The effects of congestion pricing in New York City on air pollution*

Maxim Massenkoff
Naval Postgraduate School

Jesse Rothstein
UC Berkeley

Reed Walker
UC Berkeley

January 4, 2025

Pre-analysis plan

Abstract

On January 5th, 2025, New York City will implement the first congestion pricing program in the US. This document pre-registers an analysis of the effects of the law on pollution.

1 Introduction

New York City's congestion pricing program is the first in the US. In this document, we pre-register an analysis of the law's effects on air pollution in the congestion pricing area and in surrounding areas in New York and New Jersey. We use data from PurpleAir sensors, the Environmental Protection Agency (EPA), and New York City.

We pre-register our analysis because:

1. This is a setting where analysis requires a large number of researcher decisions, including, for example, whether to adjust for weather, how to treat missing data, and what time horizons to use to summarize the effects. In particular, the needed data comes from multiple sources with different quality and completeness; it may improve the reliability of estimates to set the inclusion criteria beforehand.

2. The consequences are likely to be debated.

*Massenkoff: maxim.massenkoff@nps.edu, Rothstein: rothstein@berkeley.edu, Walker: rwalker@berkeley.edu. For helpful feedback and comments, we thank Nick Delahaye. All errors are our own.

3. Since there is only one treated unit, the best causal inference strategy *ex ante* might be similar to the best strategy *ex post*. In other words, a study of the outcomes of the policy in 2026 might make fairly similar choices in specifying the counterfactuals and estimating the causal effects.[1]

Together, these point to large benefits to pre-registration, in terms of reduced scope for potential "p-hacking." On the other side, the costs of pre-registration are likely to be relatively small, as the analytic decisions we make before the policy goes into effect are likely similar to those we would make later.

We summarize our analytical choices in the pre-analysis checklist in Table 3. Our core results will measure effects over two time horizons: the first three months and the first year of the policy. Our primary pollution outcome is daily average PM2.5 measured using government and PurpleAir sensors. Our secondary outcomes are peak hour PM2.5 and daily carbon monoxide (CO) and nitrogen dioxide (NO2), measured from EPA sources. We measure pollution effects over several treated areas: the Manhattan Central Business District (where the congestion fee applies, abbreviated CBD), Brooklyn, the Bronx, New Jersey, Queens, Staten Island, and Manhattan north of the Central Business District.

To estimate treatment effects, we compare pollution in the treated areas to that in a "synthetic control," a weighted average of outcomes from sensors outside the New York area that is selected to track outcomes in the treated area in the period leading up to the policy implementation. We use the `multisynth` estimator of Ben-Michael et al. (2021b) and the `augsynth` estimator of Ben-Michael et al. (2021a) to construct the synthetic controls. See those papers for discussions of the advantages of these estimators relative to the "vanilla" synthetic control estimator. We show below that `multisynth` tends to give much more precise estimates, in placebo treatment effect analyses, than does another alternative, the `GSynth` estimator of Xu (2017).

2 Pollution Build

Here we describe the different sources we use for measuring PM2.5 and our secondary pollution outcomes, NO2 (measured in parts per billion) and CO (parts per million). From these different sources, we make a date-by-sensor dataset giving daily average pollution and daily average pollution restricting to peak hours.[2]

[1]Future researchers may gain access to sensor data not currently made available to the public, however.

[2]Peak hours are 5 AM to 9 PM on weekdays and 9 AM to 9 PM on weekends. If these change during our treatment period, we will update the average so that the daily hour range always tracks the peak hours that are in effect.

In what follows, **comparison states** refers to: Pennsylvania, New Jersey, Connecticut, Maryland, Rhode Island, Washington, DC, and Massachusetts when discussing our main outcome, PM2.5. For CO and NO2, we draw all EPA AQS sensor data in the United States since coverage is much sparser. To avoid spillovers, the comparison set always excludes sensors from any of our treated areas in New York City and nearby New Jersey as well as any sensors within 20 miles of the New York Central Business District.

