An independent Report on aspects of Heathrow's asset beta

Robin Mason
21 May 2026

0. Context

I have been asked by IATA to provide expert academic support to evaluate the methodology for estimating the asset beta for Heathrow Airport Limited (HAL) during the H8 price control period (2027–2031). The terms of reference for the work ask for a robust, evidence-based critique of the Civil Aviation Authority's and CEPA's beta estimation framework. I have been asked to look in particular at:

1. Whether the estimated beta series for the European airport comparators can be treated as draws from a statistically stable process and, if not, how that should affect the way the CAA sets a range for Heathrow's asset beta.
2. The quantification of relative risk differentials for comparator airports.

For the CAA's approach, I rely on its H8 Initial Proposals CAP 3232C. I have been provided with a number of reports by CEPA: in particular, their 18 November 2025 Report *Estimating the H8 cost of capital*. The views expressed in this Report are my own: this is an independent piece of work.

1. Summary

1.1. This Report addresses two linked questions in the CAA's H8 beta analysis for Heathrow. First, how should the range and central estimate for the pre-adjustment asset beta be chosen when the comparator beta series are unstable? Secondly, how should Heathrow-specific features—especially capacity constraints, single till and the Traffic Risk Sharing (TRS) mechanism—be reflected in the adjustment from listed airport comparators to Heathrow?

1.2. A central empirical point is that the listed airport comparator beta series do not behave as if they are stationary around a stable mean. Statistical tests provide very little support for treating the rolling beta estimates as draws from a single stable historical distribution. That finding does not necessarily mean that OLS beta estimation should be abandoned; nor does it justify selective deletion of particular periods or structural-break modelling. It means that the estimates should be understood as being sensitive to the data window and not interpreted as observations around a fixed long-run parameter.

1.3. This has a direct consequence for the CAA's beta range. The CAA is right to use recent spot estimates and to consider 2-year, 5-year and 10-year windows. But its pre-adjustment range of 0.47–0.70 gives decisive weight to the 10-year weighted-average estimate in setting the upper bound. In a non-stationary setting, that is too strong: a 10-year estimate is informative, but it also blends older regimes that may no longer be representative of H8. A materially better pre-adjustment asset beta range is **0.47-0.64**, with a **central** (not midpoint) **estimate of 0.53**.

1.4. The second part of the Report concerns the relative-risk adjustment from listed European airport comparators to Heathrow. The CAA's framework can be expressed as a reduced-form movement from the airport comparator beta toward a lower network-utility benchmark. This movement depends on two distinct issues: Heathrow's underlying exposure to traffic risk, and the extent to which regulation mitigates the traffic-related risk borne by equity investors.

1.5. That distinction is important. Capacity constraints are mainly an exposure issue. Heathrow's slot scarcity, excess demand and ability to refill traffic reduce the pass-through from market-wide demand shocks to realised passenger volumes. While comparator airports may be more constrained than in earlier periods, the gap with Heathrow has not closed fully. A central exposure parameter of **0.85**, rather than the CAA's midpoint of 0.80, is therefore justified.

1.6. Single till is mainly a mitigation issue. It does not primarily change Heathrow's physical exposure to passenger demand. It changes how far traffic-driven commercial gains and losses are retained by equity over time, because sustained commercial outperformance or underperformance feeds into future aeronautical allowances. CEPA's direct decomposition approach is too assumption-sensitive to support a precise standalone beta adjustment, but the underlying economic effect is real and should inform the mitigation parameter. A central mitigation parameter of **0.75**, rather than the CAA's 0.50 for TRS alone, is a conservative way to recognise the additional mitigation provided by single till.

1.7. Putting the two parts together produces a **recommended central post-adjustment asset beta of 0.42**, compared with the CAA midpoint of **0.50**. The conclusion is not that the CAA's approach is conceptually wrong, but that its implementation is **inconsistent** (in terms of beta estimation) and **incomplete** (in terms of relative risk adjustment). It is inconsistent because the CAA recognises instability in the beta evidence but then gives too much weight to long-window estimates in setting the range. It is incomplete because the CAA recognises traffic-risk mitigation through TRS but does not give sufficient effect to Heathrow's capacity position and single-till framework. The recommended approach is therefore a disciplined adjustment to the CAA's framework, not a wholesale departure from it. It results in an estimate of beta that is a material improvement on the CAA's.

I. Non-stationarity of the comparator beta series and the choice of range for Heathrow's asset beta

2. Why stationarity matters for beta estimation

2.1. Empirical beta estimation in regulatory cost of capital work aims to recover a forward-looking allowed return from historical market data. A long-run average beta has a clear interpretation only if the underlying statistical process is sufficiently stable: that is, only if observations can reasonably be treated as draws from a distribution with a stable mean.

2.2. If that condition is absent, an average over a long period does not recover a well-defined "true" beta. It becomes instead a mixture of different periods, each potentially reflecting

different market conditions, risk exposures or states of the world. The data remain informative; but a long-run average does not serve well as a forward-looking estimate.

2.3. This matters particularly where a regulator places weight on both short-window and long-window beta estimates. Short-window betas can be informative about movement over time. Long-window betas can reduce noise. But if the underlying process is not mean-stationary, recent level, historical range and long-run average are not interchangeable objects—they answer different questions.

3. What the tests show

3.1. I apply two standard diagnostics to asset-beta series (which were supplied by CEPA, with a data period of 01/01/2010-31/03/2026[1]) for the European comparator airports used by the CAA. (I have not assessed the suitability of those comparators and take them as given throughout this Report.) The Augmented Dickey–Fuller (ADF) test takes non-stationarity as the null, so failure to reject is evidence against a stable mean. The KPSS test works the other way round: stationarity is the null, so rejection is evidence against a stable mean. Read together, the two tests are more informative than either test on its own. A series that is plausibly drawn from a stable process should generally reject the ADF null and fail to reject the KPSS null. The results are shown in Table 1.

Table 1. Stationarity test results for the comparator asset-beta series

Horizon	AdP	AENA	Zurich	Fraport
2-year	Non-stationary	Non-stationary	Non-stationary	Non-stationary
5-year	Non-stationary	Mixed: rejects unit root and level-stationarity	Non-stationary	Non-stationary
10-year	Mixed: rejects unit root and level-stationarity	Non-stationary	Mixed: rejects unit root and level-stationarity	Non-stationary

3.2. The 2-year results are clear. For each of the four comparators, the ADF test does not reject non-stationarity and the KPSS test rejects level stationarity. On that evidence, the time series of 2-year betas should not be treated as fluctuating around a single stable mean.

3.3. The results for 5-year and 10-year series are not so stark but nevertheless show limited support for stationarity. Only the 5-year AENA, and 10-year AdP and Zurich betas show mixed results, with the ADF test rejecting a unit root, but the KPSS test still rejecting level stationarity. All other beta series show agreement between the ADF and KPSS tests. In one sense, it is not surprising that non-stationarity is (slightly) less evident with longer estimation windows, given

[1] I am aware that this may be a slightly later date than used by the CAA. The CAA does not state explicitly what end date is used in its analysis: it appears to be end October/beginning November 2025.

the smoothing that they involve. Nevertheless, in my view, even the 5-year and 10-year data provide very limited evidence of stationarity.

3.4. The ADF and KPSS results should be interpreted as diagnostic evidence rather than as formal hypothesis tests with standard critical values. The beta series used in this report are based on rolling OLS estimates, so adjacent estimates share a large proportion of the underlying daily return observations. This overlap mechanically induces persistence in the estimated beta series and means that the conventional ADF and KPSS critical values are unlikely to be directly applicable. In particular, the KPSS test may over-reject stationarity when applied to rolling beta estimates unless the critical values are adjusted for the overlap structure. This caveat does not remove the substantive point, but it does moderate the weight that should be placed on the tests in isolation. One important comparison is that equivalent tests applied to water-sector beta evidence do not point to the same instability: the estimated betas for Severn Trent and United Utilities pass both the ADF and KPSS tests, consistent with stable long-run means. The airport comparator beta evidence shows instability across windows and over time.

