

Santander Cycles: Getting More from Every E-bike

Product and BI case study · Transport for London open data · April 2025 to March 2026

July 2026

Data: 9,048,212 Santander Cycles hires, TfL journey files and BikePoint API

Methods: difference-in-differences, mix-adjusted benchmarking, unit economics, experiment design

Srujan S. Pailwar

Contents

1. Executive summary	3
2. Situation, complication and question	4
3. Data and approach	4
4. Findings	6
4.1 The rollout added fleet, not demand	6
4.2 The economics of the marginal e-bike	7
4.3 What drives e-bike choice	8
4.4 Which stations are under-served	8
5. Options assessed	10
6. Recommendations and test design	10
7. BI dashboard specification	10
8. Limitations	11
9. References	11
Appendix A. Preparing the trip measures (Python)	12
Appendix B. Rollout test and figures (Python)	13
Appendix C. Station benchmark and test power (Python)	15

1. Executive summary

In August 2025 Santander Cycles put about 230 more e-bikes on London's streets, lifting the number of e-bikes in use by 12.2%. This paper treats that rollout as a product launch and asks the questions an operator would ask afterwards: did the extra fleet pay for itself, what actually drives e-bike demand, and where should the next pound of fleet investment go?

The answer in one line: adding e-bikes everywhere added little; the next gains come from placing them where long trips start and from fixing the price barrier that keeps casual riders at park and tourist stations off e-bikes.

1. **The rollout diluted rather than grew demand.** Against classic bikes as a control, e-bike trips rose 3.5% over two months (95% CI +0.4% to +6.7%), and the effect is not significant over four or six weeks. Trips per e-bike per day fell from 2.7 to 2.3. Each new e-bike earns roughly £270 to £880 a year, about a third of what an existing e-bike earns.
2. **Trip length, not neighbourhood, drives e-bike choice.** Riders choose an e-bike for 8% of trips under 1 km and 35% of trips over 5 km, while e-bike share barely moves with distance from central London (17% to 19%).
3. **Some stations are under-served once trip mix is accounted for.** A mix-adjusted benchmark finds park and tourist stations, led by Hyde Park Corner and Kensington Gardens, running 5 to 7 points below the e-bike share their trips predict. Commuter stations at Angel and King's Cross run up to 15 points above it. The pattern points to price sensitivity among casual riders, not a lack of bikes.
4. **The next decision should be tested, not rolled out.** A 25-versus-25 station test over eight weeks can detect a lift of about 10% to 12%, which is the right bar for a fleet investment.

Recommendations: (1) re-deploy, rather than buy, e-bikes towards the 50 stations that start the most long trips; (2) pilot a casual-rider e-bike offer at park stations; (3) run both as controlled tests with the KPIs and dashboard set out in sections 6 and 7.

2. Situation, complication and question

Summary	
Situation	Santander Cycles carries around 25,000 hires a day. E-bikes make up about a sixth of the fleet but 18.4% of hires, and cost riders twice as much per 30 minutes (£3.30 versus £1.65).
Complication	E-bikes cost far more to buy, charge and maintain than classic bikes. The August 2025 expansion added 12% more e-bikes, but daily e-bike hires show no visible step change (Figure 2).
Question	Where does an extra e-bike earn its keep, and how should the operator decide before buying more?

Table 1. Situation, complication and question.

The analysis is structured as an issue tree. E-bike revenue per day is the product of e-bikes in service, trips per e-bike and revenue per trip (Figure 1). The rollout changed the first lever; this paper measures what happened to the other two.




Figure 1. KPI tree for e-bike revenue. The August rollout moved the first branch; placement and pricing sit on the other two.

3. Data and approach

Source	Content	Use
TfL journey files	Every hire, 1 April 2025 to 31 March 2026: start and end time, stations, bike model, minutes. 9,048,212 rows after removing duplicates across weekly files.	Daily counts, fleet, segmentation
TfL BikePoint API	Station names and coordinates; 796 terminals matched, 98.5% of hires.	Trip length and location
TfL fares	Classic £1.65 per 30 minutes; e-bike £3.30 per 30 minutes for casual riders; members pay £1 extra per 60 minutes on e-bikes.	Revenue per trip range

Table 2. Data sources.

Length, location and timing analyses use the 8,681,646 hires that last between 1 minute and 3 hours and start and end at matched stations. **E-bikes in use** is the number of distinct e-bikes making at least one trip in a month; TfL does not publish the fleet size, so bikes in workshops are not counted. TfL does not publish rider identifiers or membership status per hire, so revenue is shown as a range rather than a point estimate.

4. Findings

4.1 The rollout added fleet, not demand



Figure 2. Daily hires by bike type, 7-day centred mean. The shaded band marks when new e-bikes entered service (17 to 31 August 2025).