Ultimately, we will have six builds:

1. **Main build**: daily average PM2.5 data, including PurpleAir sensors.
2. Main build excluding PurpleAir from treated areas (for robustness of the results from the build above).
3. **Peak hour build**: daily PM2.5 averages including only peak hours (as a secondary analysis).
4. Peak hour build excluding PurpleAir from treated areas (for robustness of the results from the build above).
5. Nitrogen dioxide build (as a secondary analysis): Daily average NO2 using EPA AQS data.
6. Carbon monoxide build (as a secondary analysis). Daily average CO using EPA AQS data.

Cleaning steps For all PM2.5 sources, hourly readings above 250 (0.03% of observations in PurpleAir, for example) are set to 250. Daily PM2.5 averages are capped at 150. For NO2, any daily average above 35 ppb is set to 35 (0.2% of observations in 2024 EPA AQS data) and hourly measurements are capped at 50 ppb (0.05% of EPA AQS hourly observations). CO daily averages are capped at 1 ppm and hourly measurements are capped at 2 ppm. Any hourly or daily reading below zero is set to zero.[3] We use a linear interpolation, discussed below, to fill in missing daily observations for each sensor.

PurpleAir We will download hourly sensor readings from Purple Air sensors from New York and the comparison states (PurpleAir, 2023). We use the CF3-ALT estimates of PM2.5 (Wallace and Ott, 2023). These sensors are low-cost and show some systematic biases compared to higher-quality EPA

[3]We specify the cleaning for both daily and hourly because our main analysis is of daily averages but we also perform a secondary analysis looking at specific hours.

sensors (Barkjohn et al., 2021). We do not use the corrections from Barkjohn et al. (2021) since these should be differenced out in our analysis, but we make several restrictions to improve data quality. We drop sensors that have an r-squared less than 0.6 when we regress channel A readings on channel B readings at the weekly level.[4] Further, we restrict to sensors whose position rating is either 4 or 5 stars, indicating that the sensor is close to the reported latitude and longitude based on Google's Geolocation API. As a robustness check, we will perform the same analyses leaving out PurpleAir sensors from the focal treated area.

New York City sensors The city posts some PM2.5 data from the New York City Community Air Survey at the New York City Environment and Health data portal (NYC, 2024). This currently includes 11 sensors located in New York City and is provided hourly.[5] We refer to these as the NYC Permanent sensors below.

EPA Sources We will download data from the Environmental Protection Agency's (EPA) Air Quality System (AQS) and AirNow service. The AirNow data is available almost instantly and contains measures of CO, PM2.5, ozone, and PM10. The AQS data are released with roughly a six-month lag.

We will pull data from both sources and may use only AirNow for preliminary analyses. However, the final results will be estimated once the EPA data is available. Whenever there is overlap (i.e., AirNow and EPA each have enough observations for a specific sensor to be included) we default to the EPA data because of its quality control processes. But if only AirNow data is available for a certain sensor, we still keep it since the sensor coverage is limited in our small treated areas.[6]

[4]This is based on a recommendation from PurpleAir. See, e.g., Merrin and Francisco (2022) for a discussion of quality control techniques with the PurpleAir data. We estimate these sensor-level regressions using the time window being employed since the quality of sensor ratings may change over time.

[5]In addition to this, New York City Department of Health and Mental Hygiene runs New York City Community Air Survey (e.g., NYCDOH, 2022), which incorporates data from mobile sensors, but these are released with a long lag. The data from 2023 will not be available until April 24, 2025. Moreover, our framework does not easily accommodate changing sensor location. We will not use the NYCCAS data.

[6]The AirNow fields we use are "PM2.5-24HR" for daily data and "PM2.5" for hourly data. The EPA parameter codes we use are (in order of preference) 88101, 88502, 88500, and 88501 (U.S. Environmental Protection Agency, 2006). In a recent such pull for Connecticut, Massachusetts, New Jersey, New York, and Pennsylvania, 73% of sensor-dates had a measurement for 88101, the highest-quality PM2.5 measure. The remaining observations were evenly split between 88501 and 88502. For the other pollutants, the parameter codes we use are 42101 for CO and 42602 for NO2.

