APPENDIX E.ii.b – WACC – Views on cost of equity

Views on cost of equity

In this Appendix, we discuss a number of topics from Section 2.2 in detail:

1. TRS adjustment
2. Understanding drivers of airport risk
3. Comparing airport and utility risk
4. Comparator relative risk
5. Total market return
6. Illiquidity and beta estimation

1. TRS adjustment

The CAA has included an adjustment to asset beta due to the Traffic Risk Sharing (TRS) mechanism in the H8 Initial Proposals. In this section we set out our approach to estimating the TRS adjustment in H8, in the following steps:

- How the TRS adjustment works
- Counterfactual definition
- Characterisation of the risk that TRS mitigates
- Choice of outturn measure
- Model of TRS adjustment
- Model estimation

How the TRS mechanism works

The CAA determines a passenger forecast at each price control review, typically every five years.

Starting in H7, the CAA introduced the Traffic Risk Sharing mechanism:

- At the end of the calendar year, passenger outturn is compared with the CAA's passenger forecast, and a forecast error in percentage terms is calculated
- The percentage of risk sharing is calculated:
    - If the percentage error is within a deadband of +/- 10% of the forecast, then there is 50% risk sharing. For forecast errors greater than +/- 10%, there is 105% risk sharing for the additional error outside the deadband.
- The quantum of TRS adjustment is the percentage risk sharing (positive or negative) multiplied by the aero charge in that year (the MAY), multiplied by the passenger forecast error
- Recovery of the TRS adjustment begins in calendar Year t+2 (the next available year after the aero charge is reset)

Counterfactual definition

In this analysis, the counterfactual is the scenario in which the TRS mechanism does not exist, all else is equal. The factual is the scenario in which the TRS mechanism is in place, in the form included in H8 Initial Proposals. The counterfactual and factual are in a 2R scenario.

Characterisation of the risk that TRS mitigates

Risk describes the random variability of a particular observable outturn. Risk is characterised from the perspective of a particular economic agent, and at a particular time horizon. For example, *revenue risk at a one year horizon from the perspective of shareholders* would describe the variability in the next year of total revenue that shareholders cannot predict, given their information set at the start of the year. For TRS adjustment estimation, risk should be assessed from the perspective of shareholders (over varying time horizons).

For a regulatory mechanism to mitigate a particular sort of risk, it must:

- Constrain an observable outturn;
- The constraint must be contractual. If the contract is incomplete and there remains a discretionary element, the insurance is incomplete. Residual risk remains that the variable will outturn outside of the constraint

The TRS mechanism can only mitigate risk of the outturns that it affects, on the time horizon that the outturn is affected. The TRS mechanism affects:

- Revenues from the start of year t+2 until the end of the price control period;
- The aero revenue in the upcoming price control period derived from the adjustment of the opening RAB of the following regulatory period, fully depreciated by t+11.

The TRS mechanism does not constrain total cost recovery nor aero revenue at any point. The net effect of the TRS mechanism therefore depends on whether and how the CAA and airlines incorporate TRS information into their discretionary decisions that affect cost recovery.

It is our understanding that in a 2R scenario, the CAA sets the capex envelope on the merits of the capex business cases provided during the price control review, independently of TRS adjustments. In a 3R context, the CAA is currently taking the level of the aero charge into consideration in determining the regulatory framework and the 3R capex envelope: *"HAL's aeronautical charges per passenger are high relative to average adjusted charges for comparator airports. The gap in charges between HAL and comparators is also likely to grow as a result of HAL's planned additional capital expenditure on expansion. This supports our focus on ensuring that HAL's charges are no higher than necessary"* (CAP3251, para 1.12).

It is our understanding that airlines take the current level of the aero charge heavily into consideration when determining their views about an appropriate capex envelope. Airlines have repeatedly said so themselves, citing the overall level of the charge as a key determinant of their views on the appropriate size of the capital envelope:

Capital Envelope Threshold

Summary
HAL's proposed £10bn capital envelope has been rejected by airlines as neither deliverable, warranted nor in keeping to address the overall level of charge. The high-level nature of the submission and absence of a clear 'golden thread' linking investment to measurable outputs make assessment extremely difficult ~ **this requires a reset.**

Figure E.ii.b.1 – Extract of Airline Community Alternative H8 Business Plan. Source: Alternative H8 Business Plan

Airlines have also frequently cited that the current aero charge level influences their views on the appropriate expansion capital envelope.[1] Increments/decrements in revenue due to TRS may well be offset by other regulatory decisions, and lead to negligible net change to the overall charge level.

Model of TRS adjustment

The purpose of this section is to more formally state what the TRS mechanism is from the perspective of equity and to estimate the appropriate adjustment to asset beta.

In this analysis, we have attempted to keep as closely to CAPM as possible. In particular, we have designed this approach to strengthen the evidence about risk components that are systematic.[2]

Target CAPM term

In this section we set out the CAPM term of interest.

