squares, giving a single global pair $(\alpha^*, \beta^*)$ with no city-specific parameters. Uncertainties and covariance come from the covariance matrix of the fit. Because the coefficients are strongly correlated ($\rho = 0.99$), the exponent was additionally assessed with a profile fit: $\beta^*$ was fixed over $[0.5, 6.0]$, $\alpha^*$ re-optimized at each value under the same objective, and the residual error recorded (Supplementary Note 6, Supplementary Fig. S5).

**Variation of the BPR coefficient.** Boston equilibria were recomputed for $\alpha = 0.15, 0.25, 0.40, 0.60$ and $0.80$, with the retained link set re-evaluated for each. The linearization giving $v_h = \beta/(\beta + 1)$, and the fact that $\Phi_{\text{eff}}$ is a lower bound on the exact cost, are derived in Supplementary Note 2.

**Demand scaling and high-loading fraction.** Baseline origin–destination matrices were uniformly scaled by $\lambda = 0.5, 0.6, 0.7, \ldots, 1.4$, preserving their relative spatial structure, and a new equilibrium was computed independently for each level with all other inputs at baseline. The retained link set was re-evaluated at each $\lambda$ using the same criteria; relative to baseline, its size decreases by $5–12\%$ at $\lambda = 0.5$ and increases by $2–5\%$ at $\lambda = 1.4$ (Supplementary Table S5). For Fig. 4b, link loads were transformed with Eq. (9) and normalized by $T(\lambda) = \langle \Phi_{\text{eff}}(v) \rangle_\lambda$, the quantity also shown in Fig. 4c. The high-loading fraction is $f_h(\lambda) = \frac{1}{N_\lambda}\sum_{e \in \varepsilon_\lambda} 1(v_e > v_h)$. In Fig. 4d all cities and demand levels are pooled; the decay constant $E$ of $exp(-E/T)$ is fitted to the pooled data, while Eq. (10) and its local power-law form are drawn with the predicted $v_h = 0.8$ and no refitting.

**Congestion index comparison.** City-level indices were taken from the TomTom Traffic Index[28] for 2017 and expressed as fractional increases in travel time relative to uncongested conditions. The index is averaged over the full year and over the provider's perimeter, whereas our temperatures come from morning-peak equilibria, so the comparison tests ordering rather than point-by-point agreement; agreement was quantified with Spearman's rank correlation and an exact permutation p-value. The published value for San Francisco covers the city rather than the Bay Area network analyzed here, so that point is shown with an open symbol in Fig. 3c but retained in the correlation. The curve in Fig. 3c follows from inverting Eq. (8): with $delay = \frac{t-t_o}{t_o} = \alpha^* T^{\beta^*}$ , evaluated with the coefficients of Fig. 3b and no refitting. The band is the $1\sigma$ uncertainty propagated through that transformation with the full covariance.

**Radial profiles.** Each retained link was assigned a radial distance $r$ from a fixed city-centre reference point (Supplementary Table S1) to its midpoint, normalized as $\tilde{r} = {}^{r}/_{r_{max}}$ within the retained network, and grouped into equally spaced bins in $\tilde{r}$. Within each bin we computed the unweighted mean volume-to-capacity ratio normalized by the citywide temperature, $\langle v \rangle(r)/T$. Error bars in Fig. 3d are the standard error of the mean across links in the bin; both the radial means and the citywide temperature are link-count weighted.

**Distributional comparisons and threshold sensitivity.** Exponential fits were compared against lognormal, gamma and Weibull alternatives by truncated maximum likelihood on the retained subnetwork (Supplementary Table S3). Lognormal is rejected decisively in all six cities. Gamma has the lowest AIC in five cities and Weibull in Boston, but at 9,000–36,000 links per network ΔAIC is a poor guide. We therefore use the fitted gamma shape and the log-likelihood gain per link, which equals the Kullback–Leibler divergence $D_{KL}(P_\gamma | P_{exp})$ and ranges from $0.0006$ to $0.04$ nats per link (Supplementary Note 4). The shape is close to one in Boston, the San Francisco Bay Area and Rio de Janeiro ($0.90-1.08$) and below one in Beijing, Lisbon and Porto ($0.53-0.71$). These low shapes reflect an excess of arterial links with almost no assigned flow, left by deterministic assignment on coarsely resolved demand; raising the loading cut to $v_{min} = 0.05$ brings Beijing's shape to 1.03. Sub-unity shapes also appear in several LODES-based equilibria (Supplementary Fig. S2), which likewise assign zone-level demand deterministically; across those 20 networks $\boldsymbol{D_{KL}}$ has a median of 0.007 nats per link and decreases with mean load (Spearman ρ = −0.52). Varying the lower cut-off between $v_{min} = 0.01$ and 0.05 shifts $T$ by 0.3% in Boston and up to 15% in Porto, without reaching a plateau (Supplementary Note 5, Supplementary Table S4).