3.5. These findings are also consistent with the CAA's own description of the evidence. At paragraph 9.70 of its H8 Initial Proposals, the CAA states that "Equity and asset betas for the listed comparator airports have shown significant instability over the previous two price control periods." At paragraph 9.71 it adds that "Since the pandemic, equity and asset betas for comparator airports have continued to show significant instability." And at paragraph 9.78, the CAA attributes the "wide range" of comparator asset-beta estimates in part to "instability in the underlying estimates over time." That the comparator betas are unstable is, in short, largely common ground.

4. Consequences for the CAA's present analysis

4.1. There are several possible responses to non-stationarity. One approach would be to exclude particular periods or impose structural breaks. That can be appropriate where there is a clear economic break and a clear ex ante rule for identifying it. (A variation of this is the approach that the CAA took in H7: of attaching different weights to different periods of data.) In the present case, however, the data are highly volatile and affected by several overlapping shocks. An exclusion or structural-break approach would therefore introduce too many degrees of freedom and would risk replacing statistical judgement with selective data choice. That is not an attractive basis for regulatory estimation.

4.2. A second approach is to use more sophisticated time-varying beta models, such as Generalised Autoregressive Conditional Heteroscedasticity (GARCH) models (as discussed in UKRN (2018)), state-space models or Kalman-filter estimates. Such methods can be informative, especially where the objective is to model the beta process explicitly. These techniques have not yet been used extensively in regulatory settings.

4.3. The approach I adopt here is simpler: continue to use straight, unweighted OLS beta estimates, but to interpret them differently. The estimates should not be treated as observations drawn from a stable historical distribution. They should be treated as regime-sensitive estimates

that describe the empirical envelope of plausible beta values. The task is then not to identify a single historical mean, but to use the available OLS evidence, with judgement, to define a reasonable forward-looking range.

4.4. This approach implies several practical rules. First, particular weight should be given to recent spot estimates, because they are more likely to reflect the current risk environment. Where beta estimates are unstable over time, recent estimates are likely to be more informative about the current risk environment than older observations. In this environment, rolling averages can be used as diagnostics of instability, dispersion and regime dependence, but not as the main basis for the forward-looking beta estimate.

4.5. In line with this observation, the CAA uses spot estimates. How it does so is, however, problematic. As noted above, the time series of beta estimates show that comparator betas do not behave as if they are stationary around a stable mean. That matters for interpretation. If there is no stable underlying mean, then a long-window beta estimate cannot be treated simply as an equally reliable estimate of the same parameter. It is a different estimate, based on a different weighting of historical regimes.

4.6. The second practical rule is that, in a non-stationary setting, greater weight should normally be placed on shorter and medium-length estimation windows. A 2-year estimate is more responsive to the current regime, although it may be noisier. A 5-year estimate provides a compromise between relevance and reliability. A time series of 10-year estimates is smoother, and so a spot 10-year estimate effectively assumes that the beta process moves slowly. It gives substantial weight, however, to older regimes that may no longer be representative. It should therefore be used as a cross-check, not as the mechanical anchor for the top of the range.

4.7. This is where the CAA's analysis is least convincing. The CAA sets the lower bound of its pre-adjustment range by reference to the simple average of spot 2-year asset betas but sets the upper bound by reference to the weighted average of spot 10-year asset betas. A 10-year beta has the advantage of reducing short-term sampling noise. But it does so by averaging across a long period that includes materially different market and airport-risk conditions, including the pre-pandemic period, the pandemic shock, the reopening period, the inflation and interest-rate shock, and more recent market developments. Where beta is non-stationary, that smoothing property is not an unqualified advantage. It may reduce noise, but it may also dilute the relevance of the estimate for the current period.

4.8. Thirdly, the 5-year estimates should not be treated merely as a cross-check because they happen to fall inside the CAA's range. In this context they have independent evidential value. They are less exposed to short-term noise than 2-year estimates, but less dependent on stable historical regimes than 10-year estimates. If the objective is to estimate a forward-looking beta for the H8 period, the 5-year evidence should carry more interpretive weight than the CAA appears to give it.

4.9. Fourthly, the range should not be constructed by selecting endpoints from different estimation approaches unless there is a clear reason for doing so. The CAA's lower bound comes from one approach: a simple average of 2-year spot betas. Its upper bound comes from another:

a weighted average of 10-year spot betas. The end result combines different estimation windows and different averaging conventions without a clear principle explaining why those two estimates define the reasonable bounds. The result is a wide range, but not a well-justified one.

4.10. The stationarity results therefore have a specific consequence. They do not require OLS beta estimation to be abandoned. Nor do they require deletion of the pandemic period or the use of structural breaks. Rather, they mean that caution should be exercised in treating long-window estimates as more reliable simply because they use more data—in a non-stationary setting, more data are not always better data. Older observations may add precision around a value that is no longer relevant.

4.11. A more robust approach would be to use the full set of OLS spot estimates as evidence on the empirical envelope, but to apply explicit judgement to that evidence. Recent estimates should receive particular weight. The 2-year and 5-year estimates should be treated as especially relevant to the current regime. The 10-year estimates should be retained as a cross-check but should not alone determine the upper bound unless there is a clear explanation of why the older data remain representative.

4.12. The implication is not that the beta range should be narrow. The evidence remains volatile, and a degree of caution is appropriate. In addition, the CAA is right to consider recent spot estimates and right to recognise uncertainty in the beta evidence. Its present method does not, however, fully follow through on the implications of non-stationarity. Once beta estimates are shown not to be stable around a constant mean, the CAA should place greater weight on shorter and medium-length recent estimates; treat 10-year estimates as cross-checks; and avoid mechanical averaging across the range. That would provide a more transparent and economically grounded basis for setting the beta range, bounded by estimates that are most relevant to the current risk environment.

5. Consequences for the choice of beta

5.1. The CAA's pre-adjustment asset beta range of 0.47–0.70 is based on current spot estimates for listed airport comparators: the lower bound is the simple average of the 2-year daily asset betas, while the upper bound is the weighted average[2] of the 10-year daily asset betas. It notes that the corresponding 5-year values lie inside this range. It then uses the midpoint of this range for its point estimate of beta. Hence the CAA effectively places 50% weight on 10-year betas (for the upper bound); 50% weight on 2-year betas (for the lower bound); and no weight on 5-year betas.

5.2. The issue is not, therefore, that the CAA has ignored short-window evidence. The issue is that the upper end of the range is determined by a long-window estimate that is less informative when beta is not stable over time; and that the central estimate is then derived mechanically from the resulting range. This approach is inconsistent with the fact of non-stationarity.

[2] The CAA's case for using a combination of simple and weighted averages is not clear. In the following, I use simple averages.

5.3. The empirical evidence shows a clear term structure across estimation windows. The CAA's pooled comparator averages are 0.47 on a 2-year window, 0.60 on a 5-year window, and 0.67 on a 10-year window. The size, direction and persistence of the difference are consistent with the stationarity-test evidence and with the economic history of the sample.

5.4. This is the central methodological issue with the CAA's approach. The usual case for long estimation windows depends on a trade-off between relevance and reliability. Recent data may be more relevant; longer samples may be statistically more reliable. That trade-off is recognised in regulatory guidance and in the CAA's own discussion. But where the evidence suggests that beta has not had a stable mean, the trade-off changes. A long-window estimate may be more precise as an estimate of its own historical average; but it is a poorer guide to the forward-looking parameter required for a price control.

The upper bound

5.5. The 10-year estimate should not be discarded. It remains relevant evidence, especially because airport investment is long-lived and because short-window estimates can be noisy. But it should not be allowed to set the upper bound by itself. Giving decisive weight to the 10-year estimate would effectively assume that older observations remain as informative about H8 as recent observations. That is a strong assumption, and it is not supported by the evidence on instability in the beta series.

5.6. The shorter-window evidence points to a lower forward-looking range than the CAA's 0.47–0.70. A forward-looking range should reflect all three windows (2-year, 5-year and 10-year), but should attach less weight to the part of the 10-year estimate that reflects older regimes rather than current systematic risk.

