Supplementary Information for the article "Cities have an effective temperature: A statistical mechanics description of user equilibrium"

Supplementary Note 1

All six road networks were extracted from OpenStreetMap following the parser of [1]. Road classes irrelevant to vehicular traffic are discarded, as are residential roads, which are peripheral to the congestion problem and expensive to carry. Chains of links with a single incoming and a single outgoing connection in the same class are collapsed into one segment. The resulting network is the complete road network E over which user equilibrium is solved (Eq. 1); all origin–destination demand is assigned on it, and the cross-sectional subset $\mathcal{E} \subset E$ is selected afterwards (Supplementary Note 5). Where OSM lacks lane counts or speed limits, both are imputed from the road class following the Highway Capacity Manual[2]: motorways are assigned 60 mph and three lanes per direction, primary roads 40 mph and two lanes, and so on down the hierarchy.

Capacity per lane follows the same source and depends on the imputed speed limit u(km/h):

$$C_{\text{lane}} = \begin{cases} 950, & u < 40, \\ 1500 + 30\,u, & 40 \le u < 60, \\ 1700 + 10\,u, & u \ge 60, \end{cases}$$

multiplied by the number of lanes. Free-flow travel time on a segment of length ℓ_e is $t_e^0 = f_p \cdot \ell_e/u_e$, where f_p is a city-specific factor accounting for intersection delay and other within-segment losses not represented in the free-flow speed: 1.4 in Boston, 1.3 in the San Francisco Bay Area, Rio de Janeiro, 1.15 in Lisbon, Porto, and Beijing. Link length is recovered from the stored travel time as $\ell_e = t_e^0 \cdot u_e/f_p$. No further adjustment is applied to the equilibrium travel times.

Given this assignment, the threshold $C_e > 1000\ veh/h$ used throughout removes exactly the single-lane roads below $40\ mph$, whose capacity is $950\ veh/h$; every multi-lane road, and every single-lane road at $40\ mph$ or above, is retained.

Origin–destination matrices are defined between traffic-analysis zones and describe the morning peak hour (07:00–08:00). For Boston, the San Francisco Bay Area, Lisbon, Porto and Rio de Janeiro they were inferred from mobile-phone call-detail records following [1]; the Beijing matrix and network were obtained from [3]. City-centre reference points used for the radial profiles of Fig. 3d are listed in Supplementary Table S1.

Supplementary Note 2

The effective loading cost of Eq. (3) is

$$\Phi(v) = v + \frac{\alpha}{\beta+1} \cdot v^{\beta+1},$$

whose nonlinear part $q(v) = \frac{\alpha}{\beta+1} \cdot v^{\beta+1}$ is negligible at low loading and grows rapidly near capacity. Linearizing q about nominal capacity, $v = 1$, gives

$$q(v) \approx q(1) + q'(1) \cdot (v-1) = \frac{\alpha}{\beta+1} + \alpha\,(v-1) = \alpha\left(v - \frac{\beta}{\beta+1}\right)$$

which is positive only for $v > v_h$, with $v_h = \frac{\beta}{\beta+1}$. Retaining the tangent where it is positive yields the hinged cost of Eq. (9), $\Phi_{\text{eff}}(v) = v + \alpha\, max\,(0, v - v_h)$.

Two properties follow. First, $\Phi_{\text{eff}}(v_h) = v_h$, so the sets $\{v_e > v_h\}$ and $\{\Phi_{\text{eff}}(v_e) > v_h\}$ coincide and the high-loading fraction f_h is the same whether measured in v or in Φ_{eff}. Second, q is convex, so its tangent at $v = 1$ lies below it everywhere: $\Phi_{\text{eff}}(v) \leq \Phi(v)$for all v, with equality at $v = 1$. The hinged cost is therefore a lower bound on the exact cost, and the bound loosens above nominal capacity — the direction in which the tail of Fig. 4a departs from the fitted exponential.

The breakpoint depends only on β: for $\beta = 4$, $v_h = 0.8$ whatever the value of α. Below v_h the hinged cost neglects the exact correction, which at v_h amounts to $\frac{\alpha}{\beta+1} \cdot {v_h}^{\beta+1}$, of the linear term: 1.2 % for α = 0.15 and 4.9 % for α = 0.6. Above v_h the hinged correction $\alpha(v - v_h)$ grows to 3 % ($\alpha = 0.15$) and 12 % ($\alpha = 0.6$) of the linear term at nominal capacity, where it coincides with the exact one. Over $v \leq 1.3$ the hinged and exact costs differ by at most 2.6 % ($\alpha = 0.15$) and 8.3 % ($\alpha = 0.6$).

