Supplementary Information for

Provisional COVID-19 infrastructure induces large, rapid increases in cycling

Sebastian Kraus and Nicolas Koch

Kraus.
E-mail: kraus@mcc-berlin.net

This PDF file includes:

Supplementary text
Figs. S1 to S6
Tables S1 to S5
SI References

Supporting Information Text

Data cleaning

Bicycle count data. We assemble a new data set of daily bicycle counts from municipal bicycle counters. We connect to national and municipal open data portals for bike counter data sets. We connect directly to the API of those cities that use the Eco-Counter standard (see the data and code repository at https://zenodo.org/record/4015974). We also obtain longer time series of bike counts going back to 2012 directly from the mayor's staff for road planning and data in Paris.

Our raw data set contains roughly a million daily counts starting in 2007. We drop the lower and upper percentiles from this raw sample since counters can record very low values, when they are not functioning properly or very high values, when there is a cycling event that drives up counts. We drop the counter 100041252 from Bergen that varies between very low values and some of the highest daily counts in the sample. Our results are robust to keeping these extreme values in the sample. The bulk of the bike counts are from most recent years and we focus most of the comparisons made in our regressions on the years 2019 and 2020. For certain treated cities, such as Paris and Berlin, the raw data already indicates an increase in the annual peak in June 2020 compared to June 2019. However, many of the control cities show a similar pattern (see Figure S4). In our regression analyses we find a robust effect of new infrastructures, both when taking the difference in these differences between treatment and control cities, but also when focusing on variation in treatment timing exclusively (i.e. when discarding control city information) (see Figure 3 in the paper).

Except for robustness checks at the city level, the unit of observation in our regression analyses is the bike counter and counts vary daily. An average counter detects 1457 cyclists per day. The mean number of counters active in the same city on a given day for the counters in the sample is 22.9 (mean over all counter-day observations). The mean size of cities in our sample is 33000 ha (see Figure S1).

Pop-up infrastructure data. We use project-level data on provisional infrastructure in European cities as a reaction to the COVID-19 pandemic collected by the European Cyclists' Federation (1). In the data we typically see the street, where the project is implemented, its size measured in kilometers, the date of announcement, and the date of implementation. The data also contains the type of project. 80% are categorized as bike lanes and 16% as traffic calming. Our data includes all projects recorded as of July 8, 2020. Our sample does not include infrastructure built after that date and excludes some bigger cities, for which adequate open bike counter data is missing.

We aggregate this data at the city-day level to construct a variable of daily implemented kilometers of pop-up bike lane. We use the city definition and corresponding polygons from the European Urban Audit 2020 (2). Typically areas defined by the European Urban Audit include suburbs. For instance, the Paris polygon includes many areas beyond the ring highway that surrounds the municipality of Paris ("Ville de Paris"). This allows us to capture commuting enabled by new bike lanes from the suburbs into the city center, which constitute an important share of all projects (see for instance https://carto.parlons-velo.fr/#10.13/48.8312/2.5506 for a detailed map of projects in France).

Our estimation sample contains 22 treated cities and 84 control cities, both of which some are dropped from our Poisson regressions depending on the specification because of a lack of variation after removing fixed effects or because we do not have observations for our control variables (the transit (3) and overall mobility (4) controls). Dublin and Berlin have been the earliest adopters of pop-up bike lanes in the sample and Paris has been the city with the largest program (see Figure 1 in the paper). We use variation in both timing and the extent of the implemented infrastructure to estimate our effects. We have a large sample from both France and Germany. This allows us to estimate our effect based on within-country variation removing time-varying factors related to the pandemic that could create bias in our estimates. While important cities such as London, Milan, Lisbon and Rome had either announced or already implemented a pop-up bike lane program at the time of the analysis, they are missing from the sample due to insufficient spatial or temporal coverage of the bike count data. The average length by city of all bike infrastructures in our sample combined is 11.5 km, the length of bike lanes is 8.2 and the number of measures implemented 19.8 (see Figure S2).

We check the sensitivity of our results to different specifications of the treatment, for instance as an indicator variable that is 1, if there is any cycling related infrastructure change in a city and 0 otherwise (see Figure S3).

Measurement error. The unit of observation in our preferred specifications (see Equation 1 in the Materials and Methods section) is the counter. Our estimates give the average effect over all counters in all cities in the sample. We assign the treatment to counters at the city-level, but our research design ensures that conditional on fixed effects and control variables treatment is as good as random. Therefore, measurement error in our natural experiment can be analyzed similarly to the stylized case of a cluster randomized controlled trial (RCT), for instance with a treatment that is randomized and assigned at the village or class-room level, but outcomes are measured for individuals.

We use different treatment definitions to investigate how different ways of conceptualizing and therefore mismeasuring the treatment influences our estimates. We could think that only the fact that a city rolls out *a* pop-up bike lane program could already create more cycling. We, therefore, start our investigation by looking at a binary treatment definition, separating the sample into treated and control cities for a standard difference-in-differences analysis. We could, however, also think that the number of kilometers built will make an important difference for the media echo that an announcement of such a policy gets. It is further likely that a kilometer of pop-up bike lane will have a larger impact in a small city than in a larger city. Thus, we also estimate dose-response relationships (effect for each kilometer built) in a generalized difference-in-differences setup and look at effects when the dose is expressed in absolute kilometers and in per capita and per km^2 terms. This helps correct for the fact that in larger or more populous cities, just like in the case of an individual in a cluster RCT in a larger village or school class, we tend to overestimate the dose received at each counter.

All these different cases are presented in Figure 3 in the paper. We show that our estimates get attenuated by measurement error when the dose is expressed in absolute kilometers and that treatment effects are higher, when the dose is expressed relative to population size and relative to the area of a city. Note that our fixed effect already removes any variation between cities in terms of population and area, ensuring that there is no omitted variable bias in the estimates.