5.7. On that basis, the CAA's upper bound of 0.70 is difficult to justify as a forward-looking estimate—it is too high to be treated as the top of the forward-looking range once the non-stationarity evidence is taken seriously. A more defensible upper bound is 0.64. That still gives some weight to the 10-year evidence, but it avoids treating the current 10-year estimate as the sole determinant of the upper bound (which is clearly inappropriate given non-stationarity).

The lower bound

5.8. The lower bound requires a different assessment. The CAA's 0.47 lower bound is close to the current 2-year pooled average and is therefore grounded in recent market data. The full time series of 2-year, 5-year and 10-year estimates suggests that values modestly below 0.47 are not outliers. In particular, readings around 0.45 are consistent with the lower part of the recent empirical envelope (it lies at around the 11th percentile of historic values). On balance, a lower bound of 0.47, while conservative, is defensible.

The central estimate

5.9. The central estimate should not be chosen by taking the midpoint of a range whose endpoints are generated by different estimation windows. A midpoint is a reasonable default

when the range is a symmetric uncertainty band around a stable parameter. That is not the case here. The lower and upper bounds are not two equally informative estimates of the same fixed beta. They are outputs from windows with different sensitivities to current and older data. Under non-stationarity, the central estimate should be a forecast for the regulatory period, not simply the arithmetic midpoint of some historical range.

5.10. A more defensible central estimate can be derived by computing a simple weighted average that gives greater weight to shorter windows: first 2-, and then 5-year estimates, while retaining some weight on the 10-year estimates. This produces results in the range 0.50–0.58, with an average and median of 0.53.[3] These figures can be compared with the midpoint of the recommended range 0.47–0.64, i.e. 0.56.

5.11. This approach is deliberately simple. Its purpose is not to create a new mechanical formula, but to discipline the judgement required under non-stationarity. It shows that a central estimate in the band 0.50–0.58, with a point estimate of 0.53, is supported by the evidence. In all cases, the appropriate central estimate sits below the CAA's midpoint of 0.585.

6. Summary

6.1. In summary, I recommend an asset beta range (before any adjustments for relative risk) of **0.47–0.64, with a central estimate of 0.53**. This recommendation gives weight to the CAA's recent 2-year evidence, retains the informational value of the 5-year and 10-year windows, and avoids treating the spot 10-year estimate as the sole determinant of the upper bound. The main divergences from the CAA are to the upper bound and to the method used to select the central estimate.

6.2. The key issue underpinning this recommendation is resolving an inconsistency in the CAA's approach. The CAA recognises that beta estimates are unstable—that is why it uses spot estimates. It then undermines this fact, however, by placing (effectively) 50% weight in 10-year beta estimates, which necessarily use data relating to periods when the mean beta was significantly different. To resolve this inconsistency, greater weights have to placed on the 2-year and 5-year beta estimates.

6.3. This approach is also consistent with the wider regulatory task. Where beta estimates are unstable, the regulator should avoid both special pleading for the deletion of inconvenient observations and mechanical reliance on estimates whose interpretation depends on stationarity. The better approach is to recognise the instability, use several window lengths, give greater weight to estimates that are more informative about the H8 period, and select a central

[3] This range is generated by taking a distribution of weights on the 2-year, 5-year and 10-year spot beta estimates, with the weights satisfying three criteria: (i) the 2-year weight is the largest of the three, but no larger than 0.80; (ii) the 5-year weight is the second largest; (iii) at least some weight is given to the 10-year spot. Note that the mean/median of the distribution of betas generated by this process is not the same as the midpoint of the range: the induced beta distribution is slightly skewed and is not simply uniform.

estimate that is robust to reasonable variations in weighting. On that basis, 0.53 is a more defensible central estimate than the CAA's 0.585.

6.4. Finally, this recommendation should be understood as applying to the comparator-derived pre-adjustment beta. It does not, by itself, resolve the subsequent questions of Heathrow-specific relative risk and the calibration of any adjustments. It provides the correct starting point for the further steps considered in Part II.

II. Heathrow's relative risk and the calibration of beta

7. Introduction

7.1. The CAA's H8 approach to beta starts from listed European airport comparators and ends with a Heathrow beta that is lower than the raw comparator beta because Heathrow is judged to bear less traffic risk than a standard airport, but more traffic risk than a network utility. The CAA's treatment of TRS makes this clear. It starts from a pre-adjustment beta and then asks how far Heathrow should be moved toward a utility benchmark, given the role of traffic in Heathrow's risk profile and the degree to which traffic-related risk is mitigated by the regulatory framework.

7.2. It recognises that there is no perfect listed comparator for Heathrow, while also recognising that Heathrow cannot sensibly be treated as equivalent to a network utility. Heathrow is plainly neither. It remains exposed to traffic and demand conditions in ways that utilities are not. At the same time, Heathrow differs materially from the European airport comparators. I am in full agreement with the CAA that an adjustment, reflecting Heathrow's specific features and circumstances, is both appropriate and proportionate.

7.3. Once the problem is framed in this way, the questions become relatively simple—at least in principle, even if calibration is not always straightforward. Which Heathrow-specific features make Heathrow less traffic-exposed than a standard airport comparator? Which Heathrow-specific features make Heathrow more protected by regulation than such a comparator? And how should those features be reflected in beta?

7.4. There is a good deal of agreement that the Traffic Risk Sharing (TRS) mechanism is a Heathrow-specific feature that affects both of these aspects. I argue that there are two other features that are material and which the CAA has not taken adequately into account: capacity constraints and till structure.

7.5. Capacity constraints matter because Heathrow is unusually slot constrained, protected by excess demand, and able to refill traffic. Those features reduce the pass-through from market-wide demand shocks to realised passenger volumes. That is a point about Heathrow's **underlying exposure** to traffic risk. Till structure matters because Heathrow operates under a single till, whereas the comparator airports do not. That difference does not mainly change Heathrow's physical exposure to passengers. It changes the extent to which the equity

consequences of traffic-driven commercial shocks are retained by Heathrow. That is a point about **mitigation** of traffic-related risk by regulation.

7.6. These two points are analytically distinct. Capacity constraints are an exposure issue. Single till is a mitigation issue. Some of the CAA's reasoning already reflects this distinction, at least implicitly; some of CEPA's work also reflects it. The advantage of making the distinction explicit is that it provides a clearer route through the competing arguments.

7.7. The rest of this part develops that framework and applies it to the calibration of Heathrow's asset beta. It first distinguishes exposure from mitigation, then applies that distinction to the CAA's risk drivers, capacity constraints and single till. It then draws out the implications for calibration. Technical appendices provide the algebra and derivations.

8. Conceptual framework

8.1. The CAA H8 reasoning (not always expressed in the clearest way) is well captured by the reduced-form expression implicit in Table 9.7 in its H8 Initial Proposals (Section 3):

$$\beta = \beta_0 - tr(\beta_0 - \beta_N).$$

Here, β_0 is Heathrow's beta inferred from airport comparators; β_N is a lower-beta network-utility benchmark; t is the extent to which Heathrow should be moved away from the raw airport comparator because Heathrow is not fully airport-like in its traffic exposure; and r is the extent to which Heathrow's traffic-related equity risk is mitigated by the regulatory framework, rather than borne by Heathrow investors.

8.2. This expression is called "reduced form" since it captures two key principles, without explicitly (i.e., structurally) modelling them. The first is that Heathrow's beta depends on its underlying exposure to traffic risk. That follows directly from the CAA's treatment of demand risk as a source of systematic risk, and from its repeated contrast between airport comparators and network utilities. Airports are more traffic-exposed; utilities less so; Heathrow sits somewhere between them. The second principle is that Heathrow's beta depends on the extent to which traffic-related risk is mitigated by regulation. Heathrow's beta is not determined solely by the airport's raw business model. It is also shaped by how much of traffic-related shocks is borne by Heathrow investors once the regulatory framework is taken into account. In summary, Heathrow's beta depends on both the traffic-related risk to which Heathrow is actually exposed and the proportion of that traffic-related risk that Heathrow investors bear once the regulatory framework is taken into account.