In CAPM, the only risk that affects asset return is a systematic single factor:

$$\beta_i = \frac{Cov(R_i, R_m)}{Var(R_m)}$$

Where β_i is asset beta, R_i is the asset return and R_m is the market return.

We want to estimate beta under the TRS mechanism, assuming it has some (zero or non-zero) incremental impact R_{TRS} on asset return:

[1] For example, see Heathrow Reimagined press releases.

[2] This was particularly in view of the CMA's discussion about the requirement for superior evidence about systematic vs. idiosyncratic risk: "We acknowledge that the CAA has estimated the 50% reduction in volume risk from the TRS using its estimates of total cash-flow protection from the TRS. We accept that not all traffic shocks will necessarily be systematic, and as such, this assumption is inevitably an approximation. However, it was a judgement the CAA was entitled to make based on its reasonable estimates of the relevant protection in the absence of other clearly superior evidence." Ibid, para 6.244.

$$\beta_{i,TRS} = \frac{Cov(R_i + R_{TRS}, R_m)}{Var(R_m)} = \frac{Cov(R_i, R_m) + Cov(R_{TRS}, R_m)}{Var(R_m)}$$

$$\beta_{i,TRS} = \beta_i + \frac{Cov(R_{TRS}, R_m)}{Var(R_m)}$$

The last term is the term of interest: the co-movement between (1) the incremental return from TRS in the factual and (2) market return.

Model estimation

We examine the key term, the covariance between TRS returns and market returns. TRS returns are positive to Heathrow when passenger outturn is below expectations, and vice versa. Another way of saying this is that the covariance of TRS returns with market returns has the opposite sign of the covariance of passenger forecast errors with market returns:

$$\frac{Cov(R_{TRS}, R_m)}{Var(R_m)} > 0 \text{ iff } \frac{Cov(\tilde{p}, R_m)}{Var(R_m)} < 0$$

Where $\tilde{p}$ is a passenger forecast error.

We have extensively examined covariance between passenger outturn and market returns across Heathrow and the comparators. Across the entire range of specifications that we investigated, the relationship is economically insignificant, indicated by a consistently slightly negative sign (passenger forecast is too high when market returns are relatively low):

Airport	Time period of observations	Coefficient
Aena	Feb 2015 - Mar 2026	-0.00898
Fraport	Jan 2016 - Mar 2026	-0.00961
Aeroport de Paris	Jan 2015 - Feb 2026	-0.01548
Zurich	Jan 2015 - Mar 2026	-0.01312
Heathrow	Jan 2015 - Mar 2026	-0.01382

Figure E.ii.b.2 – Table showing regressions of stock returns on % passenger change (monthly data)

We have investigated many other combinations of passenger outturn measures and equity return and beta measures, documented in Appendix E.ii.c: Passengers and market return. These results consistently show the same relationship.

To illustrate why this is the case, we can examine Stoxx 600 returns against Heathrow forecast errors:

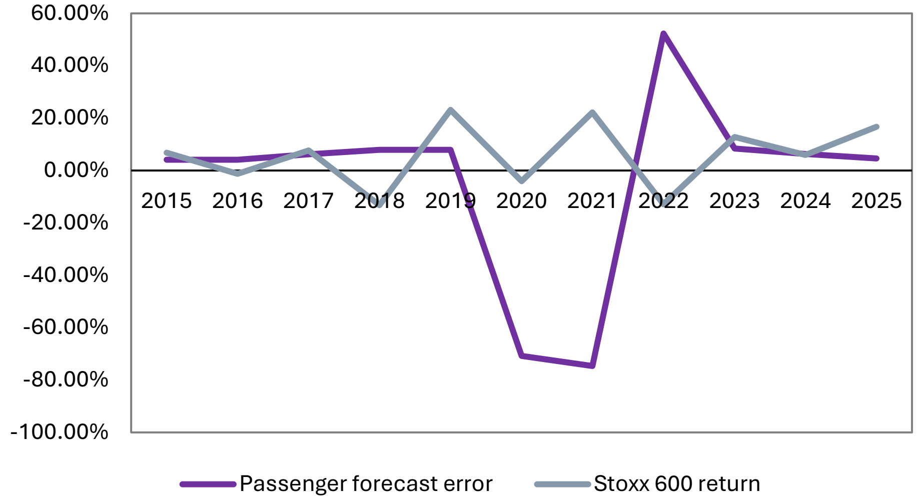

Figure E.ii.b.3 – Graph showing CAA passenger forecast error and Eurostoxx 600 market return (correlation = -0.3)

- In non-pandemic years, as market returns fluctuated, these did not systematically translate to changes in Heathrow forecast errors
- 2020-2022 had large passenger forecast errors:
  - 2020 saw both negative passenger forecast error and near-zero market returns (pandemic market crash and subsequent recovery)
  - 2021 saw a strong market recovery and the market effects of large fiscal stimulus, while the aviation market remained weak
  - 2022 saw stronger-than-expected passenger recovery, combined with a market crash related to the monetary policy reaction to the inflationary effects of fiscal stimulus

The aviation market responds to many types of macro shocks within the span of weeks to a few months,[3] whilst passengers book tickets up to six months in advance. This significantly weakens the co-movement between market expectations and passenger outturn.