Within cities there will still be counters that are farther away from pop-up bike lanes than others and some will only measure the expansion of cycling from pop-up bike lanes partly or not at all. We argue that our fixed effects remove any variation between counters and cities that could lead to systematic relationships between this measurement error and our treatment. With the counter fixed effect we control for factors related to city size, road network, and topography that could be both determinants of counter placement and treatment. The counter fixed effect also controls for local institutions and political majorities that could be driving both where there are counters and where treatment happens. In our remaining variation, we expect some counters to get a treatment measured with an upward bias and some with a downward bias without there being a systematic tendency, i.e. we are left with classical measurement error.

The counter fixed effect also removes systematic measurement differences between counters, for instance when a counter is placed near a route used for recreational cycling. Further, it ensures that only variation from counters that are present both before and after treatment will be used for the estimation of our treatment effect.

A remaining concern could be that cyclists change their routes and that this could potentially even imply that there is a problem with double-counting. However, we do not expect route changes to create a skewed measurement of cycling traffic. It is likely that cyclists will change their routes in reaction to pop-up bike lanes and this route change means they can come past a counter that they did not pass on their old routes. However, it can also be the case that they switch their route choice away from a counter. On balance, we do not expect this to create systematic measurement error. Further, because our unit of analysis is the counter and we look at average changes in these counts rather than the sum of cyclist counts in a city, there cannot be any double-counting. We also show robustness checks, for which we take the mean of all counters in a city to obtain a measure for city-level cycling traffic and run a regression at the city level (see Figure 3 in the paper).

Empirical strategy (continued from main body)

This section provides additional elements regarding our empirical strategy.

Our preferred specifications are presented in Equation 1 and Figure 2 (marked in blue) in the main body. In Table S5 we report on these specifications in more detail, varying the treatment definition (km, km per capita, and km per km^2) and including or dropping the public transit control.

In the following section we discuss our choice of using the outcome (cycling counts) in the logarithmic form. In the next section we explain the additional empirical specification on which Figure 2 in the main body is based.

Functional form. We use the natural logarithm rather than the level of the count of cyclists as the outcome because we expect cycling to grow in a multiplicative way between cities but also between counters. The country-day FE ensures that we focus on variation between cities in the same country, which are typically at similar stages in their "market penetration" of cycling and where thus cycling can be assumed to grow at similar "natural" rates in the absence of treatment. We expect pop-up bike lanes to have a multiplicative effect across counters in a city because they typically remove prominent bottlenecks from the network, leading to improved routes and increased (perceived) safety. The cycling counters that measure our outcome are placed next to roads and bike lanes that have different roles in the overall network. Central counters will pick up larger *absolute increases* than more peripheral ones. However, the *growth* rates they measure (approximated by the Δ of the natural logarithms) will be more similar.

We also show that our results are robust in a non-parametric setting using the *Matrix Completion* method for panel data (5), a machine-learning method to construct a counterfactual. The Matrix Completion approach frames the problem of causal inference as a missing data problem: For treated units we observe the potential outcome Y^1, i.e. cycling given the introduction of pop-up infrastructure, but we do not observe the potential outcome Y^0 representing cycling in the treated city had the treatment not happened. If we did, the difference would be the treatment effect. We treat our panel dataset as a matrix with missing values, which are the missing potential outcomes Y_i^0. The Matrix Completion method imputes these missing values via regularization based prediction (5).

Let Z^0 be the estimated missing elements of the Y^0 matrix. Analytically, the objective of the Matrix Completion approach is to "optimally predict the missing elements by minimizing a convex function of the difference between the observed matrix of Y^0 and the unknown complete matrix Z^0 by using nuclear norm regularization" (6):

$$\widehat{Z^0} = \arg\min_{Z^0} \sum_{(i,j)\in\Omega} \frac{\left(Y_{it}^0 - Z_{it}^0\right)^2}{|\Omega|} + \Lambda \|Z^0\ | \quad [1]$$

where $||Z^0||$ is the "nuclear norm (sum of singular values of Z^0)" (6) and Ω denotes the rows i and columns j of the non-missing entries. The regularization parameter Λ is chosen using 10-fold cross-validation.*

* For a longer explanation of the method in an applied context see (6), who we follow closely here.

The coefficients shown in Figure S5 confirm that the treatment effect of bike lanes builds up fast enough for specifications with city-week fixed effects to capture it (Equation 1, main body).

Additional empirical specification (Figure 2 in main body). Estimates shown in Figure 2 (main body) are based on the following model:

$$\ln \text{Count}_{id} = \Sigma_{\tau}\delta_{\tau}\left(\text{Treated}_{c} \times D_{m}^{\tau}\right) + \mathbf{X}_{cd} + \mu_{c} + \varphi_{nd} + \varepsilon_{id} \qquad [2]$$

where i indexes a counter, c a city, n a country, d a day, and m a month.

The data varies at the counter-day. μ_c is a city fixed effect and φ_{nd} is a country-day fixed effect that captures any daily changes common to all cities in a country.

The coefficients of interest plotted in Figure 2 are the δ_{τ}. They capture the effect of the pop-up bike lane treatment on bicycle counts over time. For this purpose our treatment variable $Treated$ is defined as a binary indicator for treatment that is 1 for treated cities and 0 for control cities. The D_{m}^{τ} are binary indicators that are 1 if month m is τ months before or after March 2020, when the pre-treatment period ends. In these specifications, the reference month and begin of the sample is February 2019, when τ equals -13.

Figure 2 in the main body and Figure S6 present the transformed estimate: $100 \times (\exp \delta_{\tau} - 1)$.

We also run this specification including weather controls $\mathbf{X}_{cd}$ to investigate, if seasonality could be driving these results (Figure S6). The results look virtually the same.