Supplementary Note 3

Let $\mathcal{E}$ be the retained link set, $w_e = t_e^0 C_e$, and $\langle\cdot\rangle$ the unweighted mean over $\mathcal{E}$. The weighted mean loading of Eq. (4) is

$$\Gamma_w = \frac{\Sigma_{e\in\mathcal{E}}\, t_e^0 C_e\, v_e}{\Sigma_{e\in\mathcal{E}}\, t_e^0 C_e} = \frac{\langle wv\rangle}{\langle w\rangle}$$

Writing the numerator through the covariance, $\langle wv\rangle = \langle w\rangle\langle v\rangle + \mathrm{Cov}(w, v)$, and dividing by $\langle w\rangle$ gives

$\Gamma_w = \langle v\rangle\left[1 + \frac{\mathrm{Cov}(w,v)}{\langle w\rangle\langle v\rangle}\right] = T\cdot(1 + \chi_w)$, with $\chi_w \equiv \frac{\mathrm{Cov}_\varepsilon(w,v)}{\langle w\rangle_\varepsilon\langle v\rangle_\varepsilon}$ and $T = \langle v\rangle$. This is an identity: no assumption is made about the distribution of either v or w, and it holds for any retained subset.

Because χ_w is a normalized covariance, it depends on the dispersion of w as well as on the strength of the association between w and v: $\chi_w = \rho_{wv}\frac{\sigma_w}{\langle w\rangle}\frac{\sigma_v}{\langle v\rangle}$. The same value of χ_w can therefore arise from a strong correlation in a homogeneous hierarchy or from a weak one in a heterogeneous hierarchy. Both ρ_{wv} and the two coefficients of variation are reported per city in Supplementary Table S2; the San Francisco Bay Area combines the lowest $\frac{\sigma_w}{\langle w\rangle}$ with a positive correlation, which is why its χ_w is the largest of the six.

Supplementary Note 4

Among all distributions on $v_e > v_{min}$ with prescribed mean load $\langle v\rangle$, the maximum-entropy one is the exponential, $P(\mathrm{v}) \propto e^{-v/T}$. Prescribing also $\langle ln\, v\rangle$ yields the gamma, $P(v) \propto v^{\kappa-1}e^{-v/\theta}$, with $\kappa - 1$ the Lagrange multiplier conjugate to $\langle ln\, v\rangle$; the exponential is its $\kappa = 1$ case. At the maximum-likelihood solution $\kappa\theta = \langle v\rangle$, so the temperature T is the same under both models and the gamma differs only in shape near v_{min}. Unlike $\langle v\rangle$, which is conjugate to T and enters the Beckmann potential, $\langle ln\, v\rangle$ has no macroscopic counterpart; since $\frac{d\ln v}{dv} = \frac{1}{v}$, it is dominated by the least-loaded links.

The exponential was compared against lognormal, gamma and Weibull on the retained subnetwork of each city (Table S3). Gamma and Weibull contain the exponential as their shape-equal-to-one case, so the fitted shape is itself a measure of how far the data depart from it. Lognormal is rejected everywhere. Among the rest the comparison is close, and ΔAIC is a poor guide at N_e between 9,000 and 36,000 links: in the Bay

Area and Rio the gamma wins on likelihood while fitting slightly worse in KS distance than the one-parameter exponential. The informative quantity is the fitted gamma shape, since the exponential is its $\kappa = 1$ case. It lies between 0.90 and 1.08 in Boston, the Bay Area and Rio. In Beijing it rises from 0.71 to 1.03 when the loading cut is raised to $v_{min} = 0.05$ (Supplementary Note 5), the expected response when the links that dominate $\langle ln\ v \rangle$ are removed. Porto (0.53) and Lisbon (0.57) remain below one, consistent with an excess of nearly empty parallel links left by deterministic route choice on coarser zone systems.

The information carried by the second constraint can be quantified directly. For maximum-likelihood fits within exponential families, the log-likelihood gain per link equals the Kullback–Leibler divergence between the fitted distributions, $\Delta\ell/N_e = D_{KL}(P_\gamma | P_{exp})$. In the six cities it ranges from 0.0006 (Bay Area) to 0.04 (Porto) nats per link (Table S3). Across the 20 cities of Supplementary Fig. S2 it has a median of 0.007, never exceeds 0.06, and decreases with the mean load (Spearman $\rho = -0.52$). The exponential thus fixes the scale of the distribution in every city, while the gamma shape adds a small correction near v_{min} that weakens as networks become more loaded.