8.3. On this interpretation, t is best understood as the factor that captures Heathrow's effective exposure to traffic risk relative to the airport and utility benchmarks. A higher t means Heathrow should be moved further away from the raw airport comparator and further toward the utility benchmark because Heathrow is not fully airport-like in its traffic sensitivity. Similarly, r is best understood as the factor that captures the effective regulatory mitigation of Heathrow's traffic-related equity risk. A higher r means more of

Heathrow's traffic-related equity consequence is absorbed by the framework rather than borne by Heathrow investors.

8.4. While the CAA develops this formulation to address the effect of the TRS mechanism, it is more general than that and should be applied consistently. This interpretation matters because it provides a disciplined way of deciding what should affect each parameter.

- A feature should affect t if it changes Heathrow's **underlying exposure** to traffic shocks.
- A feature should affect r if it changes the degree to which Heathrow's **traffic-related equity risk is mitigated** by regulation.

This distinction is the organising principle of the rest of the note.

Table 2. Definitions of t and r

Parameter	Definition
t	The extent to which Heathrow should be treated as less traffic-exposed than a standard airport comparator, and therefore more like a network utility
r	The extent to which Heathrow's traffic-related equity risk is absorbed by the regulatory framework rather than borne by Heathrow investors

9. Heathrow and the European airport comparators

9.1. The CAA uses European airport (group) comparators to estimate beta (before any adjustments). That is understandable as a starting point. The listed comparator set is limited, and the comparator airports are plainly relevant as a first empirical reference for Heathrow. It does not follow, however, that the comparator betas can simply be read across to Heathrow. The comparator firms are not Heathrow. Their betas reflect the risks of different airport businesses in different settings.

9.2. There are four Heathrow–comparator differences that matter most in the present context, in addition to the TRS mechanism (which is broadly agreed to be a material difference).

Capacity constraints

9.3. The CAA accepts that capacity constraints and demand risk can matter in principle, but concludes that Heathrow is no longer sufficiently distinct from the comparator airports because some of them, especially certain Spanish hubs, are now also highly constrained. That is an evidential position rather than a conceptual one. The issue is not whether capacity constraints matter. The issue is how much they matter, given current comparator conditions.

9.4. Heathrow is unusually slot constrained and protected by excess demand. Heathrow's scarcity differs from that of the comparator airports not merely in degree, but also in character. Heathrow is formally co-ordinated because demand significantly exceeds capability; it operated at 99% of runway capacity in 2025; it continues to report that demand exceeds infrastructure limits; access rights are extremely valuable; and its role as Europe's busiest hub and the UK's

most valuable port by trade supports the inference that displaced demand can be refilled unusually readily.

9.5. The CAA's response is that some comparator airports have become more constrained, and that this narrows the Heathrow premium. Several features distinguish, however, Heathrow's capacity position from that of even the most constrained comparators.

9.6. First, Heathrow is formally co-ordinated at IATA Level 3, a structural classification: it is not a feature of a busy period but of an airport whose total declared capacity is insufficient to meet demand in normal conditions. No comparator airport operates under equivalent structural scarcity across both runway and terminal infrastructure simultaneously.

9.7. Secondly, the market value of Heathrow slot rights provides direct evidence of excess demand that cannot be manufactured by assertion. Secondary slot trades at Heathrow have historically cleared at prices in the hundreds of thousands of pounds per pair, reflecting the value investors attach to access rights that cannot be created by the airport or obtained except through the slot-allocation process. That kind of structural scarcity premium is not observed at the comparator airports to any comparable degree.

9.8. Thirdly, Heathrow's traffic recovery following the COVID-19 disruption was faster than the comparator average, consistent with a deep demand queue able to refill capacity relatively quickly once supply is restored. A market with persistent excess demand absorbs adverse shocks differently from one in which demand and supply are more closely matched. The CAA's analysis may imply that the gap between Heathrow and the comparator airports has narrowed; but it does not follow that the gap has fully closed, or that the remaining difference is too small to affect the calibration of the parameter *t*.

9.9. Further evidence provided to me by CEPA in their May 2026 Report reinforces the point that the CAA's comparison of Heathrow with comparator hubs is insufficiently granular. The key question is whether the airport faces a hard and persistent constraint at the margin that materially dampens exposure to demand shocks. CEPA's analysis distinguishes between runway and airfield constraints, planning and environmental constraints, peak-hour constraints, terminal constraints and operational constraints, each of which differs in how binding and how readily remediable it is. On the available runway movement evidence, Heathrow remains much more tightly constrained than the main comparator hubs cited by the CAA. This evidence supports the view that the CAA has not shown that comparator airports are constrained in the same economically relevant sense as Heathrow.

9.10. I return to this issue below.

Till structure

9.11. Heathrow operates under a single till; the comparator airports do not. This is one of the clearest and most economically relevant differences. It changes the way sustained commercial gains and losses are retained by equity and offset through future aeronautical charges. The CAA does not deny that till structure may matter in principle. Its objection is that CEPA's method in

its November 2025 Report of quantifying the effect is too sensitive to assumptions. I return to this issue below.

Business and group structure

9.12. Several comparator firms are airport groups rather than single-airport references. Their observed betas reflect a mixture of holdings, geographies and activities. The CAA spends time (see §9.79ff) considering whether those broader group structures warrant adjustment and concludes that they do not, at least not robustly enough. Even if one accepts that conclusion provisionally, then it remains true that a comparator-group beta is not identical to the beta of a Heathrow-like asset within the group.

Competitive and regulatory setting

9.13. The comparator airports operate under different combinations of regulation and competition. The CAA rightly notes that it is difficult to disentangle the two. Still, Heathrow's setting is not the same as Frankfurt's, Zurich's or the AENA group's. There is a wider package of regulatory protections applying to Heathrow. These include RAB-based recovery of efficient investment, the TRS mechanism, asymmetric risk allowances, inflation protection, cost pass-through or adjustment mechanisms for selected input costs, capex uncertainty mechanisms, early expansion cost recovery and the single till reset of commercial revenue forecasts. Taken together, these features are likely to reduce Heathrow's exposure to several risks faced by listed airport comparator groups. That matters because it means that differences in observed beta are not just differences in "airport risk"; they are differences in a broader package of airport economics and regulation.

Table 3. Heathrow–comparator differences that matter for beta

Heathrow-specific feature	Why it differs from comparator airports	Main relevance
Stronger slot constraints and excess demand	Realised traffic may be less responsive to market-wide demand shocks	Exposure
Single till	Commercial shocks are retained differently than under dual till	Mitigation
Single-airport reference rather than diversified group	Group betas may reflect a broader asset mix	Exposure / empirical comparability
Different competitive and regulatory context	Risk drivers do not map one-for-one	Exposure and mitigation

10. Exposure to traffic risk and mitigation of traffic-related risk

10.1. In my view, the distinction between exposure and mitigation is critical. A feature affects **exposure** if it changes the degree to which Heathrow's traffic, revenues or cash flows respond to traffic shocks before explicit regulatory mitigation is considered. A feature affects **mitigation** if it

changes how much of the equity consequence of traffic-related shocks is retained by Heathrow investors after the regulatory framework has done its work.

10.2. This distinction can be illustrated by the two main Heathrow-specific features considered here. **Capacity constraints** affect exposure: they determine how strongly market-wide demand shocks pass through into realised passenger numbers. Heathrow is more utility-like to the extent that such shocks are filtered by slot scarcity and excess demand. **Single till** affects mitigation. It determines how much of the commercial consequence of a traffic shock remains with Heathrow equity over time, rather than being offset through future aeronautical allowances.

10.3. That distinction is useful because it prevents too much from being loaded into one parameter. It also helps to show why some of the arguments have been talking past one another. A point about Heathrow's physical traffic sensitivity is not the same thing as a point about the treatment of commercial revenue reset. A point about commercial reset is not the same thing as a point about whether Heathrow faces the same traffic volatility as a standard airport comparator.