	Mean	Std. Dev.	25%	50%	75%	95%	Min.	Max.
Daily number of cyclists	1457.2	1895.7	255	744	1923	5151	1	13339
City size (ha)	32893.8	42393.6	14163.3	22018.5	40659.9	89180.2	455.6	251517
Year	2017	2	2016	2018	2019	2020	2007	2020
Number of active counters in the same city	22.9	23.2	4	14	32	82	1	90
Facebook mobility index	-0.16	0.21	-0.27	-0.11	-0.0044	0.088	-0.81	0.51
Observations	995818							

Table S1. Summary statistics at the counter-day level. The unit of observation of our analysis is the counter and data varies daily. Count data is from municipal bike counters and is obtained from different municipal APIs. Treatment and control variables are assigned to counters based on their city attribute. City definitions are from the EU Urban Audit (2). The Facebook mobility index is only available from March 2020. It measures aggregate movement activity by Facebook users in a given administrative area (districts or states).

	Mean	Std. Dev.	25%	50%	75%	95%	Max.
Total length of bike infrastructures	11.5	20.0	1.39	2.57	16.6	57.9	85.1
Total length of bike lanes	8.24	18.3	0.24	2.05	7.35	24.8	84.3
Number of measures	19.8	48.1	1	4	17	52	226
Observations	22						

Table S2. Summary statistics of most recent state of infrastructure at the city level. We use data from the European Cyclists' Federation (1). The raw data includes information on individual infrastructure projects announced or implemented. We aggregate it to the city-day level using city definitions from the EU Urban Audit. Our analysis includes data up to July 8, 2020. The newest data can be found at: https://ecf.com/dashboard

	Outcome: Cyclist count			
	(1) All km	(2) Bike lane km	(3) Num of measures	(4) Any treatment
Pop-up treatment	0.006^{**}	0.007^{**}	0.004^{*}	0.061^{*}
	(0.003)	(0.003)	(0.002)	(0.036)
City clusters	78	78	78	78
N	59904	59904	59904	59904

Table S3. Different treatment specifications. Each column shows the effect of treatment with pop-up infrastructure on a city's cycling count compiled from city APIs. The data on daily pop-up bike lane additions is from the European Cyclists' Federation (1). The newest data can be found at: https://ecf.com/dashboard. The unit of observation is the cycling counter. Time variation is daily. Coefficients are from Poisson regressions. Column (1) shows the effect of a kilometer of any bike infrastructure, (2) shows the effect of a kilometer of bike lanes, (3) the effect of any single measure in a city, and (4) the overall treatment of an implemented pop-up infrastructure program in a city. All regressions include counter and day fixed effects and controls for overall mobility (measured with Facebook user movements) (4), weather (temperature, wind, sunshine, precipitation) (7), and the number of counters active on a given day in a city. We cluster standard errors (in parentheses) at the city level, where treatment is assigned. Significance levels are * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

	(1) Preferred	(2) Transit sample	(3) Large sample	(4) Transit sample w/o Paris	(5) Large sample w/o Paris
Pop-up treatment	30.183^{***}	30.170^{***}	11.197^{**}	34.079^{***}	15.582^{***}
	(8.855)	(8.895)	(5.610)	(6.729)	(3.330)
N	29596	29596	63617	28719	62726

Table S4. Additional robustness checks for our preferred estimate. Each column shows the effect of treatment with pop-up infrastructure on a city's cycling count compiled from city APIs. The data on daily pop-up bike lane additions is from the European Cyclists' Federation (1). The newest data can be found at: https://ecf.com/dashboard. The unit of observation is the cycling counter. Time variation is daily. Coefficients are from Poisson regressions. Column (1) shows our baseline estimate including the transit control, (2) shows an estimate in the same sample as (1) but without the transit control, (3) the effect in the large sample including cities for which the Apple transit variable does not exist, (4) the same specification as (2) but dropping Paris from the sample, and (5) the same specification as (3) but without Paris. All regressions include counter, city-week, and country-day fixed effects. They also include controls for overall mobility (measured with Facebook user movements) (4), weather (temperature, wind, sunshine, precipitation) (7), and the number of counters active on a given day in a city. We cluster standard errors (in parentheses) at the city level, where treatment is assigned. Significance levels are * p $<$ 0.1, ** p $<$ 0.05, *** p $<$ 0.01.

	(1)	(2)	(3)	(4)	(5)	(6)
Pop-up bike lanes (km)	7.585***	7.595***				
	(2.545)	(2.538)				
Pop-up bike lanes (km per km^2)			30.183***	30.170***		
			(8.855)	(8.895)		
Pop-up bike lanes (km per capita)					23.295***	23.362***
					(6.124)	(6.148)
Temperature	0.229***	0.231***	0.229***	0.231***	0.229***	0.230***
	(0.048)	(0.049)	(0.048)	(0.049)	(0.048)	(0.049)
Precipitation	-0.035***	-0.036***	-0.035***	-0.036***	-0.035***	-0.036***
	(0.010)	(0.010)	(0.010)	(0.010)	(0.010)	(0.010)
UV Radiation	0.158***	0.158***	0.159***	0.159***	0.159***	0.159***
	(0.023)	(0.023)	(0.024)	(0.024)	(0.024)	(0.024)
Wind (u component)	-0.010	-0.010	-0.011	-0.010	-0.011	-0.011
	(0.009)	(0.010)	(0.009)	(0.010)	(0.009)	(0.010)
Wind (v component)	-0.008	-0.008	-0.008	-0.008	-0.007	-0.008
	(0.009)	(0.009)	(0.009)	(0.009)	(0.009)	(0.009)
Active Counters (day)	-0.009	-0.012	-0.010	-0.012	-0.010	-0.012
	(0.011)	(0.011)	(0.011)	(0.011)	(0.011)	(0.011)
Overall mobility	1.081***	1.114***	1.076***	1.109***	1.073***	1.105***
	(0.108)	(0.106)	(0.105)	(0.102)	(0.104)	(0.100)
Transit	0.001		0.001		0.001	
	(0.001)		(0.001)		(0.001)	
Counter FE	Y	Y	Y	Y	Y	Y
City-week FE	Y	Y	Y	Y	Y	Y
Country-day FE	Y	Y	Y	Y	Y	Y
City clusters	18	18	18	18	17	17
N	29,596	29,596	29,596	29,596	29,210	29,210
Pseudo R^2	0.932	0.932	0.932	0.932	0.932	0.932