2.1 Interpolation

The sensors are frequently missing data. To interpolate missing values, we will perform the following steps for each of the sensors, performed separately for post-treatment and pre-treatment periods:

- Estimate a regression predicting the pollution outcome with day of week dummies, month dummies, a quadratic in date, a quadratic in precipitation, linear controls and decile dummies for the minimum and maximum temperatures, and a holiday dummy.
- Calculate interpolated pollution outcome as the fitted value from the regression plus the simple linear interpolation of the regression residuals from the complete-data points on either side of the missing observation.

In other words, if we have data from July 17 and July 19, but not from July 18, we will interpolate the value for July 18 as the predicted value for July 18 from the regression from the first step, plus the simple average of the regression residuals for July 17 and July 19.[7]

2.2 Weather data

We will download daily weather station data from the Global Historical Climatology Network daily (GHCNd) (Menne et al., 2012b,a). For each pollution sensor in our data, we assign to it the daily weather data from its closest weather station (after dropping weather stations with too many missing values). Missing data is less common compared to the pollution sensors. On average, 3-5% of daily observations are missing in our demonstration build, with an average distance of 5 miles to the pollution sensor. We use maximum daily temperature, minimum daily temperature, and precipitation from the weather stations. We interpolate these with a simple linear regression,

$$y_{it} = \beta_0 + \sum_{m=2}^{12} \beta_m \text{Month}_m + \delta \text{Date}_t + \gamma \text{Date}_t^2 + \alpha_i + \varepsilon_{it}, \tag{1}$$

where y_{it} is a weather outcome on day t at station i, β_0 is a constant, Month_m gives month dummies, $\delta \text{Date}_t + \gamma \text{Date}_t^2$ is a quadratic in the numeric date, and α_i is a station fixed effect. As in the interpolation for the pollution sensors (Section 2.1), the imputed values are equal to the fitted value from Equation 1 plus the simple linear interpolation of the residuals.

[7] If several consecutive days are missing, all will be interpolated, using weighted averages of the surrounding complete-data residuals. Thus, if July 18 and July 19 are both missing, the July 18 interpolation will equal the predicted value for July 18 plus 2/3 of the July 17 residual and 1/3 of the July 20 residual.

3 Treated areas

We will study impacts on pollution in the following areas:

Figure 1: Map of sensors

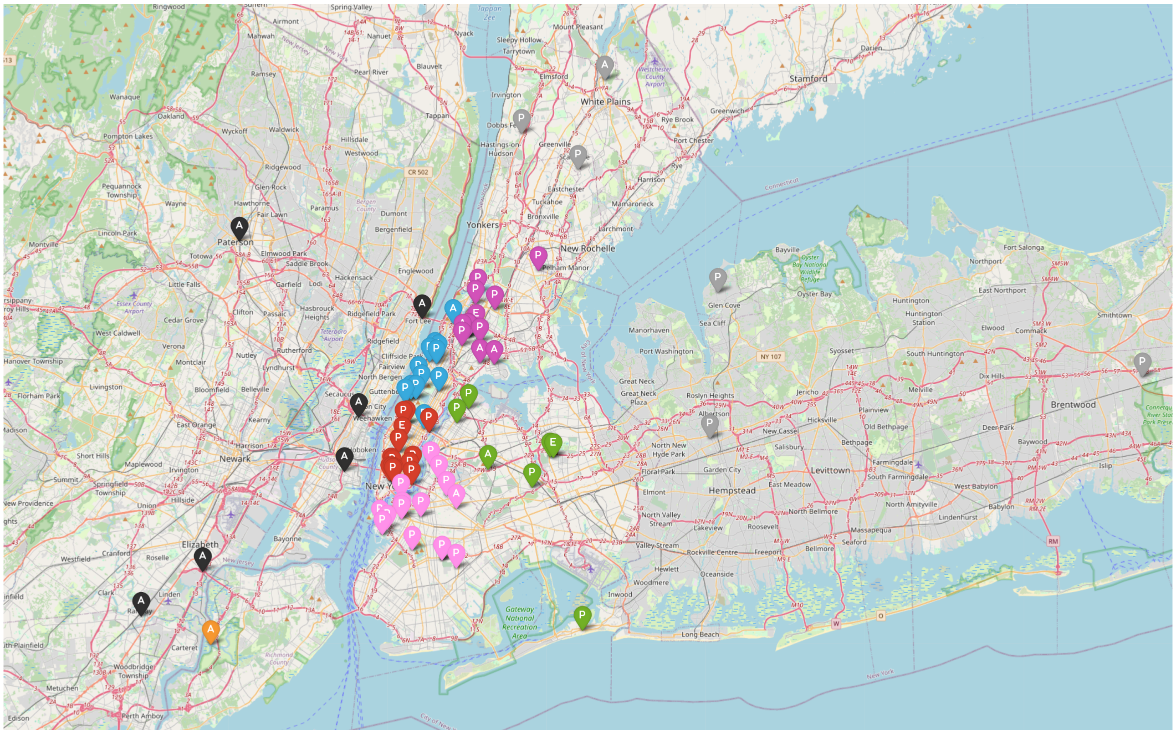

Notes: This shows a subset of sensors available for analysis as of December 2024 assuming a start date of March 2023 and before dropping sensors that have too many missing values. The colors indicate treatment status. Red: Manhattan CBD, Blue: Manhattan north of the CBD, Purple: Bronx, Green: Queens, Pink: Brooklyn, Black: New Jersey, Orange: Staten Island.

- The Manhattan Central Business District (CBD): This is where the congestion pricing fee applies.
- New Jersey within 20 miles of the center of the CBD[8]
- Manhattan north of the CBD
- Brooklyn County
- Bronx County
- Staten Island
- Queens

[8] The exact coordinates we use are (40°44'39.7"N 73°59'21.5"W).

Our comparison set for all analyses consists of all observations not in any of the above areas and at least 20 miles from the center of the CBD.

4 City permutation to choose analytical sample

We have two major build decisions to pin down. First, we need to choose a start date for the pre-treatment period. How far back should the data go? (The end of the post-treatment period is fixed at April 5, 2025 and January 5, 2026 for our three-month and 12-month analyses, respectively.) Second, we need to set a threshold for how much interpolation we are willing to do for a sensor with a great deal of missing data. Our sources for both pollution and weather often have periods with missing data, especially the NYC sensors.

This presents a tradeoff: a longer pre-treatment period will likely lead to a better synthetic control fit and more precision. However, since many NYC Permanent and EPA sensors have outages, and the PurpleAir network has been growing rapidly over the past few years, we will lose sensors to our missing data criterion as we move the start date back in time.

In order to choose these two parameters to maximize precision, we will conduct a series of placebo tests using other geographies in our comparison states. The placebo tests measure "treatment effects" over the same time horizons in other places that have not implemented congestion pricing. The average estimate for these cities should be close to zero, so we judge the quality of the parameter set by how close its average estimate is to zero.

For the PM2.5 analysis, we select the 20 largest geographies that have at least one EPA/AirNow sensor and at least three PurpleAir sensors under the following loose parameters: up to 50% missing days with a January 1, 2024 start date. We exclude data from New York City and its surrounding areas. We will first consider cities, ranking them by 2020 Census population. If we exhaust eligible cities before reaching 20 units, we will continue the selection process using counties, excluding any counties that contain previously selected cities.

Next, we will estimate synthetic control models treating that geography as if it implemented congestion pricing on January 5, 2025, always omitting our actual treated areas plus any sensor within 20 miles of the Manhattan CBD from the analysis. We will estimate these placebo models across a grid of potential start dates for the pre-treatment period and various thresholds for allowable missing observations, potentially allowing different thresholds in the pre- and post-treatment periods.