However, shareholders cannot be certain about the value of TRS prior to the realisation of the passenger forecast error from which is estimated. It is difficult to be certain about the value of TRS prior to its revenue payout. This is because TRS is one of a suite of regulatory risk sharing mechanisms, and the CAA and other stakeholders have the discretion to adjust the suite on an ongoing basis depending on the value of each particular component. For example, during the pandemic, other policies and regulations--such as alleviation from slot requirements as well as government fiscal support to businesses and airlines—were very important for re-allocating risk in the commercial aviation market, but were also challenging to predict in advance.

[3] Airline industry equities under external uncertainty shocks

Based on the above, there is no evidence that the TRS adjustment affects systematic risk to Heathrow within a CAPM framework.

Analysis of TRS as an insurance policy

Setting aside the question of whether TRS mitigates systematic risk, we have analysed TRS as an insurance policy against idiosyncratic risk. We have estimated the quantum of TRS adjustment in the scenario where the TRS mechanism existed historically, against the scenario where it did not exist. This results in the below compensation profile:

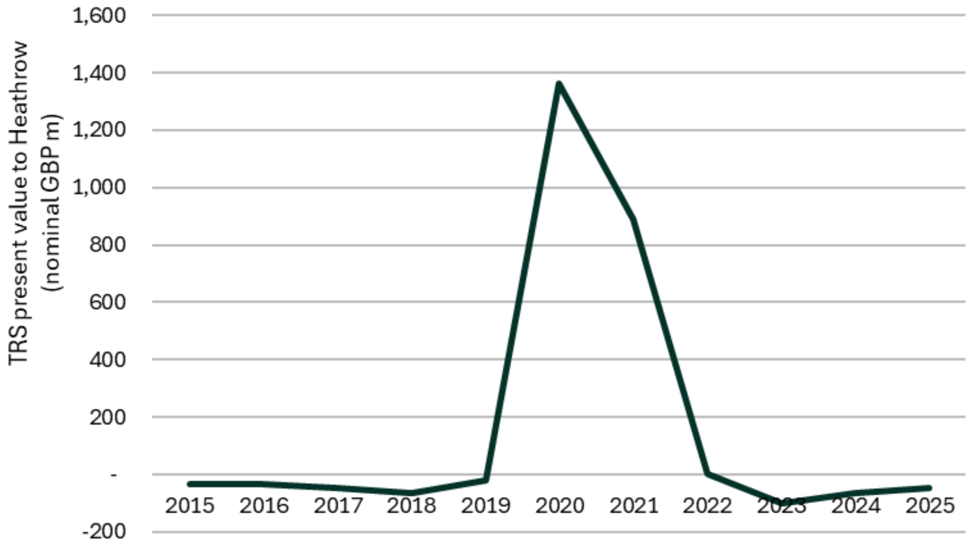

Figure E.ii.b.4 – Graph showing historical TRS value (2015 – 2025). Source: Heathrow analysis

This analysis is provided in Appendix E.ii.c, and it is the basis of the Return on Asset analysis presented in Section 2.2.

2. Understanding drivers of airport risk

In the following section we discuss contextual issues to help understand the results from Section 1:

- Drivers of total revenue risk at airports
- Drivers of systematic revenue risk at airports
- Why the TRS mechanism does not mitigate systematic risk
- TRS idiosyncratic airport risk mitigation
- Operational gearing and risk

Drivers of total demand risk at airports

Credit rating agency frameworks provide an independent and reasonably comprehensive overview of airport risks. The 2024 Fitch ratings guidance for transport infrastructure provides a useful overview of volume and price risks. This framework is shown below.

Assessment of Price Risk

Description Price Risk	Stronger	Midrange	Weaker
Legal or contractual price-setting flexibility	No or minimal contractual or statutory price caps, or ability to increase rates materially above inflation	Contractual or statutory price caps substantially tracking inflation	Contractual or statutory price caps track substantially less than inflation
Political pressure	Evidence of minimal legislative or political interference	Evidence of some legislative or political interference	Evidence of significant legislative or political interference
Long-term contracted revenues [a]	Take-or-pay contracts, long-term leases representing a substantial portion of revenues with long minimum average lives or to the end of the concession	Take-or-pay contracts, long-term leases representing a moderate portion of revenues with intermediate term average lives	Take-or-pay contracts, long-term leases representing a small portion of revenues with short-term average lives

[a] Typically, the third subfactor would not be applicable to either toll roads or US municipal airports.