Table S5. Estimates of the average effect of pop-up bike lanes on cycling. Estimates are from the preferred specifications marked in blue in Figure 3 in the main body (Equation 1). These are Poisson regressions using the level of the cyclist count. The unit of observation is the bike counter and data varies daily. Treatment is defined in kilometers, km per capita, or km per km^2 of pop-up infrastructure in service in a city on a day. Treatment effects are scaled to the mean treatment intensity in the sample. Data for the treatment is from the European Cyclists' Federation (1) and data for the outcome is from municipal bike counters (Materials and Methods). All regressions include controls for the number of active counters in a city on a given day and for the weather (temperature, sunshine, wind, precipitation; all standardized) (7). All regressions include a control for overall mobility (4). The transit control is from Apple routing requests (3). All regressions include fixed effects at the counter, city-week, and country-day level. We cluster standard errors (in parentheses) at the city level, where treatment is assigned. Significance levels are * p < 0.1, ** p < 0.05, *** p < 0.01.

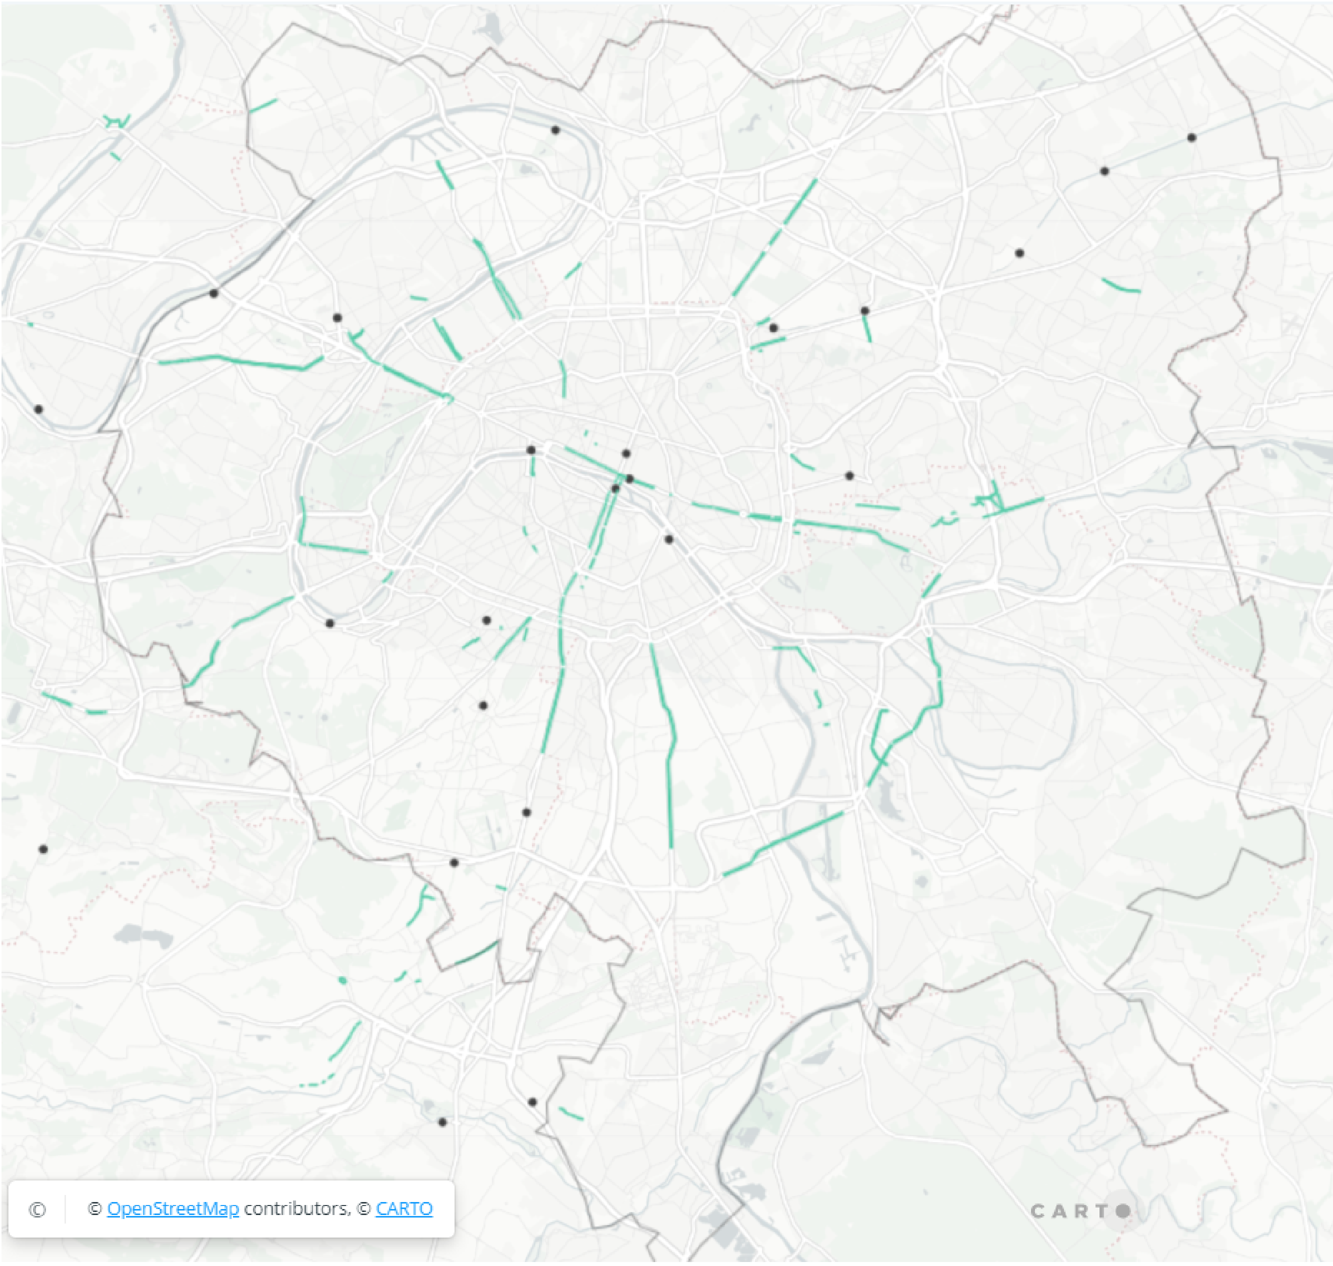

Fig. S1. Pop-up bike lanes and bicycle counters in Paris. The map shows pop-up bike lanes implemented in Paris as of July 3, 2020 (green lines) and the location of bike counters (dots) in our data set. The detailed infrastructure data has been collected by a consortium of French NGOs and researchers. It is available at: https://carto.parlons-velo.fr/#10.13/48.8312/2.5506