The full list of parameters is

- Missing percent: {10, 20, 30, 40, 50}
- Start date of pre-treatment period: {January 5, 2022 up to July 5, 2024, incremented by 3 months}

For example, for one set of placebo analyses, we will start the pre-treatment period on January 5, 2022 and sensors and weather stations will be dropped if they are missing over 10% of days in either the period from January 5, 2022 to January 5, 2025 or the period from January 6, 2025 to January 6, 2026. For some parameter sets, sensors will drop out. We necessarily exclude geographies with zero sensors, and skip a parameter set if this leaves us with fewer than 15 geographies.

For each combination of parameters, we will compute the placebo treatment effect for each placebo treatment city, averaged across all sensors in that city and across all observations. We will then compute the mean squared placebo treatment effect, giving each city equal weight. For our final analysis, we will then select the start date and missing data threshold combination that minimizes this mean squared placebo treatment effect. This process is run separately for each of the outcomes and time horizons.

It's possible that the parameter set that is optimal with this criteria will yield a sample with zero treated sensors in one or more of the treated areas (e.g., Staten Island) in a particular build. If this occurs, we will select instead the best performing parameter values among those that yield at least one selected sensor in that treated area. We perform this process for each of the builds described in Section 2. So for builds that do not use PurpleAir sensors, the initial selection step just requires a single EPA sensor in that geography with enough non-missing data.

Table 1: Sensors used for demonstration build

TreatmentArea	AirNow	EPA	PurpleAir	NYC_Perm	Mean PM2.5
CBD	0	0	15	2	6.21
Staten Island	1	0	0	0	7.83
Queens	2	1	4	1	7.97
Manhattan Above CBD	2	0	11	1	5.08
Brooklyn	1	0	23	0	6.24
New Jersey	6	0	0	0	8.90
The Bronx	3	1	9	1	9.23

Notes: Each row shows a different treated area. The middle four columns show how many sensors from the specified source are used in that treated area. For example, there are 15 PurpleAir sensors in the CBD and 23 in Brooklyn in this particular build.

5 Example of the sensor layout in demonstration build

Table 1 shows the attributes of a build which uses a mix of data from PurpleAir, the EPA, and New York City. This was performed for the purposes of demonstration and the power analysis below. It includes data from June 2022 to November 2024.

The table shows that there are few government sensors available in the CBD. There are many more PurpleAir sensors, though none in Staten Island.[9] Only one AirNow sensor was available in Brooklyn given the date range and inclusion criteria for missing values.

6 Visual assessment

In Figure 2, we show the results of a placebo exercise using the demonstration build. The red dashed line shows the daily counterfactual PM2.5 for the Manhattan CBD as determined by the synthetic control, while the blue solid line shows actual daily PM2.5 in the Manhattan CBD. Data prior to the placebo treatment date of April 1, 2024 are used to choose the synthetic control weights. The two series track each other closely, including during the large spike in pollution from the Canadian wildfires in June 2023 (leftmost side of the plot). Overall, this suggests that the design is well-positioned to measure the effects of congestion pricing. Below, we examine power more systematically.

[9] Our 2024 pull of the PurpleAir data did not include sensors in New Jersey, so the table shows zero PurpleAir sensors there. Our eventual data pull will include New Jersey sensors.