Table 4. Exposure versus mitigation

Question	Exposure	Mitigation
What is being changed?	Heathrow's underlying traffic sensitivity	Heathrow's retained equity sensitivity once regulation is taken into account
Typical Heathrow-specific drivers	Capacity constraints; some aspects of operational gearing	TRS; single till; other explicit sharing/reset features
Best associated parameter	t	r

11. The CAA's five risk drivers revisited

11.1. The CAA reviews five main risk drivers in its relative-risk discussion. Those drivers are useful, but they are more informative when sorted according to whether they affect exposure or mitigation.

Capacity constraints and demand risk

11.2. This is the clearest exposure factor. It is directly about how Heathrow's realised traffic responds to demand conditions. It therefore bears most heavily on t.

Capex and opex risk

11.3. The CAA is broadly right that simple capex and opex forecasting error is not obviously systematic in the CAPM sense. The more relevant notion is operational gearing. A fixed-cost business is more exposed to traffic shocks because a given fall in traffic translates into a larger fall in profits. That is still, however, primarily an exposure point rather than a mitigation point. It may influence t at the margin, but it is not a central parameter driver.

Till structure

11.4. This is the clearest mitigation factor. Till structure determines how much of the commercial consequence of shocks is retained. Since traffic shocks affect commercial revenues as well as aeronautical revenues, till structure affects the mitigation of traffic-related equity risk. It therefore affects r primarily.

Government support

11.5. Government support is difficult to use analytically here. The CAA is right to be cautious. Unless one can establish a clearer Heathrow-specific pattern of expected support relative to comparator airports, the factor is too ill-defined to play much role in calibration.

Regulatory framework

11.6. This category is too broad to do much work unless decomposed. Some elements of the framework may affect exposure indirectly; others may affect mitigation directly. The general category is less useful than the specific mechanisms, such as the TRS mechanism and single till, that sit within it.

Table 5. The five risk drivers mapped to the two factors

Risk driver	Mainly affects exposure?	Mainly affects mitigation?	Main parameter implication
Capacity constraints and demand risk	Yes	No	t
Capex and opex risk	Weakly	No	Weakly affects t
Till structure	No	Yes	r primarily
Government support	Unclear	Unclear	Probably neither
Regulatory framework	Sometimes	Sometimes	Case-specific

12. Capacity constraints and the determination of exposure

12.1. The strongest candidate for affecting Heathrow's exposure to traffic risk is Heathrow's slot scarcity and excess-demand position.

12.2. The economic intuition is straightforward. A market-wide demand shock only changes Heathrow's realised traffic to the extent that it passes through the constraint created by capacity. A standard airport comparator with greater slack capacity will, all else equal, allow more pass-through of demand shocks into realised passenger volumes. A more constrained Heathrow will not. To the extent that Heathrow can refill traffic or maintain throughput despite weaker demand conditions, Heathrow's realised traffic becomes less sensitive to the market than that of a more standard airport business.

12.3. This point has several implications. First, capacity constraints can reduce the traffic sensitivity of both aeronautical and commercial revenues. If fewer demand shocks translate into realised traffic changes, fewer of those shocks will reach the relevant revenue lines. Secondly, capacity constraints make Heathrow more utility-like. A network utility is, in the present context, shorthand for a business whose volumes and revenues are much less exposed to airport-style traffic volatility. Heathrow is not equivalent to such a utility, but it is more like one than the comparator airports to the extent that slot scarcity dampens traffic pass-through. Thirdly, this is exactly the sort of effect that should show up in a parameter that measures Heathrow's effective traffic exposure. It is not mainly about how regulation shares outcomes after the fact. It is about Heathrow's exposure before the sharing occurs.

12.4. The CAA's principal answer is that comparator airports have become more constrained too. While that is a legitimate point, it is not conclusive. The fact that some comparators are now more constrained than they once were narrows the Heathrow–comparator difference; it does not remove it. The question is one of degree.

12.5. A further difficulty with the CAA's present treatment is that it moves too quickly from the proposition that the comparator airports are also constrained to the conclusion that there is therefore no material Heathrow-specific effect. That does not follow. The more careful conclusion is that Heathrow's capacity position remains relevant to its exposure to traffic risk; the question becomes the extent to which the parameter t should be calibrated to reflect that fact.

Table 6. The capacity-constraint argument in steps

Step	Explanation
Heathrow is strongly slot constrained	Capacity limits the extent to which demand shocks become realised traffic shocks
Lower pass-through of demand shocks	Realised traffic is less volatile in response to market-wide conditions
Lower traffic sensitivity of cash flows	Heathrow's airport-type traffic exposure is reduced
Heathrow becomes more utility-like	Heathrow should be moved further away from the raw airport beta
Calibration implication	t should be higher than otherwise

11. Single till and the mitigation of traffic-related risk

11.1. Single till is a different kind of factor. It does not mainly change Heathrow's physical response to traffic shocks. It changes how much of the commercial consequence of those shocks is ultimately retained by Heathrow equity.

11.2. The key point is that a traffic shock affects Heathrow not only through aeronautical revenues but also through commercial revenues. More passengers mean more commercial spend; fewer passengers mean less. Under a dual-till benchmark, those effects are retained by

the airport. Under Heathrow's single-till framework, sustained differences in commercial performance feed through into future aeronautical allowances. The retained equity effect of the traffic-driven commercial shock is therefore lower—Heathrow investors retain less of the long-run equity consequence of traffic-driven commercial shocks than they would under a dual-till benchmark.

11.3. That is why single till should be treated primarily as a mitigation factor. It is also why single till should not be dismissed simply because CEPA's way of hard-wiring it directly into the comparator-derived beta is too sensitive (which seems to be the CAA's position). A standalone comparator adjustment is not the only, or indeed the best, way to reflect the single till aspect of Heathrow. Within the CAA's reduced-form approach, the cleaner route is to allow till structure to inform the calibration of mitigation i.e., the parameter r directly.

11.4. A separate CAA argument is that investors understand the single till and price it into their required returns, so that no further beta adjustment is needed. This argument, if accepted, would extend too far. By the same logic, investors understand the TRS and price that in too—yet the CAA uses the TRS as a reason to move r away from zero and reduce beta. The relevant question is not whether investors understand a feature of the regulatory framework but whether that feature changes the systematic risk they bear. The single till changes the systematic risk that investors bear by attenuating the retained equity consequence of traffic-driven commercial shocks. That is a direct effect on the beta of the equity claim, not a negotiable preference. Investor awareness of the mechanism does not neutralise it; it means that investors already reflect the lower systematic risk in the lower return they require, which is precisely what a higher r in the beta formula is intended to capture.

11.5. (There is a possible secondary effect of single till, through the exposure parameter t. If single till dampens non-traffic commercial risk to a greater degree than traffic-related commercial risk, then Heathrow's residual wedge over utilities becomes more traffic-dominated. This is likely to be second-order, however, compared to the more direct effect of single till on the mitigation parameter r: see Appendix C.)

Table 7. The single-till argument in steps

Step	Explanation
Traffic shocks affect commercial revenues	Passenger volumes influence spend-related revenues
Under single till, sustained commercial effects feed into future aeronautical allowances	Commercial gains and losses are not retained in full indefinitely
Heathrow equity retains less of the traffic-driven commercial shock	Lower long-run equity consequence of traffic shocks
This is not mainly a change in physical exposure	Heathrow remains passenger-dependent
Calibration implication	r should be higher than otherwise

12. CEPA's approach to single till

12.1. CEPA's central economic intuition in its November 2025 Report is broadly right. It is not correct to treat dual-till airports as if they are comparable to Heathrow—till structure matters because it affects the duration and retained value of commercial shocks. CEPA's simple reset-versus-no-reset illustrations capture that point well.

12.2. The problem arises when, in a worked example to illustrate the decomposition approach, CEPA tries to turn that intuition into a direct beta adjustment by decomposing comparator-group betas into narrower and narrower components and then recombining them under a hypothetical Heathrow-like single-till structure. This method is problematic for three reasons.

12.3. First, it relies on too many unobserved intermediate objects. One is no longer dealing with observed comparator betas but with inferred component betas for businesses and segments that are not separately traded.