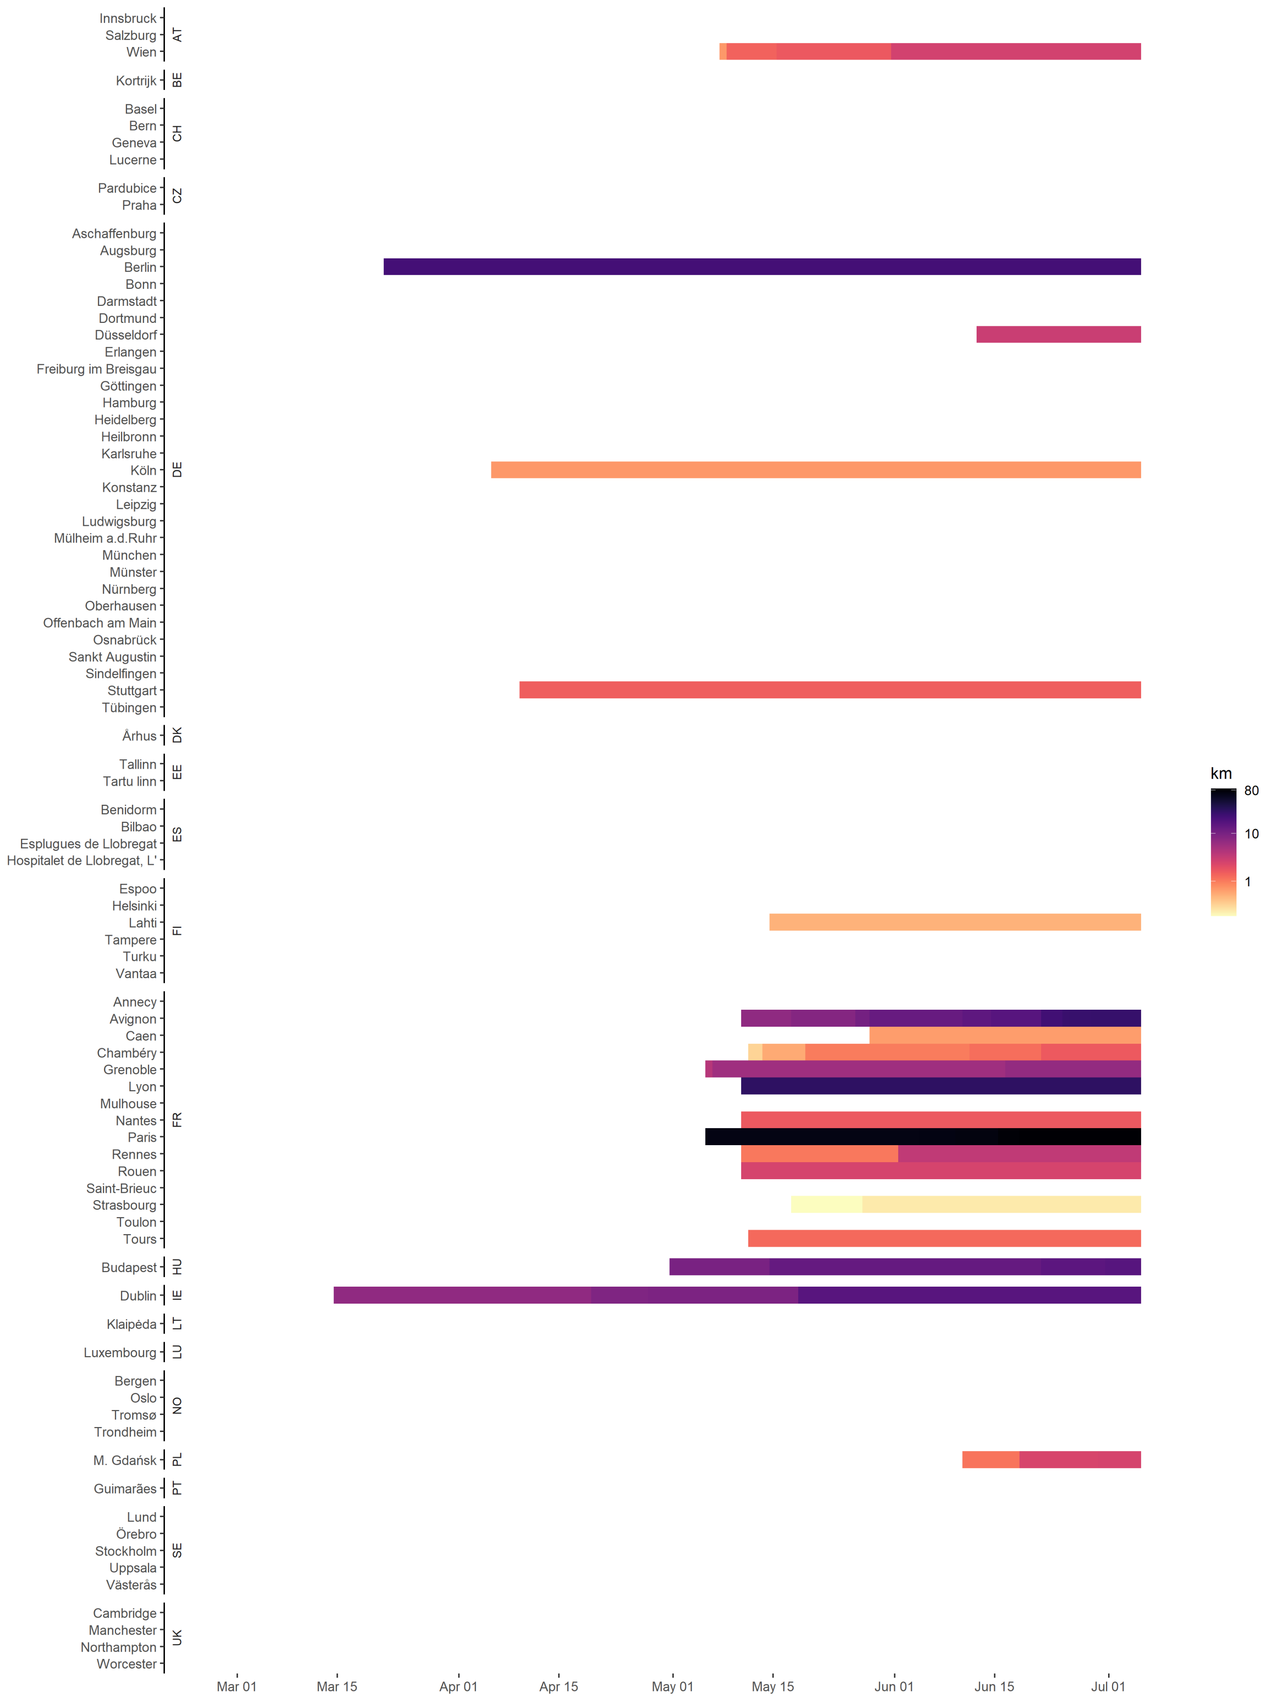

Fig. S2. Intensity of pop-up bike lane treatment over time in treatment cities and control cities. This Figure shows treated cities and their treatment intensities in terms of kilometers (coloring on a log scale) of public bike lanes in service on a given day between March and July 2020. Note, that some control cities have implemented pop-up bike lanes after July 2020. London, Milan, Lisbon and Rome are missing from the sample due to insufficient spatial or temporal coverage of the data. Information on individual pop-up bike lanes with their street location, announcement