Supplementary Note 5

Two thresholds define the retained subnetwork: a capacity cut $C_e > 1000\,\frac{veh}{h}$ and a loading cut $v_e \geq v_{min}$, with $v_{min} = 0.01$ at baseline.

Table S4 reports T as v_{min} is raised from 0.01 to 0.05. The drift is +14.8% in Porto and +11.9% in Lisbon, but only +3.0% in the San Francisco Bay Area and −0.3% in Boston. It tracks the resolution of the origin–destination data rather than the congestion level: raising v_{min} removes links carrying almost no assigned flow, which are most abundant where the zoning is coarsest. Extending the sweep to $v_{min} = 0.2$ separates two behaviours: in Boston and Rio de Janeiro T passes through a maximum and declines thereafter, the signature of those links being exhausted before the threshold begins to erode the distribution itself, whereas in the other four networks T is still rising at the end of the range. The ordering of cities by T is unchanged throughout, but differences of order 15% or less should not be over-interpreted.

Because capacity is assigned per lane by road class, the baseline cut at $1000\,\frac{veh}{h}$ falls in a gap: single-lane roads below 40 mph carry 950 veh/h and everything else at least 1900 (Supplementary Note 1). Beijing is the most sensitive, consistent with its retention of 19.3%, and should be read as a more restrictive sampling of the hierarchy than the other five.

Supplementary Note 6

The parameters α^* and β^* in Equation (8) are no independent, thus they correlate. We swept β^* across [1.0, 4.0] in steps of 0.05, re-fitting α^* at each value. Since the model is linear in α* once β^* is held fixed, each point of the sweep has an exact solution. The result is a very shallow valley: with a minimum between β^*= 1.27 and 2.73, so any exponent in that range describes the six cities about as well as the best one (Supplementary Fig. S5a).

Along the valley the two coefficients compensate. Plotting ln α^* against the fixed β^*gives a straight line, which means α^* falls off geometrically as β^*rises:

$$\alpha^*(\beta^*) = A \cdot {T_{piv}}^{-\beta^*},\ T_{piv} = 0.34$$

The reason is simple. All six temperatures sit in a narrow band, so the fit only ever evaluates $\alpha^* T^{\beta^*}$ near T_{piv}, and many pairs give the same answer there. The free fit says the same thing: its two coefficients come out correlated at 0.993.

So the quantity the data pin down is the fractional delay at the temperatures actually observed, not α^* or β^* on their own. We use this twice in the main text. The global pair is reported as evidence that one relation covers six cities, not as a measurement of its coefficients. And β^* is read as the local sensitivity of f_h to T, because over a narrow range a power law is what an exponential in $1/T$ looks like.