Figure 2: Example of the augsynth fit with the Manhattan CBD as treated area

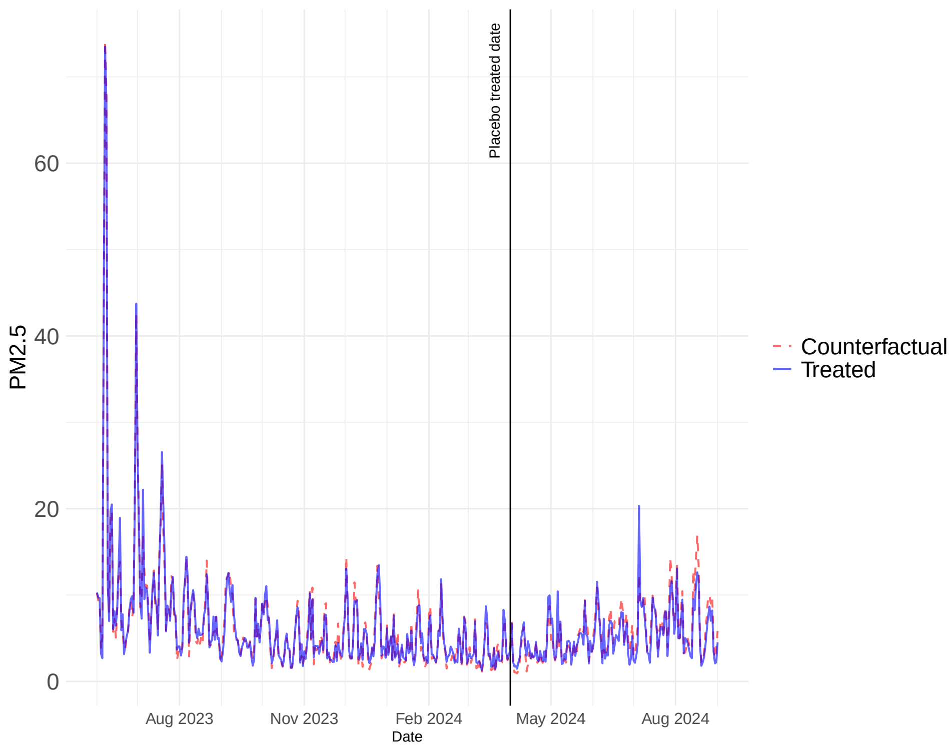

Notes: This shows an example of fitting a synthetic control on the daily average PM2.5 data, with the Manhattan CBD as the treated area. We use a placebo treated date of April 1, 2024 (shown in the vertical line).

7 Power analysis

We probed the power by estimating effects for several placebo dates in the data and several estimation methods. Table 2 shows the results. We consider three ways of forming the counterfactual: The simple average of all non-treated observations (described as method "none" in the table), the multisynth estimator of Ben-Michael et al. (2021b), and the GSynth estimator of Xu (2017). We switch to `agusynth` with ridge augmentation whenever we have only one sensor, which for the demonstration build was true only for Staten Island.

We show estimates of the distribution of placebo treatment effects in each of seven geographic areas. For each area, we sort by the width of the distribution of average ATTs from smallest to largest. The first panel of the table shows our analysis of effects in the CBD, the main treated area. Focusing on the multisynth estimates in the first row, we estimate that we should be able to detect about a 0.26 point or 4.6% change in PM2.5 pollution. The GSynth estimates are less precise.

This detectable effect is smaller than the changes predicted in Manhattan. Our main reference for the expected effects is Metropolitan Transportation Authority (2023). The report projects an 11% drop in PM2.5 pollution in the Manhattan CBD. Our analysis should be able to detect an effect of this magnitude.

Subsequent panels show the power analyses for spillover effects on other areas. In each, the multisynth estimates are much more precise than the GSynth estimates; in several, they are also more precise than the unweighted estimates. But detectable effects are larger than those predicted in Brooklyn, Queens, the Bronx, and the treated parts of New Jersey. For example, Metropolitan Transportation Authority (2023) predicts a 2% increase in Staten Island (see Figure 10-8). We will have trouble detecting effects of this magnitude.

These power calculations should be taken as only a rough guide because the projections were made assuming a $15 toll (it was subsequently decreased to $9) and because the modeling exercise is inherently uncertain. Taken literally, however, we should be powered to detect the expected decreases within the congestion zone. We are not powered to detect the expected increases outside of the congestion zone, as these changes are much smaller. However, given the policy interest in negative spillovers from congestion pricing (e.g. Howard, 2024), our results may help to rule out especially large effects.