date, and implementation status is from the European Cyclists' Federation (1). The newest data can be found at: https://ecf.com/dashboard

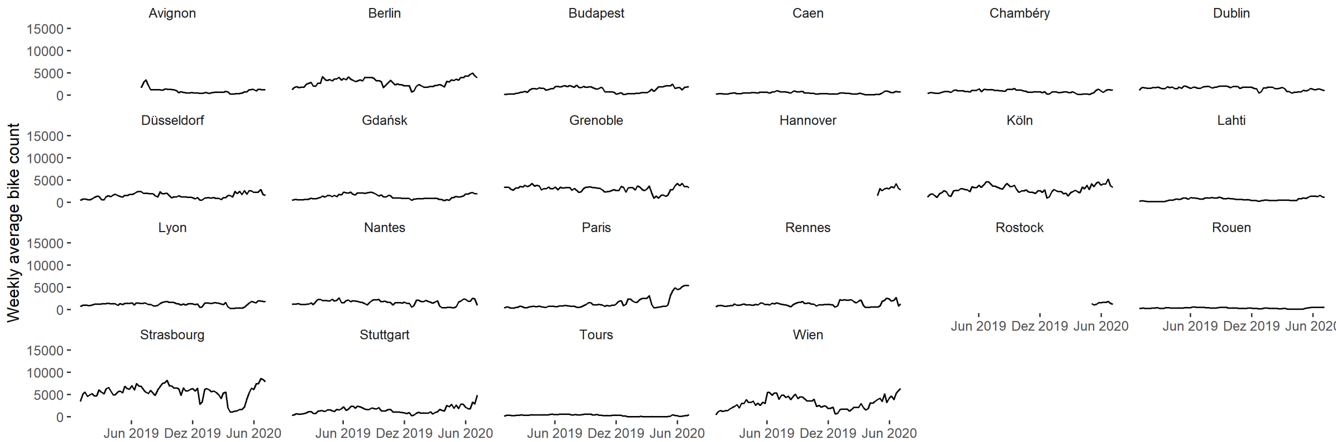

Fig. S3. Average bike count per week in treated cities. Daily bike counts are aggregated by city and averaged over the week. Bike counts are assembled from municipal open data feeds. The lower and upper percentiles of the base sample (treated and control cities combined) are removed from the estimation sample. Only measurements from 2019 and 2020 are shown. City definitions are chosen according to the EU Urban Audit (2).