12.4. Secondly, each decomposition step introduces new modelling choices: what the relevant segment weights should be; what proxy should be used for non-aeronautical beta; how to treat international assets; how to infer the aeronautical component; how to relate the dual-till and single-till treatment of commercial revenues. Each step may be arguable in isolation. Taken together, they create a long chain of contestable assumptions.

12.5. Thirdly, the end result is fragile by construction. Small changes in any one of the inputs can produce materially different final answers. The problem is not merely that some of the assumptions are debatable. The problem is structural: repeated de-averaging and re-averaging of unobserved components is bound to create sensitivity.[4]

12.6. Nevertheless, till structure does matter, even if CEPA's November 2025 route from plausible intuition to precise number is too sensitive to support a robust standalone adjustment. Below, I suggest an alternative, less sensitive way to include the effect of single till.

13. The CAA's criticisms of CEPA

13.1. The CAA's critique of CEPA is strongest where it points to the sensitivity of the final numbers to alternative plausible assumptions, as indicated above. The CAA is entitled to say that a direct comparator-beta adjustment is not robust if changing one or two reasonable inputs can radically alter the final result. That part of the response is sound. The CAA's response is much weaker, however, in what it takes from that criticism.

13.2. The first weakness is that the CAA moves too quickly from "CEPA's quantification is fragile" to "single till should play no role." That does not follow. The correct response to a fragile direct-adjustment method is to ask whether the same underlying economic effect can be handled more

[4] This arises simply as a matter of mathematics. Once CEPA uses its sum-of-the-parts approach to strip out international and commercial activities, the resulting beta inevitably has a very high sensitivity to initial parameter choices, with a coefficient proportional to $\frac{1}{(1-a)(1-b)}$, where a is the international airports share (31% in CEPA's illustrative calculation), and b is the commercial share (53%).

simply elsewhere. The reduced-form calibration through the parameter r provides exactly such a route.

13.3. The second weakness is that the CAA relies too heavily on broad arguments about regulatory discretion, investor preference for commitment, and the absence of hindsight. These arguments may not be wholly irrelevant, but they are not specific to single till. They are generic points about regulation and review structure. They do not amount to a strong answer to the more specific proposition that Heathrow's single-till treatment dampens the retained equity consequences of traffic-driven commercial shocks.

13.4. The third weakness is that the CAA often treats the issue as binary. Either CEPA has a robust standalone single-till adjustment, or till structure gets no effective role. That is too stark. Once the CAA has already chosen to think about Heathrow's beta in reduced-form terms, there is no reason to force every Heathrow–comparator difference into a standalone adjustment to the comparator-derived beta. Some of those differences—especially capacity constraints and till structure—sit more naturally in the calibration of the reduced-form parameters.

13.5. The fourth weakness is that the CAA applies an inconsistent evidentiary standard. Its own calibration of *r* at 0.5 for TRS involves exactly the same kind of judgement it criticises in CEPA's single-till work: namely, the inference of a quantitative parameter value from economic reasoning rather than a direct statistical estimate. The 0.5 figure for TRS is not formulaically derived from e.g., a regression: it reflects a considered view about the degree of risk sharing that the mechanism achieves. If the CAA accepts that route for TRS, then it cannot object to the same route for single till solely on grounds of imprecision. The legitimate criticism of CEPA is that its direct comparator-beta adjustment is fragile, not that judgement-based calibration through the parameter *r* is inherently unreliable. A well-grounded single-till increment through *r* should be held to the same standard as the TRS calibration the CAA has already accepted—and by that standard, it passes.

Table 8. The CAA's response to CEPA: strengths and weaknesses

Element of response	Assessment
Critique of CEPA's assumption sensitivity	Strong
Observation that CEPA's numbers can be implausibly sensitive	Strong
Positive alternative route for recognising till structure	Weak
Reliance on generic arguments about regulatory discretion	Weak
Conclusion that till structure should play little or no role	Not well supported
Inconsistent evidentiary standard (judgement accepted for TRS, rejected for single till)	Weak

14. Implications for calibration

14.1. The CAA has effectively set t in the range 0.7–0.9 and r at 0.5 on the basis of the TRS alone. That is, in effect, a calibration based on the CAA's view of the role of traffic and the extent to

which TRS mitigates it. This creates an obvious question. If the CAA has already accepted that Heathrow's beta depends both on traffic exposure and on regulatory mitigation of traffic-related risk, should the calibration not also reflect the two other Heathrow-specific features that bear most directly on those factors: capacity constraints and single till? My answer is yes.

14.2. In answering in this way, I recognise that the choices of both t and r are, to some extent, subject to regulatory judgement: they are not parameters that can be read directly off the data. Neither are they, however, free policy choices. The choices still have to be disciplined by economic structure, internal consistency, and plausibility bounds. I have argued that capacity constraints and single till operate in a clear direction, to increase t and r respectively. I will also argue below that both effects are material enough that zero is not the natural default; indeed, that plausible estimates can be arrived at.

Capacity constraints

14.3. The CAA's 0.7–0.9 range already reflects the view that traffic is an important determinant of Heathrow's airport–utility wedge. The remaining question is where within that range the best judgement lies. Heathrow's scarcity position still points to a value above the midpoint: Heathrow remains exceptionally tightly constrained and more refillable than the comparator set taken as a whole. I start by recognising that Heathrow is no longer uniquely constrained in Europe, but it still appears more chronically and systemically constrained than the average comparator.

14.4. The CAA's own evidence indicates that Barcelona is operating at about 100% of capacity and Madrid at about 95%, while Zurich retains some spare capacity overall, though with peak constraints. Heathrow, however, remains exceptionally tight: Heathrow states that its two runways are the most intensively used in the world, that it operates at or close to runway capacity for long periods of the day, and that terminal capacity is also heavily constraining growth; it handled 83.9 million passengers in 2024 (see Heathrow Airport). Frankfurt appears materially less constrained, having handled only 61.6 million passengers in 2024, 87.3% of 2019 traffic, while opening a new terminal in 2026 with capacity for up to 19 million passengers per year (see Fraport AG).

14.5. This evidence is consistent with the traffic volatility figures presented by CEPA in its November 2025 Report, which I reproduce below. Its figures indicate that Heathrow is similar to the core airports of the comparator groups, but different from the non-core airports, with the latter showing higher levels of traffic volatility.

Table 2.7: Heathrow and airport group traffic volatility measured as the std. dev. of changes in passenger traffic volumes (2013-2024)[19]

Core / Non-core	Heathrow	AdP	Fraport	AENA[20]
Non-core airports	2.3%	14.4%	10.9%	24.1%
Core airports[21]		2.8%	2.2%	6.1%
Simple average[22]	2.3%	8.7%	10.8%	16.8%
Weighted average[23]	2.3%	13.2%	14.6%	6.6%

Source: CEPA analysis of airport published traffic statistics.

14.6. Nevertheless, the comparator airports are now themselves more capacity-constrained than they were in earlier periods. Any further adjustment for Heathrow's lower traffic-risk exposure should therefore be modest. A central estimate of 0.85 is reasonable. This sits slightly above the midpoint of the CAA's 70–90 per cent traffic-risk assumption. That midpoint already recognises that traffic risk accounts for most of the difference between airport comparators and network utilities. The issue is whether Heathrow's particularly acute capacity constraints, slot scarcity and ability to refill lost demand justify placing it somewhat above that midpoint. In my view, they do, but not by enough to support a value at the top of the range.

14.7. Several considerations support this judgement. First, the comparator set includes Frankfurt, which was materially under-capacity in 2024 and is actively adding infrastructure; Zurich, which retains meaningful off-peak slack; and other airports whose constraint levels are materially below Heathrow's. A weighted-average constraint level across the comparator set lies well below Heathrow's. Secondly, as argued earlier, the structural character of Heathrow's constraint—IATA Level 3 co-ordination, significant secondary slot premiums, chronic excess demand—is qualitatively different from the peak-period tightness observed at some comparators. Thirdly, a value of 0.85 implies only a modest increment over the CAA's midpoint of 0.80. It is consistent with accepting that the comparator gap has narrowed while maintaining that it has not closed.