Supplementary Figures

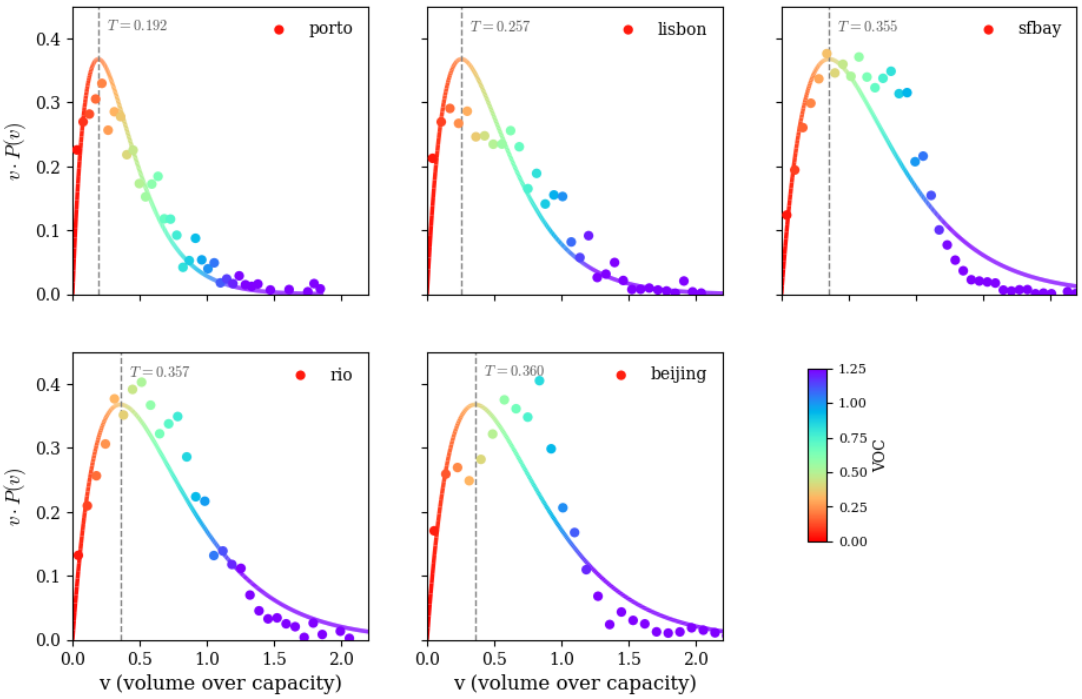
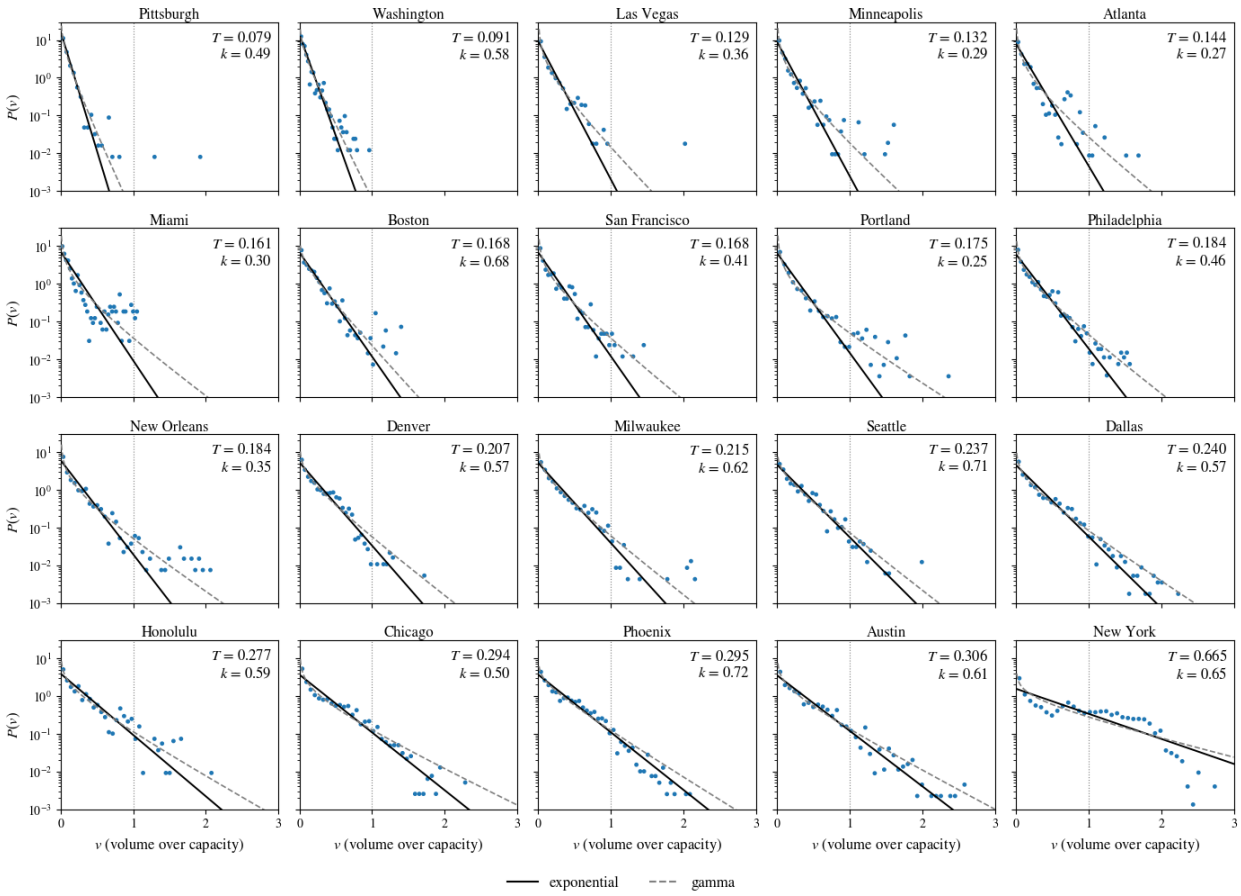


Figure S1 | Contribution of each loading level to the mean. The quantity $v \cdot P(v)$ for the five cities not shown in Fig. 2g. For an exponential distribution it peaks at $v = T$, so the effective temperature is both the mean cross-sectional loading and the loading level that contributes most to it. The peak sets the prominence weighting used in the maps of Fig. 2.