**Independent 20-city dataset.** We used the published user-equilibrium link flows of a unified traffic-assignment dataset for 20 US metropolitan areas[27]. Its road networks are extracted independently from OpenStreetMap, its demand consists of 2019 LODES commuting flows aggregated to census tracts, and its BPR functions use city-calibrated coefficients ($\alpha \in [0.15, 0.6]$, $\beta \in [1.2, 3]$). Equilibria were computed by the dataset authors with *AequilibraE*; we did not re-assign them. Volume-to-capacity ratios were evaluated with the same link selection as for the six main networks. Exponential and gamma distributions were fitted by maximum likelihood to the left-truncated sample, as for the six main networks (Supplementary Note 4, Supplementary Fig. S2).

**Dynamic simulations.** Mesoscopic dynamic simulations of the morning peak were re-run for Rio de Janeiro and the San Francisco Bay Area from the published networks, demand and simulation configuration of ref. [8], since link-level output was not released. In these simulations vehicles reroute stochastically in response to evolving travel times and queues form and spill back, so no equilibrium problem is solved. For each link and each one-hour window between 07:00 and 11:00, the volume-to-capacity ratio is the mean outflow per lane over the window (veh $h^{-1}$) divided by the speed-based lane capacity of Supplementary Note 1. Link selection uses the same capacity and loading cuts as the static analysis. $T$ is taken directly as $\langle v \rangle$ in each window, with no fitted parameter (Supplementary Fig. S3).

Large language model tools were used to assist with code for data analysis and figure generation, consistency checks, and language editing. All analyses, results and conclusions were designed, verified and are the sole responsibility of the author.

**Data availability:** The link-level equilibrium flows, the processed data and all inputs needed to reproduce the figures will be deposited in a public GitHub repository, archived with a permanent DOI, upon publication.

**Code availability.** The traffic-assignment code and the analysis notebooks used to produce all figures and tables will be made available in the same public repository upon publication.

**Author contributions.** The author conceived the study, developed the theoretical framework, performed the traffic assignments and analyses, and wrote the manuscript.

**Competing interests:** The author declares no competing interests.

**References**


1. Herman, R. & Prigogine, I. A Two-Fluid Approach to Town Traffic. *Science* **204**, 148–151 (1979).
2. Herman, R., Lam, T. & Prigogine, I. Kinetic Theory of Vehicular Traffic: Comparison with Data. *Transp. Sci.* **6**, 440–452 (1972).
3. Helbing, D. Derivation of a fundamental diagram for urban traffic flow. *Eur. Phys. J. B* **70**, 229–241 (2009).
4. Geroliminis, N. & Daganzo, C. F. Existence of urban-scale macroscopic fundamental diagrams: Some experimental findings. *Transp. Res. Part B Methodol.* **42**, 759–770 (2008).
5. Daganzo, C. F. & Geroliminis, N. An analytical approximation for the macroscopic fundamental diagram of urban traffic. *Transp. Res. Part B Methodol.* **42**, 771–781 (2008).
6. Loder, A., Ambühl, L., Menendez, M. & Axhausen, K. W. Understanding traffic capacity of urban networks. *Sci. Rep.* **9**, 16283 (2019).
7. Ambühl, L., Loder, A., Leclercq, L. & Menendez, M. Disentangling the city traffic rhythms: A longitudinal analysis of MFD patterns over a year. *Transp. Res. Part C Emerg. Technol.* **126**, 103065 (2021).

8. Ambühl, L., Menendez, M. & González, M. C. Understanding congestion propagation by combining percolation theory with the macroscopic fundamental diagram. *Commun. Phys.* **6**, 26 (2023).

9. Zeng, G. *et al.* Switch between critical percolation modes in city traffic dynamics. *Proc. Natl. Acad. Sci.* **116**, 23–28 (2019).

10. Li, D. *et al.* Percolation transition in dynamical traffic network with evolving critical bottlenecks. *Proc. Natl. Acad. Sci.* **112**, 669–672 (2015).