Table 2: Placebo runs for confidence intervals

Treatment	Method	Outcome	CI width over 2	Treated mean	N dates
CBD	multisynth	PM2.5	0.263	5.716	64
CBD	none	PM2.5	0.368	5.716	64
CBD	GSYN	PM2.5	1.097	5.716	64
Brooklyn	multisynth	PM2.5	0.105	6.236	64
Brooklyn	none	PM2.5	0.140	6.236	64
Brooklyn	GSYN	PM2.5	0.442	6.236	64
Manhattan Above CBD	multisynth	PM2.5	0.660	7.558	64
Manhattan Above CBD	none	PM2.5	0.863	7.558	64
Manhattan Above CBD	GSYN	PM2.5	1.357	7.558	64
New Jersey	multisynth	PM2.5	0.364	8.676	64
New Jersey	none	PM2.5	0.430	8.676	64
New Jersey	GSYN	PM2.5	0.699	8.676	64
Queens	none	PM2.5	0.200	8.816	64
Queens	multisynth	PM2.5	0.295	8.816	64
Queens	GSYN	PM2.5	0.846	8.816	64
Staten Island	Ridge	PM2.5	0.528	7.036	64
Staten Island	none	PM2.5	0.598	7.036	64
Staten Island	GSYN	PM2.5	1.836	7.036	64
The Bronx	multisynth	PM2.5	0.188	8.145	64
The Bronx	none	PM2.5	0.312	8.145	64
The Bronx	GSYN	PM2.5	0.625	8.145	64

Notes: This table shows the confidence intervals from several placebo runs in order to probe the power of our design. Each row describes a set of placebo runs where we vary the date that the treated area is "treated" within that build's date range, leaving periods at the beginning and end for fitting the weights and estimating an effect. From each placebo treatment date, we calculate the average ATT using the first year of post-treatment ATT estimates. The confidence interval is based on this distribution of average ATTs. Across rows, we change the synthetic control method that we use. Staten Island does not have a multisynth run because it has only one sensor.
The columns are as follows:
Treatment: The area for which effects will be estimated, more details above.
Method: the synthetic control technique used. Multisynth is the R `multisynth` package with defaults, "GSYN" uses the `gsynth` package with defaults (Xu, 2017), Ridge is `augsynth` with Ridge, "none" is vanilla synthetic control (estimated in `augsynth`).
Outcome: PM2.5, with interpolation as discussed above.
CI width over 2: half of the 95% confidence interval width based on the distribution of average one-year ATTs.
Treated mean: the mean of the outcome in the treated group.
N dates: the number of placebo treatment dates used to estimate the distribution of treatment effects.

Table 3: Pre-analysis checklist: Pollution outcomes

	Item	Brief Description	Followed plan?
1	**Sample and Data**	Daily average PM2.5 from PurpleAir sensors, NYC permanent sensors, and EPA sensors. Weather data from GHCNd. Data start determined by the power analysis above (Section 4).	
2	**Primary Outcomes**	Average daily PM2.5.	
3	**Treatment definition**	Treatment begins January 5, 2025, when NYC implements congestion pricing. We estimate effects separately for seven treated areas: Manhattan CBD, the Bronx, New Jersey, Queens, Staten Island, Brooklyn, and Manhattan north of CBD (Section 3).	
4	**Causal design / controls**	We estimate effects using `augsynth` or `multisynth` (when the treated areas have multiple sensors) with defaults. Counterfactual PM2.5 is constructed from sensors in Northeastern US excluding other treated areas and a 20-mile doughnut around the CBD.	
5	**Primary estimates**	Average treatment effects on daily PM2.5 over the first three months and first year of the policy. Effects estimated separately for each treated area.	
6	**Missing data and cleaning**	Drop sensors missing some percent of observations in either the post-treatment period or pre-treatment period, with missingness threshold determined as outlined in Section 4. Linear regression to interpolate missing weather and pollution data.	
7	**Inference**	Wild bootstrap when using `multisynth`, conformal inference when using `augsynth`.	
8	**Secondary analyses**	i.) Estimate effects using only daily average pollution during peak hours (i.e., times when congestion pricing is in effect). ii.) Estimate effects on daily average NO2 and CO using nationwide data from EPA sensors.	
9	**Robustness checks**	i.) Estimate PM2.5 effects without using PurpleAir sensors in the treated areas.	
10	**Contingencies**	If policy canceled before one of the end points, estimand shifts to average change during implementation period.	