Supplementary Fig. S2 | VOC distributions in an independent 20-city dataset. Distribution $P(v)$ of volume-over-capacity on directed links ($v > v_{min}, 1000 < C < 10{,}000\ veh\ h-1$) for 20 US metropolitan areas from the unified traffic-assignment dataset of [4]. Networks (OpenStreetMap), demand (2019 LODES commuting flows aggregated to census tracts) and cost functions (BPR with city-calibrated α in [0.15,0.6], β in [1.2,3]) are all independent of the six cities in the main text. Link flows are the published user-equilibrium solutions computed with AequilibraE. Points: histogram with 40 linear bins, normalised over all retained links. Solid line: maximum-entropy exponential, constrained by $< v >$. Dashed line: maximum-entropy gamma, additionally constrained by $< ln\ v >$. Both are fitted by maximum likelihood to the sample left-truncated at $v_{min} = 0.01$. Each panel reports $T = \langle v \rangle$ and the gamma shape k; panels are ordered by increasing T, and the dotted line marks v = 1. In every city $P(v)$ decays exponentially with scale T; the gamma correction is confined to small v and decreases with increasing T (Supplementary Note 4).

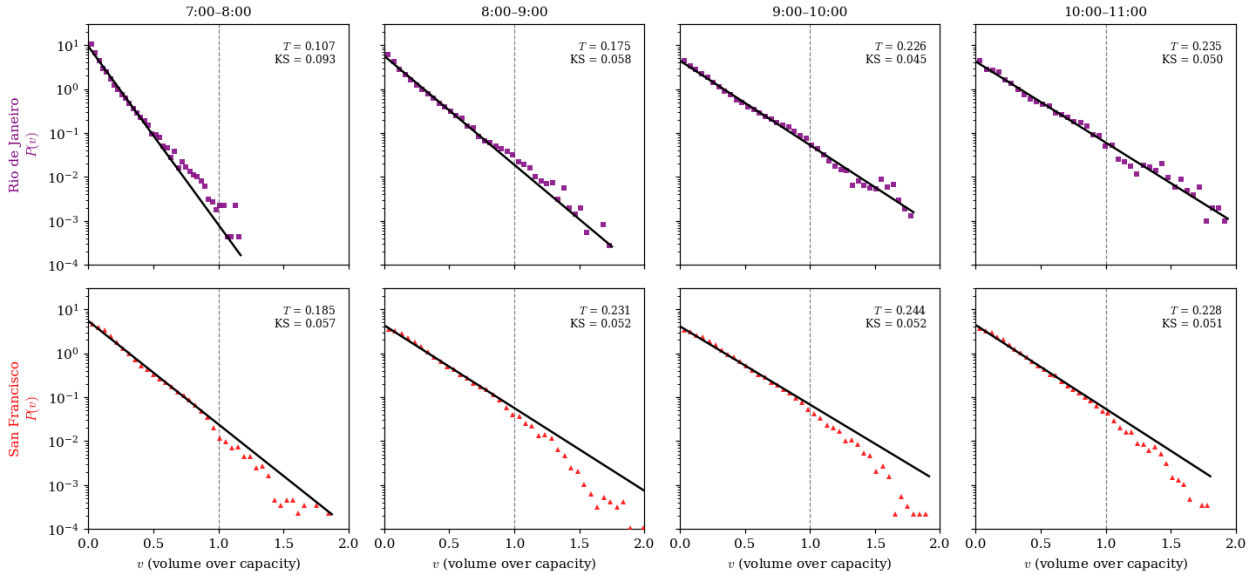


Figure S3 | Exponential loading without an equilibrium problem. Cross-sectional distributions of the volume-to-capacity ratio v in mesoscopic dynamic simulations of the morning peak, re-run from the published networks, demand and configuration of [5], for Rio de Janeiro and the San Francisco Bay Area. Columns are consecutive one-hour windows from 07:00 to 11:00; vertical axes are logarithmic. For each link and window, v is the mean outflow per lane over the window, expressed in vehicles per hour and divided by the same speed-based lane capacity used in the static analysis, with the same capacity filters (Supplementary Note 5). Black lines are exponentials $P(\mathrm{v}) = T^{-1}e^{-v/T}$ with $\mathrm{T} = \langle v \rangle$ taken directly from each panel, so no parameter is fitted; dashed lines mark $v = 1$. Temperatures range from 0.107 to 0.244 and KS distances remain below 0.093.