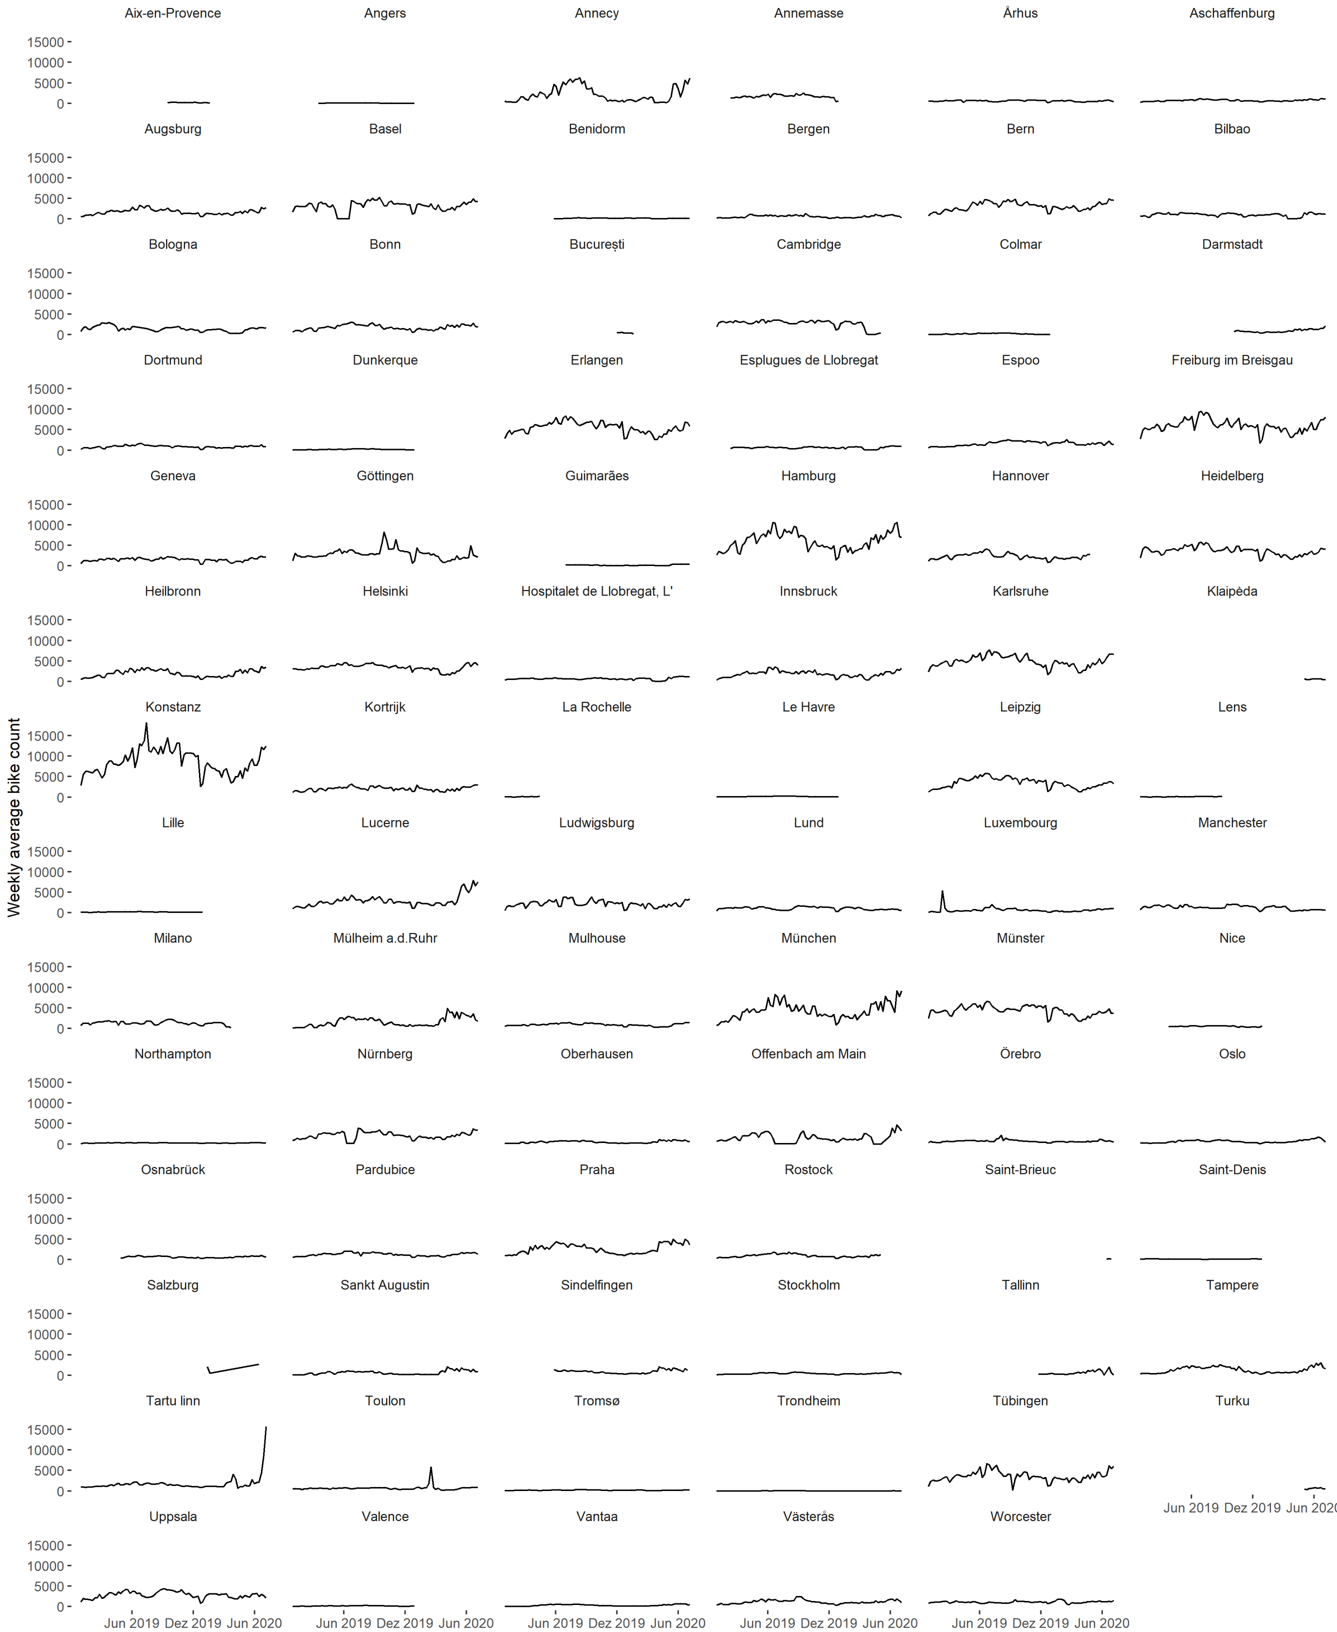

Fig. S4. Average bike count per week in control cities. Daily bike counts are aggregated by city and averaged over the week. Bike counts are assembled from municipal open data feeds. The lower and upper percentiles of the base sample (treated and control cities combined) are removed from the estimation sample. Only measurements from 2019 and 2020 are shown. City definitions are chosen according to the EU Urban Audit (2).