E-bike and classic hires rise and fall together through the year, which makes classic bikes a fair control: weather, school holidays and tube strikes move both, while the extra supply only moves e-bikes. Comparing the daily ratio of e-bike to classic hires before (June to July) and after (September to October) the rollout isolates the effect of the new bikes.




Figure 3. Change in e-bike hires relative to classic hires after the rollout, by comparison window, with 95% confidence intervals.

Window each side	Days	Effect	95% CI	p-value	Elasticity
4 weeks	28 + 28	+2.3%	-2.7% to +7.6%	0.371	0.20
6 weeks	42 + 42	+1.7%	-2.0% to +5.6%	0.369	0.15
2 months	61 + 61	+3.5%	+0.4% to +6.7%	0.027	0.30

Table 3. Difference in the daily log ratio of e-bike to classic hires. Elasticity is the share of an existing e-bike's trips that each new e-bike adds.

Only the two-month comparison is significant, and its interval almost reaches zero. The most generous reading is that 12% more e-bikes produced 3.5% more e-bike use: an elasticity of 0.3. The operator should not read this as a success. A naive trip-level test, treating each of 3.45 million hires as independent, would have reported a highly significant win ($z = 9.5$); because

hires on the same day share weather and events, days are the correct unit of analysis, and that changes the verdict.




Figure 4. Left: distinct e-bikes in use each month. Right: trips per bike per day by type.

The utilisation view confirms it. Trips per e-bike per day fell from 2.7 in June and July to 2.4 in September and 2.3 in October, while classic utilisation tracked the season. The same demand was spread over more bikes.

4.2 The economics of the marginal e-bike




Figure 5. Revenue range per e-bike per year. The low end assumes every e-bike trip is a member paying £1 extra; the high end assumes a casual rider paying £3.30.

The two-month estimate implies about 170 extra e-bike hires a day across the 233 bikes added, or roughly £270 to £880 of revenue per new e-bike a year. An existing e-bike at 2.3 trips a day earns roughly £840 to £2,770. **The marginal e-bike earns about a third of the average one.** Whatever an e-bike costs to buy, charge and maintain, payback on the August batch is about three times slower than the fleet average suggests. Fleet growth should therefore be judged on marginal, not average, utilisation.

4.3 What drives e-bike choice




Figure 6. E-bike share of hires by straight-line trip length.

Trip length is the dominant driver. E-bike share climbs from 8% under 1 km to 35% over 5 km; the median e-bike trip is 2.9 km in 15 minutes against 1.9 km in 13 minutes on a classic bike. Timing matters too (Figure 7): e-bike share peaks at 24% to 26% between 03:00 and 06:00 on weekdays and again at 20% in the evening commute, and is lowest around midday at 16%. Location matters least: share is 17% to 19% at every distance from central London.




Figure 7. E-bike share of hires by hour, weekdays and weekends.

4.4 Which stations are under-served

Raw e-bike share is a misleading guide to where bikes are needed, because stations differ in the trips that start there. I built a **mix-adjusted benchmark**: for each station, the e-bike share it would have if its riders behaved like the network average for the same trip length, weekday or weekend, and hour. The gap between actual and expected share, multiplied by trips, gives under- or over-performance in e-bike hires per day.




Figure 8. Actual versus mix-adjusted expected e-bike share for 699 stations with at least 5,000 hires. Bubble size is total hires.

Station	Actual share	Expected share	Gap, e-bike hires a day	Long trips a day
Hyde Park Corner, Hyde Park	10.2%	15.4%	7.8	20.1
Black Lion Gate, Kensington Gardens	10.9%	16.2%	4.9	17.3
Queen Street 1, Bank	13.2%	20.0%	4.8	24.8
Westminster Pier, Westminster	14.0%	19.5%	4.0	26.7
Albert Gate, Hyde Park	12.1%	15.6%	3.5	15.1
Palace Gate, Kensington Gardens	9.0%	16.1%	3.5	9.7
Natural History Museum, South Kensington	10.0%	18.2%	3.3	10.8
Orchard Place, Blackwall Tunnel	19.1%	34.3%	3.3	21.0

Table 4. The most under-served stations: e-bike share below what their trip mix predicts.

Two groups stand out. Under-served stations cluster around Hyde Park, Kensington Gardens, South Kensington's museums and Westminster: leisure and visitor territory. Over-served stations, such as Islington Green and Liverpool Road at Angel and York Way at King's Cross, are commuter stations where e-bike share runs 10 to 15 points above expectation. Visitors pay as they go, so an e-bike costs them twice a classic bike; members pay only £1 extra per hour. **The gap looks like price sensitivity, not missing bikes**, which makes it a pricing and product problem rather than a fleet problem. The benchmark cannot rule out low dock availability, which TfL does not publish, so the test in section 6 checks both explanations.