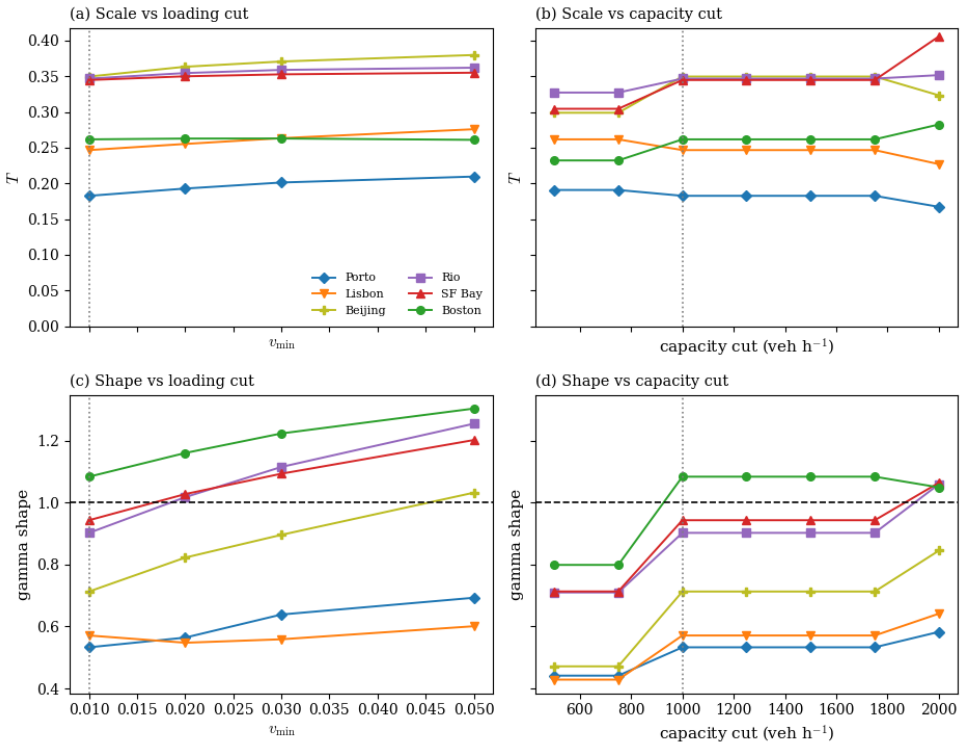


Figure S4 | Sensitivity to the selection thresholds. Estimated scale T (a, b) and fitted gamma shape (c, d) against the loading cut v_{min}, over 0.01–0.05, and against the capacity cut, over 500–2000 veh h^{-1}, for the six cities. All fits are truncated maximum likelihood. Dotted vertical lines mark the values used throughout; the dashed horizontal line in (c, d) marks shape = 1, at which the gamma reduces to the exponential. The scale is robust. T changes by less than 15 % across the v_{min} sweep (a), and responds to the capacity cut only when that cut crosses a road class (b). Capacity is assigned per lane from a step function of the free-flow speed, so link capacities take only a few distinct values — 950 veh h^{-1} per lane below 40 mph, and higher values for the faster classes. No link is therefore removed between 1000 and 1750 veh h^{-1} and T is flat over that interval, whereas the step at 2000 crosses the two-lane class at 1900 and discards 57 % of the San Francisco Bay Area network. The sign of that step differs across cities, since it depends on how heavily the removed class is loaded relative to the rest. The reference value of 1000 veh h^{-1} lies immediately above the single-lane class and was not tuned. The shape is not robust, and the way it moves is itself informative. It increases as the loading cut deepens (c) in every city except Lisbon, which is nearly flat (0.55–0.60). It also increases at the capacity step from 750 to 1000 veh h^{-1} (d), when the single-lane class is removed. Boston lies above 1 at every loading cut, Rio de Janeiro and the San Francisco Bay Area cross 1 at $v_{min} \approx 0.02$, and Beijing reaches 1 only at $v_{min} = 0.05$. These sub-unity shapes at the reference thresholds therefore reflect the population of near-empty links rather than a property of the loaded network. Porto and Lisbon remain below 0.7 throughout, so their departure from the exponential reference is not a threshold effect. In the five cities whose shape is below 1 at the reference thresholds, $v_{min} = 0.01$lies at or near the largest departure from 1, so the comparisons in the main text are made in the least favourable part of this sweep. Only T enters the analysis, and it is stable across both sweeps.