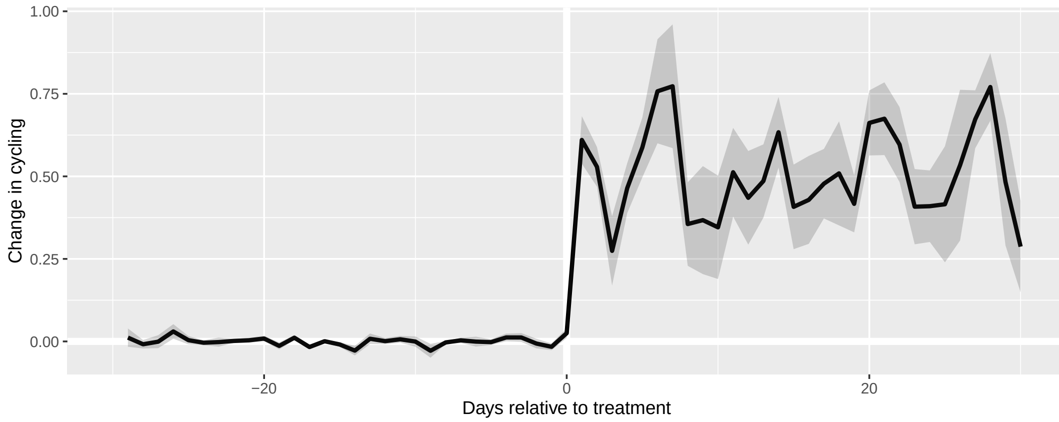

Fig. S5. Change in daily cycling after the first day of treatment for each city (see Equation 1). The Figure shows the average treatment effect on the treated (ATT) of treatment with pop-up bike lanes based on the *Matrix Completion* method (5). We implement Matrix Completion at the counter-day level and include controls for the number of active counters in a city on a given day and for the weather (temperature, sunshine, wind, precipitation) (7). The lighter grey area shows 95% confidence intervals based on 5000 bootstrap runs clustered at the city level. The sample for this Figure is restricted to 30 days before and after the first treatment day for each counter/city. Note, that estimates are not converted to % changes. The Figure is implemented with the *gsynth* package (8).

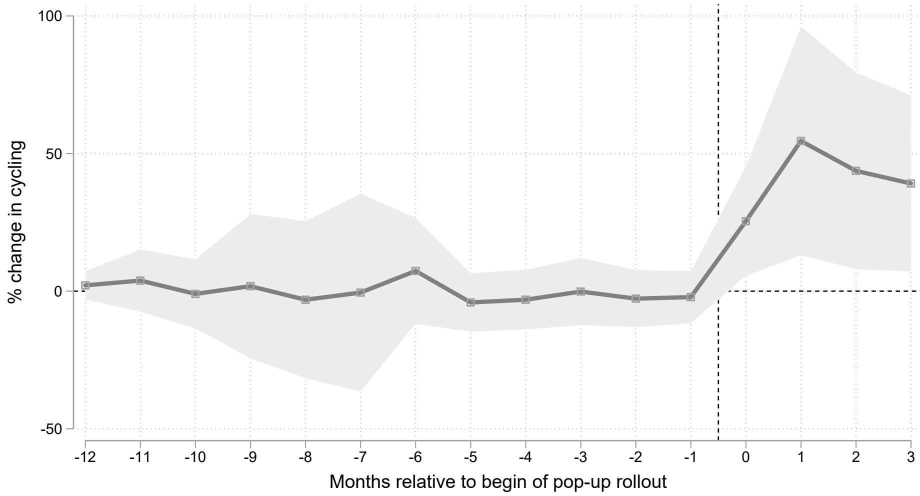

Fig. S6. Treatment effect (difference between treated and control cities) in months before and after the beginning of the pop-up bike lane policy. This is the same specification as shown in Figure 2 in the main body (Equation S2) except for the inclusion of weather controls (7). Observations are binned into months. Treatment for this plot is hard-coded to March 2020 and the baseline category and the begin of the sample are set to February 2019. Estimates are from Poisson regressions that include city and country-day fixed effects. The shaded area shows the 95% confidence interval. Data for the outcome variable is from the European Cyclists' Federation (1) and data for the treatment variable is from municipal bike counters (Materials and Methods).

References

1. ECF, COVID-19 Cycling Measures Tracker (2020) https://ecf.com/dashboard.
2. Eurostat, Urban Audit geodata (gisco) (2020).
3. Apple, Mobility Trends Reports (2020) https://www.apple.com/covid19/mobility.
4. Facebook, Movement Range Maps (2020) https://data.humdata.org/dataset/movement-range-maps.
5. S Athey, M Bayati, N Doudchenko, G Imbens, K Khosravi, Matrix completion methods for causal panel data models, (National Bureau of Economic Research), Working Paper 25132 (2018).
6. S Cunningham, G DeAngelo, J Tripp, Craigslist Reduced Violence Against Women, Working paper (year?).
7. H Hersbach, et al., The era5 global reanalysis. *Q. J. Royal Meteorol. Soc.* **146**, 1999–2049 (2020).
8. Y Xu, Generalized Synthetic Control Method: Causal Inference with Interactive Fixed Effects Models. *Polit. Analysis* **25**, 57–76 (2017).