Figure E.ii.b.5 – Table showing Fitch credit rating assessment criteria for revenue risk

Airports – Guidance for Volume Risk Sub-Factor Assessments

Airports	Volume Risk – Sub-Factor – Description
1. Reference market	✓Considered Stronger if the airport operates within a large, economically strong, diverse and mature reference market. ✓It would be Midrange if one or more of these attributes may be weaker – such as a smaller, more concentrated or less economically strong reference market. ✓It would be considered Weaker if the airport operates in a relatively small or economically weak/volatile reference market.
2. Strategic importance	✓Considered Stronger if the airport has a key role in the broader aviation system with a strategic or preferred location with fast, direct and quality road and mass transport links connecting the reference market to the airport[a]. ✓It would be Midrange if the airport is medium-sized, plays an important role in the broader aviation system and has good location and road and mass transport links connecting the reference market to the airport[a]. ✓It would be considered Weaker if the airport is small, has a secondary role in the aviation system or is outside the reference market or is difficult to access[a].
3. Diversification	✓Considered Stronger if the airport exhibits most of the following factors: is exposed to a single carrier concentration of 30% or less, offers extensive nonstop and international services and has a balanced mix of business and leisure traffic. For the assessment of issuers with multiple assets, stronger would apply to an issuer with a portfolio of geographically diversified assets with limited economic correlation and that is predominantly located in countries with well-established regulatory frameworks. ✓It would be Midrange if it exhibits most of the following factors: is exposed to a single carrier concentration of 30% to 60%, offers broad services and has high levels of leisure traffic to support volumes. For issuers with multiple assets, midrange would apply to an issuer with a small portfolio of assets with limited geographical diversification or some economic correlation or some assets in countries with less developed regulatory frameworks. ✓It would be considered Weaker if the airport exhibits most of the following factors: is exposed to a single carrier concentration greater than 60% or offers limited services or caters for predominantly leisure traffic. For issuers with multiple assets, 'Weaker' would apply to an issuer with a small portfolio of assets with high geographical concentration and/or high economic correlation and/or most of the assets located in countries with less developed regulatory frameworks.
4. Competition	✓Considered Stronger if the airport benefits from minimal competition from other airports or alternative modes of transport. ✓It would be Midrange if the airport was exposed to some competition from other airports or alternative modes of transport. ✓It would be considered Weaker if the airport is exposed to meaningful competition from other airports or alternative modes of transport.
5. Demand volatility	✓Considered Stronger if the airport exhibits all or most of the following: a) Demand is stable and resilient demonstrating historical volume and low volatility; b) Has connecting traffic of up to 20% for domestic airports and higher for international gateways; c) High barriers to entry[b] in the issuer's reference market; and d) Low risk that structural changes will affect issuer volumes in the long term also support the stronger assessment. ✓It would be Midrange if the airport exhibits all or most of the following: e) Demand exhibited moderate historical and prospective volatility; f) The airport has connecting traffic of between 20% and 60% or supporting a primary connecting operation, or a major carrier base of operations; g) Moderate barriers to entry[b] in the issuer's reference market; and h) Moderate risk that structural changes will affect issuer volumes in the long term also support the 'Midrange' assessment. ✓It would be considered Weaker if the airport exhibits all or most of the following: i) Demand is highly volatile, both historically and prospectively; j) Over 60% of traffic is connecting traffic; k) Low barriers to entry[b] in the issuer's reference market; and l) High risk that structural changes will affect issuer volumes in the long term also support the weaker assessment.
6. Relative cost to end users	✓Considered Stronger if the airport's aeronautical and retail yield per passenger is higher or the cost per enplanement is lower than competition and peers. ✓Considered Midrange if the aeronautical and retail yield per passenger and cost per enplanement is average relative to competition and peers. ✓It would be considered Weaker if the airports' aeronautical and retail yield per passenger is lower or the cost per enplanement is considered high compared to competition and peers.

[a] Size is a secondary indicator of essentiality.
[b] To determine whether the barriers to entry assessment is high, moderate or low, the influences Fitch considers will include: a) access to market for new entrants based on regulatory or market surroundings, b) scarcity of alternative locations/physical assets, c) planning/permitting processes, and d) significance of required capital-related costs.
Note: Indicative high level considerations are not just based on mechanical counting of the sub-factor assessments, as the final overall volume assessment considers the weighting of each sub-factor which can very between transactions.

Figure E.ii.b.6 – Table showing airport guidance for volume risk sub-factor assessments. Source: Fitch Ratings

At the highest level, we can simplify these risks into:

1. **Airline price risk**: this largely driven by regulation, the bargaining power of airlines relative to airports, and passengers' elasticity of price demand. Heathrow has relatively little ability to flexibly raise prices in the short-term, due to regulation.
2. **Airline volume risk**: this is overwhelmingly driven by airline profitability risk. The more that the airport is situated in a location with robust demand, with barriers to entry for airline customers, the less profitability risk the airlines have, and the less volume risk the airport has. A few drivers modulate the extent to which the airport is exposed to airline profitability risk: diversification (does the airport have other streams of revenue), is the customer base diversified (no single airline can

disproportionately influence airport volume), and do other airports offer better value (will airlines leave in a downturn). Although Heathrow's capacity constraint may reduce upside volume risk, there are counteracting forces that increase revenue risk: lack of revenue diversification, high airline concentration, and high proportion of premium passengers whose demand is relatively highly pro-cyclical.