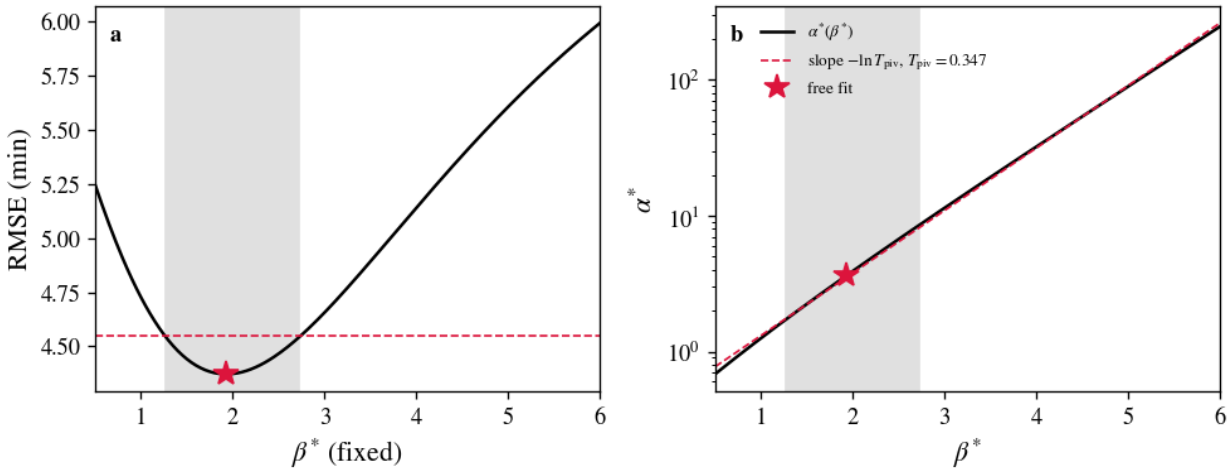


Figure S5 | Identifiability of the travel-time parameters. (a) Root-mean-square error of the fit of Eq.8 to the $N = 48$ flow-weighted travel-time observations (six cities, eight distance bins each) with β^* fixed and α^* re-optimized at each value. The dashed line marks the approximate 95% criterion, $\Delta SSE \leq \Delta SSE_{min}\chi^2_{1,0.95}/(N-2)$, and the shaded band the compatible range β^* in [1.27, 2.73], across which the error rises by at most 4.1%. (b) The corresponding compensation ridge, $\alpha^*(\beta^*)$ on logarithmic axes; the dashed line is the linear fit $ln\,\alpha^* = \ln A - \beta^*\,ln\,T_{piv}$, with $T_{piv} = 0.347$ ($R^2 = 0.999$). The star marks the free fit $\alpha^* = 3.5$, $\beta^* = 1.92$; correlation of the estimates 0.99). Along the ridge the data constrain the combination $\alpha^* T^{\beta^*}$ near T_{piv}, which varies by less than 6.2% within the shaded band, rather than either coefficient separately.

Supplementary Tables

Table S1 | Network properties and city-centre reference points.

City	Links (total)	Retained N_e	Retained (%)	Centre (lat, lon)
Porto	18 895	9 223	48.8	(41.1496, -8.6110)
Lisbon	32 756	15 661	47.8	(38.7169, -9.1395)
Boston	58 212	36 271	62.3	(42.3601, -71.0589)
SF Bay	73 929	33 200	44.9	(37.7749, -122.4194)
Rio	27 661	13 181	47.7	(-22.9068, -43.1729)
Beijing	81 524	15 641	19.3	(39.918546, 116.390913)

Table S2 | Hierarchy coupling and its decomposition. χ_w is the normalized covariance of Eq. (5); ρ_{wv} is the Pearson correlation between link weight and relative loading, and the two coefficients of variation complete the decomposition $\chi_w = \rho_{wv}\,(\sigma_w/\langle w\rangle)(\sigma_v/\langle v\rangle)$.