14.8. I do not suggest that Heathrow should be placed at the absolute top of any conceivable range simply because it is capacity constrained. I do suggest that the capacity-constraint argument points away from the lower end of the current range and towards the upper part of it. The choice of 0.85 reflects that.

Single till

14.9. Single till should point to a value of r above 0.5. If 0.5 reflects the CAA's estimate of the mitigation delivered by the TRS alone, then it is incomplete once Heathrow's single-till treatment is recognised. The single till does not eliminate traffic-related commercial risk; it does, however, attenuate it. That should increase the degree of effective mitigation of traffic-related risk.

14.10. A value of r above 0.5 is supported by the scale of Heathrow's commercial revenues and by the magnitude of the single-till offset in the price control. In Heathrow's 2024 accounts, retail

revenue was £772 million and "other revenue" was £481 million, against total revenue of £3,482 million. On that basis, retail alone accounted for about 22% of total revenue, and a broader commercial/non-aeronautical measure accounted for about 36%. Those are not marginal revenue streams. Nor is the regulatory reset trivial in size. In the CAA's H7 Final Proposals, commercial revenues including ORCs deducted from the gross revenue requirement were £852 million in 2022, rising to £1,122 million by 2026, or £5.1 billion in total over H7, equivalent to about 37% of the gross revenue requirement across the period. The CAA's H8 mid-case has commercial revenues including ORCs of £6.7 billion against a gross revenue requirement of £19.4 billion, or about 34.5% of gross revenue.

14.11. A simple way to calibrate the single-till increment is to start from the CAA's TRS calibration and then apply a single conservative haircut. TRS applies to roughly two-thirds of Heathrow's revenues but only shares 50% of traffic risk. Single till applies to roughly one-third of Heathrow's revenues; and, for sustained gains and losses, operates much closer to a full reset at the next review. On a raw revenue basis, the two mechanisms are therefore of similar order: $(2/3) \times 50\% \approx (1/3) \times 100\%$. Single till is, however, weaker than that raw comparison suggests, because the reset is lagged and because some of the relevant risk overlaps with traffic-related risk. A simple and conservative approach is therefore to apply a 50% haircut to the raw equivalence result. That implies an increment of about 0.25, meaning that r should be set at 0.75.

14.12. Three qualifications to the haircut deserve brief note. First, the overlap between TRS and single till on traffic-related commercial risk is limited in practice. TRS operates primarily on aeronautical revenues; single till resets commercial revenues at the price review. The revenue bases are largely distinct, so the raw equivalence comparison is not significantly contaminated by double-counting, and the 50% haircut is already conservative in this respect.

14.13. Secondly, single till operates symmetrically: it resets both commercial over-performance and under-performance through future aeronautical allowances. It therefore functions as a genuine two-way mitigation mechanism, attenuating the consequence of traffic-driven commercial shocks in both directions, not merely capping upside. This reinforces the view that the mechanism should be reflected in a higher r rather than ignored.

14.14. Thirdly, the formal model in Appendix B, evaluated at a Heathrow-type discount rate and a standard five-year review period, implies a substantially larger single-till discount than the 0.25 increment assumed here. With a cost of capital of approximately 6.5% and a review period of five years, the present-value retention factor φ is approximately 0.28, meaning that Heathrow retains only about 28% of the present value of a permanent commercial shock under single till compared with a dual-till benchmark. The 0.75 figure therefore represents a conservative estimate: it applies a haircut well in excess of what the formal model alone would imply, reflecting appropriate caution about the speed and completeness of the reset mechanism across different types of commercial shock.

Table 9. Practical calibration implications

Parameter	CAA current calibration	Effect of Heathrow-specific features	Recommended central estimate
t	0.7–0.9	Heathrow capacity constraints reduce traffic pass-through	0.85
r	0.5	Single till increases effective mitigation of traffic-related equity risk	0.75

14.15. One objection to this recommendation is that it makes the framework too discretionary. That is true only up to a point. Any reduced-form framework necessarily requires judgement. Indeed, the CAA's own approach to the TRS is largely judgement-driven; and in previous price control decisions, it has made adjustments to comparator airport asset betas to reflect a regulatory judgement that Heathrow had lower demand risk (see §9.94). The issue is whether the judgement is being applied to the right things. In my view, capacity constraints and single till are not peripheral candidates for such judgement: they are two of the most important Heathrow-specific differences from the comparator set.

15. Overall effect

15.1. The consequences of bringing these arguments together are shown in Table 10. My recommendations result in a central estimate of Heathrow's asset beta of 0.42 (compared to the CAA's 0.50).

Table 10. The effect of capacity constraints and single till on asset beta

	CAA Low	CAA High	CAA Mid	**Recommended central estimate**
Pre-adjustment asset beta	0.47	0.70	0.585	0.53
Network utilities beta	0.36	0.36	0.36	0.36
Exposure %	0.90	0.70	0.80	0.85
Mitigation %	0.50	0.50	0.50	0.75
Adjusted asset beta	0.42	0.58	**0.50**	**0.42**

15.2. Figure 1 decomposes the difference between CAA's midpoint asset beta and the one recommended in this Report. (The "Interactions" column reflects that the three factors—of a lower pre-adjustment beta, and higher exposure and mitigation parameters—have a combined effect as well as separate ones.)

Figure 1. The difference between CAA's point asset beta and the recommended central estimate

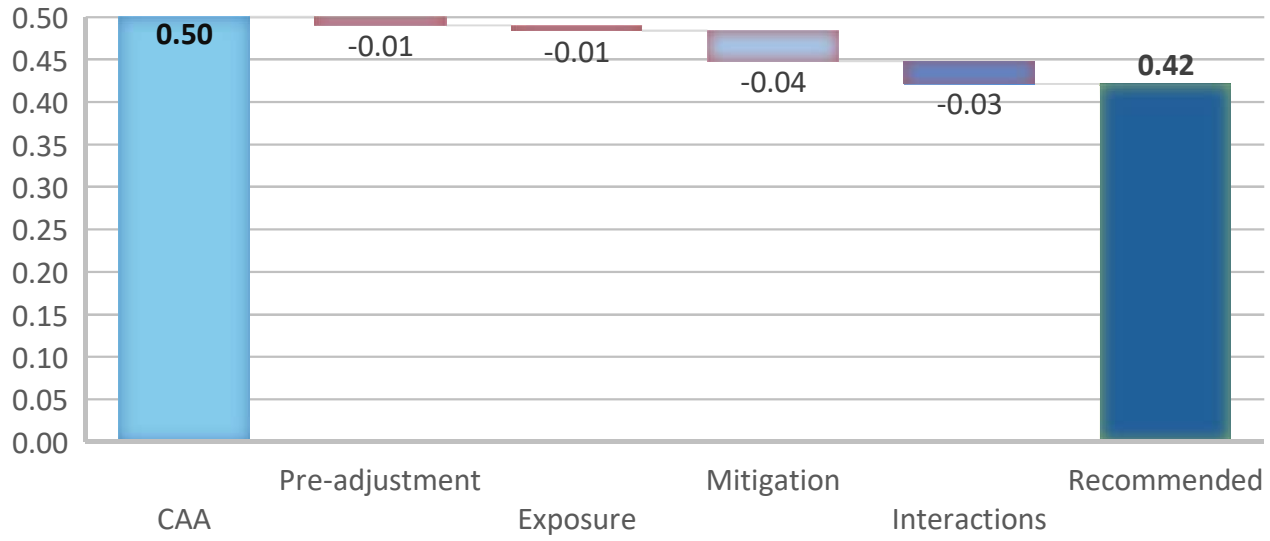

16. Conclusions

16.1. The CAA's H8 beta analysis has a sensible starting point. Heathrow is not a listed company, so listed airport comparators are the natural empirical reference. Heathrow is also not identical to those comparators, so some movement toward a lower-risk benchmark is appropriate. The issue is not whether the CAA's framework is directionally plausible. The issue is whether the range, central estimate and Heathrow-specific adjustments have been calibrated consistently and completely in line with the evidence.