From this we can hypothesise the key drivers of variation in airport revenue risk as:

1. In the medium- to long-term: airport exposure to and the intensity of the profitability risk of their airline customers
2. In the short-term: ability to quickly raise airport charges

Explaining variation in airport beta

If it is the case that airport beta is driven by expectations about long-run airline profitability, then we should see this pattern reflected in the relationship between airport and airline betas.

Among airlines, we see that there is not a very stable ordering among equity betas over time.

Figure E.ii.b.7 – Graph showing airline equity betas (2-year)

We find a relatively stable, positive, and material relationship between the equity betas of airports and their respective airline customers. We find that lower mutual concentration (the airline is a small fraction of the airport's traffic and vice-versa) dilutes the relationship. We would interpret the coefficients with relatively low coefficients in the below figure to represent a 'baseline' aviation sector co-movement, with airport-airline pairs that have high mutual concentration exhibiting higher co-movement.

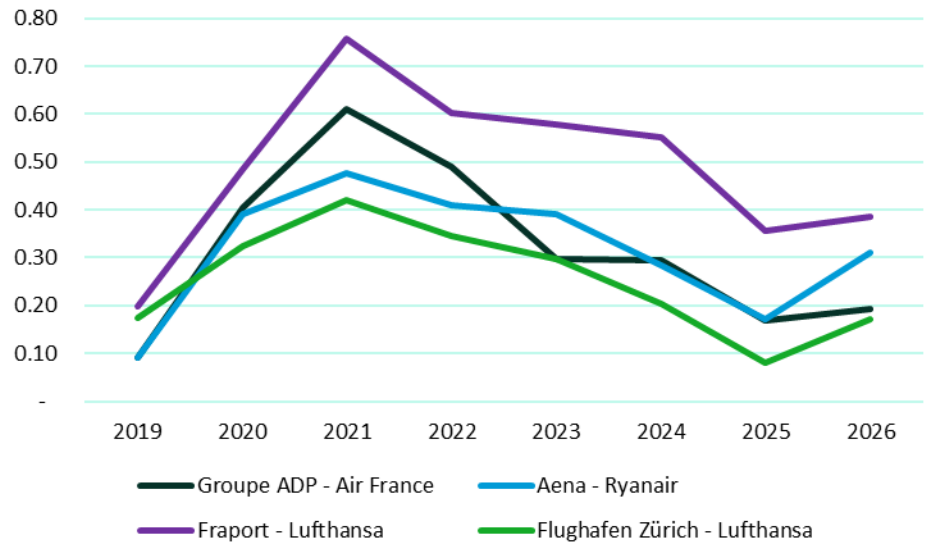

Figure E.ii.b.8 – Graph showing contribution of airline beta to airport beta (2-year)

The full results of this exercise are included in Appendix E.ii.c.

We find that there is reasonable empirical evidence showing that there is variation in airport asset beta that can be explained by exposure to airline betas.

The TRS adjustment does not mitigate systematic risk

As discussed above, airport systematic revenue risk can principally be reduced by:

- Mitigating price risk: improving expectations about the airport's ability to raise prices in response short-run events
- Mitigating volume risk: improving expectations about the long-run profitability of its main airline customers
- Shareholders' expectations must be updated in a timely way, to reduce co-movement with market-wide expectations.

The TRS adjustment does not improve expectations about long-run airline profitability. It does not reduce the sort of price risk that is described in the Fitch extract above (Figure E.ii.b.5 & E.ii.b.6), because TRS does not affect Heathrow's ability to raise prices in the very short-term. It makes logical sense that the TRS adjustment would not affect systematic risk.

3. Comparing airport and utility risk

In their analysis of the effect of TRS on WACC[4], CEPA make a series of key arguments, to which we respond below.

"H7 saw a change in the demand risk faced by Heathrow, with the introduction of the TRS mechanism."

There is no evidence that the TRS mechanism affects demand for airport services at Heathrow; therefore the TRS cannot not affect demand risk. The TRS mechanism affects cost recovery through the aeronautical charge. The distinction between demand risk and revenue risk matters, because the relationship between demand and commercial performance varies by comparator, as Heathrow's revenues are more sensitive to fluctuations in passenger volumes than the comparators' revenues (see Section 2.2).

"Where investors bear more non-diversifiable risk, they will require a higher return (and vice versa)."

We agree that the above statement reflects the CAPM framework, however CEPA subsequently do not incorporate diversify-ability of risk in their calculation. This is a critical omission in their logic. Heathrow has relatively little diversification of revenue due to its focus on core airport activities that are directly tied to passenger volumes (see 'Comparator relative risk' section of this appendix), and as a result has materially higher risk for shareholders than the comparators in absence of TRS (see Return on Asset calculation).