11. Çolak, S., Lima, A. & González, M. C. Understanding congested travel in urban areas. *Nat. Commun.* **7**, 10793 (2016).

12. Olmos, L. E., Çolak, S., Shafiei, S., Saberi, M. & González, M. C. Macroscopic dynamics and the collapse of urban traffic. *Proc. Natl. Acad. Sci.* **115**, 12654–12661 (2018).

13. Nagel, K., Wagner, P. & Woesler, R. Still Flowing: Approaches to Traffic Flow and Traffic Jam Modeling. *Oper. Res.* **51**, 681–710 (2003).

14. Helbing, D. Traffic and related self-driven many-particle systems. *Rev. Mod. Phys.* **73**, 1067–1141 (2001).

15. Chowdhury, D., Santen, L. & Schadschneider, A. Statistical physics of vehicular traffic and some related systems. *Phys. Rep.* **329**, 199–329 (2000).

16. Traffic Flow as a Simple Fluid: Toward a Scaling Theory of Urban Congestion - Jorge A. Laval, 2024. https://journals.sagepub.com/doi/abs/10.1177/03611981231179703.

17. Wang, P., Hunter, T., Bayen, A. M., Schechtner, K. & González, M. C. Understanding Road Usage Patterns in Urban Areas. *Sci. Rep.* **2**, 1001 (2012).

18. Xu, Z., Sun, L., Wang, J. & Wang, P. The Loss of Efficiency Caused by Agents' Uncoordinated Routing in Transport Networks. *PLOS ONE* **9**, e111088 (2014).

19. Wardrop, J. G. SOME THEORETICAL ASPECTS OF ROAD TRAFFIC RESEARCH. in vol. 1 325–378 (1952).

20. Beckmann, M. J., McGuire, C. B., & Winsten, C. B. *Studies in the Economics of Transportation*. (Yale University Press, 1956).

21. Sheffi, Y. *Urban Transportation Networks: Equilibrium Analysis with Mathematical Programming Methods*. (Prentice-Hall, Englewood Cliffs, N.J, 1984).

22. Monderer, D. & Shapley, L. S. Potential Games. *Games Econ. Behav.* **14**, 124–143 (1996).

23. US Department of Commerce. *Traffic Assignment Manual. Bureau of Public Roads.* (1964).

24. Jaynes, E. T. Information Theory and Statistical Mechanics. *Phys. Rev.* **106**, 620–630 (1957).

25. Liu, C. *et al.* Revealing spatiotemporal interaction patterns behind complex cities. *Chaos Interdiscip. J. Nonlinear Sci.* **32**, 081105 (2022).

26. Xu, Y. *et al.* Unraveling environmental justice in ambient PM2.5 exposure in Beijing: A big data approach. *Comput. Environ. Urban Syst.* **75**, 12–21 (2019).

27. Xu, X., Zheng, Z., Hu, Z., Feng, K. & Ma, W. A unified dataset for the city-scale traffic assignment model in 20 U.S. cities. *Sci. Data* **11**, 325 (2024).

28. Traffic Index ranking. *TomTom Traffic Index* https://www.tomtom.com/traffic-index/ranking/.

29. Kirkpatrick, S., Gelatt, C. D. & Vecchi, M. P. Optimization by Simulated Annealing. *Science* **220**, 671–680 (1983).

30. Mézard, M. & Parisi, G. A replica analysis of the travelling salesman problem. *J. Phys.* **47**, 1285–1296 (1986).

31. Mézard, M. & Parisi, G. Replicas and optimization. *J. Phys. Lett.* **46**, 771–778 (1985).

32. Yeung, C. H., Saad, D. & Wong, K. Y. M. From the physics of interacting polymers to optimizing routes on the London Underground. *Proc. Natl. Acad. Sci.* **110**, 13717–13722 (2013).

33. Xiao, F., Yang, H. & Ye, H. Physics of day-to-day network flow dynamics. *Transp. Res. Part B Methodol.* **86**, 86–103 (2016).

34. Zhao, X., Wan, C. & Bi, J. Day-to-Day Assignment Models and Traffic Dynamics Under Information Provision. *Netw. Spat. Econ.* **19**, 473–502 (2019).

35. OpenStreetMap contributors. OpenStreetMap database. (2026).

36. Bar-Gera, H. Origin-Based Algorithm for the Traffic Assignment Problem. *Transp. Sci.* **36**, 398–417 (2002).

5

## Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- SupplementaryInformationCitieshaveaneffectivetemperature.pdf