16.2. On the pre-adjustment beta, the evidence does not support a mechanical midpoint derived from a range whose endpoints come from different estimation windows and different averaging conventions. The comparator betas are unstable, and the stationarity tests provide little support for treating long-run averages as estimates of a single stable mean. The CAA recognises this fact but has been inconsistent in then placing a 50% weight on 10-year beta estimates. These estimates remain relevant, but they should operate as cross-checks rather than as the determinant of the upper bound. A pre-adjustment range of 0.47–0.64, with a central estimate of 0.53, gives appropriate weight to recent evidence while retaining the informational value of longer windows.

16.3. On relative risk, the CAA is right that Heathrow should sit between the airport comparator benchmark and a network-utility benchmark. But that judgement should be applied to the right economic mechanisms. Capacity constraints reduce Heathrow's underlying traffic exposure. Single till mitigates the retained equity effect of traffic-driven commercial shocks. These are not peripheral features. They are central to the difference between Heathrow and the listed European airport comparators. The CAA has, in my view, been incomplete in its consideration of these issues.

16.4. The recommended calibration is therefore straightforward. Heathrow's capacity constraints justify setting the exposure parameter modestly above the CAA midpoint, at 0.85. Heathrow's

single-till framework justifies setting the mitigation parameter above the TRS-only value, at 0.75. These estimates are deliberately conservative. They do not assume that Heathrow is a network utility, and they do not assume that comparator airports are unconstrained. They simply recognise that Heathrow remains more capacity-protected and more regulatorily insulated than the comparator beta would imply.

16.5. The combined effect is to reduce the central adjusted asset beta from the CAA's 0.50 to 0.42. This gives a central estimate of beta that is, in my view, materially better than the CAA's. It follows from the CAA's own logic once that logic is applied consistently and completely: instability in the comparator beta evidence should reduce reliance on the 10-year upper-bound estimate; and Heathrow-specific exposure and mitigation features should be reflected in the adjustment from airport comparators to Heathrow.

16.6. The final recommendation is therefore that the H8 asset beta should be calibrated from a pre-adjustment range of **0.47–0.64**, with a central pre-adjustment estimate of **0.53**; and that the relative-risk adjustment should use central parameter values of **0.85** for exposure and **0.75** for mitigation. This gives a central adjusted asset beta of **0.42**. It is a more coherent, and therefore materially better, reading of the evidence than the CAA midpoint, while remaining anchored in the same comparator set and the same broad regulatory framework.

Appendix A. Formal treatment of traffic-risk exposure

This appendix sets out a simple formal representation of the distinction between Heathrow's exposure to traffic risk and mitigation of traffic-related risk by regulation.

Let unconstrained passenger demand be D and effective capacity be K. Realised traffic is given by the minimum of the two:

$$N = min(D, K)$$

Assume unconstrained demand varies with a market shock m according to:

$$D = \bar{D}(1 + am + \varepsilon)$$

where a measures the sensitivity of unconstrained demand to the market, and ε is an idiosyncratic shock.

If capacity is slack, realised traffic moves with demand. If capacity binds, only a fraction of the market-driven demand shock passes through into realised traffic. A convenient reduced-form approximation is:

$$n = \kappa am + \varepsilon^*$$

where:

- n is the proportional realised traffic shock;
- $\kappa \in [0,1]$ is the pass-through parameter;
- ε^* is a residual idiosyncratic term.

The lower is κ, the less sensitive realised traffic is to market-wide demand conditions.

Suppose aeronautical revenues are proportional to passenger numbers and commercial revenues contain a traffic-related component also proportional to passenger numbers. Then the traffic-related components of both revenue streams are proportional to κam. A lower κ therefore reduces Heathrow's structural traffic sensitivity.

In a structural model, this would reduce Heathrow's beta directly. In the CAA's reduced-form framework, the same effect is represented by moving Heathrow further away from the raw airport-comparator beta and closer to the utility benchmark. That is why stronger capacity constraints imply a higher value of the exposure parameter.

A simple mapping is:

$$t = \bar{t} + \gamma(1 - \kappa)$$

where $\gamma > 0$. Lower pass-through, represented by lower κ, increases t.

This formalisation does not determine the exact numerical value of t. It clarifies only the sign and the logic. The calibration question remains empirical and judgmental: how much lower is Heathrow's pass-through than that implied by the raw comparator beta?

Appendix B. Formal treatment of mitigation of traffic-related risk

This appendix formalises the single-till argument.

Suppose a traffic shock affects Heathrow equity through two channels:

1. an aeronautical effect;
2. a commercial effect.

Let the total equity effect of a traffic shock under a dual-till benchmark be:

$$dV^{DT} = dV_A + dV_C$$

where dV_A is the aeronautical component and dV_C the commercial component.

Now consider the present value of the commercial component. Under a dual-till benchmark, suppose a one-unit traffic-driven commercial gain is expected to persist indefinitely. Then its present value is

$$PV^D = \frac{1}{\rho}$$

where ρ is the relevant discount rate.

Under a single till, suppose that gain is retained only until the next reset after T years, after which stronger commercial forecasts are offset through lower aeronautical allowances. Then the present value is

$$PV^S = \int_0^T e^{-\rho u}\, du = \frac{1 - e^{-\rho T}}{\rho}.$$

Define:

$$\phi = 1 - e^{-\rho T}$$

with $0 < \phi < 1$.

Then

$$PV^S = \phi PV^D$$

and the retained equity effect of the traffic-driven commercial shock under single till is only a fraction ϕ of the dual-till value.

Accordingly, under single till:

$$dV^{ST} = dV_A + \phi dV_C.$$

The reduction in retained traffic-related equity sensitivity is therefore

$$dV^{DT} - dV^{ST} = (1 - \phi)dV_C.$$

This gives the formal basis for saying that single till increases the mitigation of traffic-related equity risk. The key point is that the commercial component of a traffic shock is still present, but it is not retained in full indefinitely.

A simple reduced-form mapping is

$$r = r^{TRS} + \delta(1 - \phi).$$

with $\delta > 0$, subject to $r \leq 1$.
This says that the effective mitigation of traffic-related equity risk is greater than the mitigation supplied by TRS alone once the single-till treatment of commercial shocks is recognised.
Again, this does not determine the exact value of r. It clarifies only the direction and logic of the adjustment.

Appendix C. Secondary route from single till to exposure

There is also a weaker secondary route from single till to the exposure parameter.
Suppose Heathrow's wedge over the utility benchmark can be decomposed into a traffic-related component T and a non-traffic component U:

$$\beta_0 - \beta_N = T + U.$$

Then a natural measure of the traffic-related share is:

$$t = \frac{T}{T + U}.$$

Now suppose single till disproportionately reduces the non-traffic commercial component by a factor $\phi < 1$. Then the residual wedge becomes:

$$T + \phi U$$

and the corresponding traffic share becomes:

$$t^{ST} = \frac{T}{T + \phi U}.$$

Since $\phi < 1$, it follows that:

$$t^{ST} > \frac{T}{T + U}.$$

This is the formal version of the secondary argument that single till can make Heathrow's residual wedge over utilities more traffic-dominated. The argument is coherent. It is not the most natural primary route, because it is more abstract than the direct mitigation story. That is why the main text treats single till primarily as a determinant of mitigation rather than exposure.

Appendix D. Brief biography

I am currently a decision-making member of the Competition Appeal Tribunal. I also served on the Competition Commission and subsequently the Competition and Markets Authority for eight years. I have served as a member of the Financial Conduct Authority's Regulatory Decisions Committee and Competition Decisions Committee, and of the Payment Systems Regulator's Enforcement Decisions Committee and Competition Decisions Committee. I have advised both Ofwat and Ofgem on cost of capital matters. With David Miles and Stephen Wright, I was co-author of the 2003 Report on cost of capital matters commissioned by the U.K. economic regulators and the Office of Fair Trading. With Stephen Wright, Phil Burns and Derry Pickford, I was a co-author of the 2018 update of that work, commissioned jointly by the CAA, Ofcom, Ofgem and the Utility Regulator.