5. Options assessed

Option	Expected impact	Cost and risk	Assessment
A. Buy more e-bikes and spread them evenly	Low: elasticity of about 0.3	High capital cost; repeats August's dilution	Do not pursue
B. Re-deploy existing e-bikes to the top 50 long-trip stations	Medium: these stations start 16.3% of long trips	Rebalancing van time; no capital	Pursue, as a test
C. Casual-rider e-bike offer at park and visitor stations	Medium: about 345 e-bike hires a day of gap network-wide	Revenue per trip falls; may cannibalise classic	Pursue, as a pricing test
D. Do nothing	None	Low utilisation persists on the August batch	Baseline

Table 5. Options assessed.

6. Recommendations and test design

1. **Re-deploy before you buy.** Move part of the e-bike fleet towards the 50 stations that start the most long trips, prioritising those with a positive mix-adjusted gap. Do not buy more e-bikes until utilisation of the August batch recovers to the network average.
2. **Pilot a casual e-bike offer at park and visitor stations.** For example, a reduced first 30 minutes on e-bikes for pay-as-you-go riders starting at stations around Hyde Park, Kensington Gardens and Westminster. Measure net revenue, not hires, because the offer lowers revenue per trip.
3. **Make both decisions with controlled tests.** Randomly assign stations to treatment and control. Primary metric: e-bike hires per e-bike per day at the station. Guardrails: classic hires at the same stations (cannibalisation) and van rebalancing minutes (cost). Unit: station-day, with errors clustered by station. With 25 treated and 25 control stations over eight weeks, the test can detect a lift of about 10% to 12% at 80% power, based on the observed day-to-day variation in station hires.
4. **Report marginal, not average, utilisation.** Every fleet decision should show what the last bikes added are earning, not the fleet average.

Phase	Actions	Decision gate
Weeks 1 to 2	Select and randomise stations; build the dashboard; agree success thresholds	Thresholds signed off
Weeks 3 to 10	Run the re-deployment and pricing tests in parallel on separate station sets	Guardrails within limits each week
Weeks 11 to 12	Read out results; size network-wide impact	Scale, iterate or stop

Table 6. Twelve-week test plan.

7. BI dashboard specification

The tests need one shared view that operations, pricing and finance trust. The specification below is built to run on the same daily hire data in Power BI.

Page	Visuals and KPIs	Key measures
Fleet health	E-bikes in use, trips per e-bike per day (network and marginal batch), month-to-date and year-to-date hires	Trips per E-bike per Day = $\text{DIVIDE}([\text{E-bike Hires}], [\text{E-bikes in Use}] \times [\text{Days}])$
Station opportunity	Map and table of actual versus expected share; gap in hires a day; long trips a day	Expected Share = $\text{SUMX}(\text{Hires}, \text{RELATED}(\text{Benchmark}[\text{Share}])) / [\text{Hires}]$
Test readout	Treatment versus control trend; effect with confidence interval; guardrails	Lift = $\text{DIVIDE}([\text{Treated Ratio}], [\text{Control Ratio}]) - 1$
Revenue	Revenue range per trip and per e-bike; member versus casual mix when available	Revenue High = $[\text{E-bike Hires}] \times 3.30$

Table 7. Dashboard pages and measures.

8. Limitations

- **No rider identifiers or membership flag.** Retention, cohorts and exact revenue cannot be measured; revenue is shown as a range.
- **Fleet size is inferred** from bikes that made a trip; bikes in workshops and dock availability are invisible.
- **Before and after windows fall in different seasons.** The classic-bike control removes common shocks but not a seasonal shift towards longer leisure trips that favour e-bikes.
- **Straight-line distance** understates the route ridden, more so on longer trips.
- **The mix-adjusted benchmark** is descriptive: it shows where share is low relative to similar trips, not why. The pricing test is designed to find out.

9. References

- Angrist, J.D. and Pischke, J.-S. (2009) *Mostly Harmless Econometrics*. Princeton: Princeton University Press.
- Kohavi, R., Tang, D. and Xu, Y. (2020) *Trustworthy Online Controlled Experiments*. Cambridge: Cambridge University Press.
- Transport for London (2026) *Santander Cycles journey data*. Available at: <https://cycling.data.tfl.gov.uk> (Accessed: October 2026).
- Transport for London (2026) *BikePoint API*. Available at: <https://api.tfl.gov.uk/BikePoint> (Accessed: October 2026).
- Visit London (2026) *London cycle hire scheme: Santander Cycles*. Available at: <https://www.visitlondon.com/traveller-information/getting-around-london/london-cycle-hire-scheme> (Accessed: October 2026).