"The approach adopted by the CAA shows that risk allocation is not a given and the regulator should consider the allocation of risk that gives the optimal outcomes. We recommend that the CAA undertakes an assessment of potential combinations of risk allocations and revenues (in particular the cost of capital) to test that the H7 approach continues to be optimal for H8."

'Optimal' regulation delivers the best outcomes possible for consumers. The objective of the cost of capital determination is to efficiently incentivise the investment needed to deliver consumer value. The CAA have previously determined in H7 that the relevant benchmark for an efficient cost of equity (as captured by asset beta) is other European hub airports. CEPA's approach to TRS adjustment does not analyse the comparators at all. It does not tell us anything about how to apply to Heathrow an asset beta that is calculated from comparator data.

"We recognise that you cannot 'solve' for the right answer, but running illustrative financial scenarios should be useful. ....It should also be noted that there may be upside

[4] CAP3044A Virgin Atlantic consultation response, *Approach*.

and downside risks faced and the distribution of returns needs to be considered around whether this is a 'fair bet.'"

CEPA's 'illustrative' approach treats Traffic Risk Sharing as an insurance policy. As CEPA notes, the value of an insurance policy to a beneficiary depends on when the policy pays out. A policy that has positive value when the insured outcome is high will increase total risk, whereas a policy with an identical expected value, but which has positive value when the insured outcome is low will decrease risk. CEPA's analysis considers the expected return of a policy, and does not consider the covariance of this return with any other commercial or financial outcome. Their analysis therefore tells us nothing about the effect of TRS on risk.

"If investors require compensation in excess of the maximum revenue impact, it is unlikely to be a suitable risk allocation for them to bear the risk. … It is unclear to us that the risk faced by Heathrow justifies the compensation awarded for bearing partial demand risk under the TRS mechanism (or from bearing full risk in the absence of a TRS mechanism). Having a review of risk allocation would be a useful exercise to allow stakeholders to discuss optimal outcomes."

In order to assess the effect of TRS on risk allocation, and the appropriateness of the risk allocation, we have analysed the risk-return profile of Heathrow against the comparator airports over the last decade (see Section 2.2). Heathrow without TRS has had a worse risk-return profile from the perspective of investors than the average of the comparators. TRS brings Heathrow closer in line with the average of the comparators. We believe that our analysis directly addresses CEPA's request.

Capex risk

CEPA's underlying assumption is that Heathrow with no volume risk would have the same asset beta as a utility. This assumption omits other differences between Heathrow's business model and a utility business model that lead to different risk profiles. These include but are not limited to: risk due to the bargaining power of customers (key airline customers have relatively high bargaining power), tax regime risk (such as impact of changes in tax free shopping and business rates), and capex risk (including risk in capex approval, delivery, and cost recovery allowance).

On capex risk, given the scale and complexity of the expansion programme, it is important to carefully define and understand how large capital programmes affect systematic risk.

The CAA considers that the primary channel through which capex risk may influence systematic risk is operational gearing. It argues that, while variations in capex and opex relative to forecasts can materially affect cashflow volatility, such risks are diversifiable, on the basis that there is no reason to expect cost overruns at Heathrow to be more likely during periods of weak market performance. We consider this to be an oversimplification.

We agree that operational gearing is an important driver of systematic risk and expect it to be an important differentiator of R3 relative to H8 and comparator airports. However, we consider that additional factors are also relevant.

In particular, many drivers of capex and opex outturns - such as supply chain constraints, materials and labour shortages, and input price volatility - are linked to broader macroeconomic conditions affecting the UK economy as a whole. These factors are not firm-specific and therefore cannot be assumed to be fully diversifiable.

Moreover, some of these drivers are procyclical. For example, periods of strong economic activity tend to increase demand for materials and labour, placing upward pressure on costs. While the regulatory regime provides some protection, it does not insulate firms from these macroeconomic influences.

Supply chain risk is similarly not diversifiable given it is driven by economy-wide factors. Supply chain risk across infrastructure is complex and is driven not only by demand conditions, but also by structural factors and exogenous shocks.

The scale and complexity of an investment programme can amplify these exposures.

Taken together, these factors suggest that a number of capex risk drivers beyond operational gearing are not diversifiable. Consistent with this, FTI considers "capex and opex risk to be systematic as it amplifies the airports' exposure to systematic shocks… As this exposure is driven by economy-wide factors rather than firm-specific events, it is non-diversifiable and therefore contributes to beta."[5]

4. Comparator relative risk

In this section we evidence our conclusions that Heathrow bears greater revenue and regulatory risk than the comparators.

This analysis builds on the KPMG's relative risk assessments included in our BP submission, which was largely a qualitative analysis, and here we seek to quantify measures of relative risk.

Where we discuss demand risk, we examine revenue risk. This is to compare airport groups on a like-for-like basis. The different airport groups have very different mixes of activity, and comparing passenger volumes or other activity measures omits airport group business segments; whereas revenue incorporates all activities.

The main findings are:

[5] FTI Report - H8 Cost of Capital - Final.docx , para A5.5.