Notes: This table gives a brief checklist of analysis steps for pollution outcomes, with references to the full explanation of the decisions. When our analysis is complete, we will fill in the "Followed plan?" column with either "Yes" or an explanation.

References

Barkjohn, Karoline K, Brett Gantt, and Andrea L Clements, "Development and application of a United States-wide correction for PM 2.5 data collected with the PurpleAir sensor," *Atmospheric Measurement Techniques*, 2021, *14* (6), 4617–4637.

Ben-Michael, Eli, Avi Feller, and Jesse Rothstein, "The augmented synthetic control method," *Journal of the American Statistical Association*, 2021, *116* (536), 1789–1803.

_ , _ , and _ , "Synthetic Controls with Staggered Adoption," *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 12 2021, *84* (2), 351–381.

Howard, Hilary, "Congestion Pricing Could Bring Cleaner Air. But Maybe Not for Everyone.," *The New York Times*, may 2024, p. A19. Updated June 1, 2024. Print version appears in New York edition on May 30, 2024, Section A, Page 19.

Menne, Matthew J, Imke Durre, Russell S Vose, Byron E Gleason, and Tamara G Houston, "An overview of the global historical climatology network-daily database," *Journal of atmospheric and oceanic technology*, 2012, *29* (7), 897–910.

Menne, M.J., I. Durre, B. Korzeniewski, S. McNeal, K. Thomas, X. Yin, S. Anthony, R. Ray, R.S. Vose, B.E. Gleason, and T.G. Houston, "Global Historical Climatology Network - Daily (GHCN-Daily), Version 3," *NOAA National Climatic Data Center*, 2012. http://doi.org/10.7289/V5D21VHZ. Readme at this link: https://www.ncei.noaa.gov/pub/data/ghcn/daily/readme.txt.

Merrin, Zachary and Paul W Francisco, "Pilot Study for Multifamily Building Ventilation and Indoor Air Quality," Technical Report, Oak Ridge National Laboratory (ORNL), Oak Ridge, TN (United States) 2022.

Metropolitan Transportation Authority, "Central Business District Tolling Program Environmental Assessment," Environmental Assessment, Metropolitan Transportation Authority, New York, NY 5 2023. Final Environmental Assessment released on May 11, 2023.

NYC, "Real-Time Air Quality: PM2.5 in NYC," *NYC.gov Health EH Data Portal*, 2024.

NYCDOH, "The New York City Community Air Survey: Neighborhood Air Quality 2008-2020," *City of New York*, 2022. https://nyccas.cityofnewyork.us/nyccas2022/report/3.

PurpleAir, "API History Fields Descriptions," *PurpleAir*, 2023. https://community.purpleair.com/t/api-history-fields-descriptions/4652.

U.S. Environmental Protection Agency, "Technical Note on Reporting PM2.5 Continuous Monitoring and Speciation Data to the Air Quality System (AQS)," Technical Note, U.S. Environmental Protection Agency 6 2006.

Wallace, Lance and Wayne Ott, "Long-Term Indoor-Outdoor PM2. 5 Measurements Using PurpleAir Sensors: An Improved Method of Calculating Indoor-Generated and Outdoor-Infiltrated Contributions to Potential Indoor Exposure," *Sensors*, 2023, *23* (3), 1160. Forum post: https://community.purpleair.com/t/calibration-of-purpleair-monitors/482.

Xu, Yiqing, "Generalized synthetic control method: Causal inference with interactive fixed effects models," *Political Analysis*, 2017, *25* (1), 57–76.