City	T	Γ	Γ_w	χ_w	ρ_{wv}	$\sigma_w/\langle w\rangle$	$\sigma_v/\langle v\rangle$
Porto	0.193	0.171	0.163	−0.15	−0.06	2.12	1.12
Lisbon	0.257	0.205	0.194	−0.24	−0.08	2.60	1.11
Boston	0.272	0.260	0.241	−0.11	−0.08	1.58	0.87
SF Bay	0.355	0.429	0.382	+0.08	+0.05	1.58	0.88
Rio	0.359	0.329	0.322	−0.10	−0.06	1.88	0.91
Beijing	0.360	0.369	0.375	−0.15	−0.09	1.63	0.98

Table S3 | Comparison of distributional families. ΔAIC is reported relative to the best-fitting family in each city (marked *); lower is better. All four families were fitted to the same retained links by truncated maximum likelihood with location fixed at zero. Gamma and Weibull contain the exponential as their shape-equal-to-one case. KS is the Kolmogorov–Smirnov distance of the exponential fit. See Supplementary Note 4.

City	exponential	gamma	Weibull	lognormal	gamma shape k	KS (exp.)	$AIC_{exp} - AIC_{\gamma}$	D_{KL} (nats/link)
Boston	172	97	0*	4603	1.08	0.027	75	0.0011
SF Bay Area	38	0*	39	5292	0.94	0.033	38	0.0006
Rio	46	0*	42	2215	0.90	0.035	46	0.0018
Beijing	569	0*	209	2460	0.71	0.067	569	0.018
Lisbon	1204	0*	114	1124	0.57	0.100	1204	0.038
Porto	732	0*	72	657	0.53	0.092	732	0.040

Table S4 | Sensitivity of the exponential scale to the loading threshold. $T' =$ is the maximum-likelihood scale of the exponential truncated at v_{min}, $T' = \langle v \rangle - v_{min}$, with the capacity threshold held at its baseline value. For an exact exponential $T' =$ is independent of v_{min}, so its drift measures the departure from the exponential near zero loading. At the baseline v_{min} = 0.01, the effective temperature of the main text is $T = T' + 0.01$. Drift is the relative change in T' between $v_{min} = 0.01$ and 0.05.

City	$T'(v_{min}= 0.01)$	$T'(0.02)$	$T'(0.03)$	$T'(0.05)$	drift (%)
Porto	0.182	0.193	0.201	0.210	+14.8
Lisbon	0.247	0.255	0.263	0.276	+11.9
Boston	0.262	0.263	0.263	0.261	−0.3
SF Bay	0.345	0.350	0.353	0.355	+3.0
Rio	0.347	0.355	0.359	0.362	+4.4
Beijing	0.350	0.363	0.371	0.380	+8.6

Supplementary Table S5 | Size of the retained link set across demand levels. Number of retained links Ne ($Ce > 1000\ veh\ h^{-1}, v_e \geq v_{min}$) for each demand multiplier λ. The shaded column is the baseline demand (λ = 1). Δ(0.5) and Δ(1.4) give the relative change in N_e at the lowest (λ = 0.5) and highest (λ = 1.4) demand with respect to the baseline.

City	λ = 0.5	λ = 0.6	λ = 0.7	λ = 0.8	λ = 0.9	λ = 1.0	λ = 1.1	λ = 1.2	λ = 1.3	λ = 1.4	Δ(0.5) (%)	Δ(1.4) (%)
Porto	8,114	8,459	8,719	8,934	9,088	9,223	9,377	9,534	9,630	9,684	−12.0	+5.0
Lisbon	14,203	14,620	14,956	15,214	15,441	15,661	15,850	15,999	16,147	16,246	−9.3	+3.7
Boston	34,395	34,998	35,410	35,740	36,009	36,271	36,496	36,680	36,851	37,025	−5.2	+2.1
SF Bay Area	31,302	31,843	32,324	32,677	32,975	33,200	33,464	33,742	33,940	34,127	−5.7	+2.8
Rio de Janeiro	12,397	12,597	12,775	12,919	13,080	13,181	13,283	13,370	13,452	13,536	−5.9	+2.7
Beijing	14,350	14,716	15,010	15,275	15,466	15,641	15,784	15,921	16,038	16,145	−8.3	+3.2

References

1. Çolak, S., Lima, A. & González, M. C. Understanding congested travel in urban areas. *Nat. Commun.* **7**, 10793 (2016).
2. Reilly, W. HIGHWAY CAPACITY MANUAL 2000. *TR News* https://trid.trb.org/View/475202 (1997).
3. 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).
4. 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).
5. Ambühl, L., Menendez, M. & González, M. C. Understanding congestion propagation by combining percolation theory with the macroscopic fundamental diagram. *Commun. Phys.* **6**, 26 (2023).