- **Revenue diversification**: Heathrow's focus--on core airport activities and its operation of a single hub airport--exposes it to greater systematic and idiosyncratic risk than the comparators, which operate more diversified businesses.
- **Regulatory framework**: HAL's commercial model has greater exposure to regulatory risk than comparators'

Revenue diversification

The most important approach for an airport to reduce revenue risk is diversification. An airport group can make high risk investments that decrease its total and systemic risk if those investments increase its portfolio diversification. During the pandemic, the airport sector highlighted that diversification away from core aeronautical activities increased commercial and financial resilience.[6] [7] [8] We examine the following dimensions of diversification:

- Diversification within country
- Diversification internationally

Diversification within country

HAL's revenues are more focused on its aeronautical business streams than the AdP, Fraport, and Zurich. Specifically, HAL earned, on average, 60% of its domestic revenues from aeronautical charges from 2019 to 2024, whereas the comparators earned 53%, with greater reliance on other revenue streams.

Operator	Aeronautical revenue as a proportion of domestic revenues (%)	Non-aeronautical property developments
HAL	60%	-
AENA	55%	-
AdP	60%	ADP operates **Hub One**, a digital technology subsidiary providing broadband connectivity, business software and cybersecurity services to companies and public organisations, with revenues of €168m in 2024 (approximately 4% of its Parisian revenues).[9] Hub One operates in competitive commercial markets entirely

[6] diversifying-airport-revenue-through-non-aeronautical-land-development.pdf
[7] Maximizing Non-Aeronautical Revenues: Key to Airport Financial Sustainability | ACI World Insights
[8] Non-aeronautical revenues: Diversify and grow – Airport World
[9] Calculated from AdP Annual Accounts 2024.

Operator	Aeronautical revenue as a proportion of domestic revenues (%)	Non-aeronautical property developments
		disconnected from aviation demand.[10]
Fraport	46%	Fraport's Retail and Real Estate segment encompasses the commercial development of the broader Frankfurt airport site, including office buildings, logistics real estate and commercial premises beyond the terminal perimeter. These include sites such as **Gateway Gardens**[11] and **Cargo City**[12].
Zurich	50%	Zurich Airport owns 51% of **The Circle**, a major mixed-use development immediately adjacent to the terminal complex, home to more than 50 companies and over 5,000 employees and encompassing office space, two hotels, medical facilities, retail and hospitality[13]. The Circle represents a deliberate strategy of converting land value on the airport estate into recurring income streams independent of passenger throughput.

Figure E.ii.b.9 – Table showing summary of differences in revenue activities (2019-2024)

In addition to its high reliance on aeronautical revenue, Heathrow earns its non-aeronautical revenues primarily through concessions. Under a concession, the airport receives a percentage of the tenant operator's turnover, typically subject to a minimum guaranteed floor. By contrast, Fraport explicitly notes in its accounts that, while its retail leases are predominantly variable (turnover-linked), it also holds a substantial portfolio of out-of-terminal real estate, comprising offices, logistics parks and commercial premises, on conventional fixed-rent terms, which produces income that is stable and independent of passenger volumes.[14] Zurich similarly structures its commercial leases as a guaranteed base rent with a turnover-linked top-up, making its

[10] AdP Annual Accounts 2024, page 51.
[11] See: https://www.gateway-gardens.de/en/businesses/.
[12] See: press release.
[13] Zurich Integrated Report 2024, page 24. See also: https://www.thecircle.ch/en/home.
[14] Fraport Annual Accounts 2024, page 216.

income more weighted towards fixed receipts than HAL's based model.[15] AENA's commercial income incorporates a distinct Minimum Annual Guaranteed (MAG) rent component alongside its variable, turnover-linked rents, providing a more stable floor than HAL's minimum guarantee arrangements.[16]

While it does not have substantial non-aeronautical property developments, AENA operates a total of 46 airports and two heliports within Spain alone, meaning that revenues are not tied to demand within one specific hub airport.[17]

Diversification internationally

HAL's revenues exclusively come from its single hub airport, London Heathrow. It has no international concessions, equity interests in foreign airports, or management contract income from overseas operations.

By contrast, the comparators have sizeable revenues from international operations, accounting for 21% of their revenue, on average, over 2019-2024.[18]

Operator	Non-domestic operations as a proportion of total revenues (%)	Non-domestic operations
HAL	0%	-
AENA	9%	UK, Brazil, Colombia, Jamaica and Mexico
AdP	25%	Turkey, Georgia, Kazakhstan, Tunisia, Latvia, North Macedonia, Saudi Arabia, Jordan, India, Indonesia and more
Fraport	41%	Slovenia, Brazil, Peru, Greece, Bulgaria, Türkiye
Zurich	9%	Brazil, Chile, Curaçao, India

Figure E.ii.b.10 – Table showing comparators' revenues from international operations

AENA. AENA's International segment, managed through Aena Desarrollo Internacional (ADI), generated 14% of total consolidated group revenue in 2025[19]. The segment comprises three principal operating subsidiaries:

- London Luton Airport, in which AENA holds a controlling stake.
- Aeroportos do Nordeste do Brasil (ANB), which operates six airports in northeast Brazil under a concession running to 2047.

[15] Zurich Annual Accounts 2024, page 162.
[16] AENA Annual Accounts 2025, page 198, 203.
[17] AENA Annual Accounts 2024, page 269.
[18] Comparator Annual Accounts.
[19] AENA Annual Accounts 2025.

- Bloco de Onze Aeroportos do Brasil (BOAB), which operates eleven airports across four Brazilian states.

Beyond these consolidated subsidiaries, AENA holds interests in airports in Colombia, Jamaica and Mexico[20].

ADP. ADP operates an extensive international portfolio, representing 30% of total consolidated group revenue in 2024[21].

- TAV Airports, in which ADP holds a majority stake and which it fully consolidates, operates 14 airports across Turkey, Georgia, Kazakhstan, Tunisia, Latvia, North Macedonia, Saudi Arabia and other markets.[22]
- ADP International (AIG), a wholly owned subsidiary, operates Amman's Queen Alia International Airport in Jordan.[23]
- GMR Airports, an Indian airport group in which ADP holds a 45.7% economic interest (accounted for on an equity basis), operates four airports: New Delhi Indira Gandhi, Hyderabad Rajiv Gandhi, Goa Manohar, and Medan in Indonesia.[24]

Fraport. Fraport operates on four continents through concession agreements, with its International Activities and Services segment covering 43% of total group revenue in 2024[25]. Fraport's international assets are[26]:

- Ljubljana airport in Slovenia, of which it has a 100% ownership.
- Two airports in Brazil Porto Alegre and Fortaleza of which it has a 100% ownership.
- Lima Airport in Peru (80% owned, concession to 2051).
- 14 regional airports in Greece (65% owned, concession to 2057).
- Two airports in Bulgaria (60% owned, concession to 2046).
- Antalya Airport in Türkiye (with its contract due to renewal in 2026).

Zurich Airport. Zurich Airport operates an international division through Zurich Airport International Ltd., a wholly owned subsidiary responsible for the analysis, development and operational management of overseas concessions and advisory mandates, representing 9.9% of total group revenue in 2024[27]. Its portfolio comprises concessions and subsidiaries in[28]:

- Brazil (Florianópolis, 100% owned; Vitória, 100% owned; Macaé, 100% owned; São Gonçalo do Amarante, 100% owned; Belo Horizonte, 12.75% stake).

---

[20] AENA Annual Accounts 2025, page 189.
[21] Calculated from AdP Annual Accounts.
[22] See AdP Universal Registration Document 2024, page 18.
[23] See AdP Universal Registration Document 2024, page 47.
[24] See AdP Universal Registration Document 2024, page 18.
[25] Fraport Annual Report 2024.
[26] Fraport Annual Report 2024, page 44.
[27] Zurich Integrated Report 2024.
[28] Zurich Integrated Report 2024, page 17.

- Chile (Iquique and Antofagasta, 100% owned).
- Curaçao (9.69% stake).
- India (Noida International Airport, 100%-owned and was scheduled to commence operations in 2025, but in practice has not yet started its commercial operations[29]).

All of the comparators have seen growth in their international operations since 2019, except for Zurich which has not quite recovered to the shares of international revenues seen before Covid (though this is likely to change with the completion of Noida International Airport).

Effect on Heathrow's risk profile

We have analysed the correlation between group-level revenues and airport passenger volumes at each Comparator Operator's flagship airport:

- London Heathrow (LHR) operated by HAL.
- Madrid (MAD) operated by AENA.
- Charles de Gaulle (CDG) operated by AdP.
- Frankfurt (FRA) operated by Fraport.
- Zurich (ZRH) operated by Zurich airport.

Economic literature on risk and diversification

There is a limited economic literature about the effects of diversification into adjacent sectors. It suggests that modest levels of diversification into related business areas reduce systematic risk.

Lubatkin and Chatterjee (1994)[30] proposed that an important way for corporations to minimise risk is to diversify into similar businesses rather than into very similar or very different businesses. Specifically:

- Single business firms are "*tied to the environmental uncertainties of a single industry and …show high sensitivity to business specific (unsystematic) and macro economic (systematic) disturbances in returns*".
- Firms that undertake a limited or 'constrained' diversification programme benefit from: (i) economies of scope where similar technologies can be used across business streams even if there is a cyclical downturn in one business stream; (ii) the presence of a common "industrial logic" in the different business streams can empower management to adapt to problems as they arise in different subsidiaries; and (iii) a long term competitive which makes them more able to withstand downturns than their competitors.

[29] See: https://www.niairport.in/en/company/news/2026/2026-03-06.
[30] See: Extending Modern Portfolio Theory into the Domain of Corporate Diversification: Does it Apply? | Academy of Management Journal.

- However, highly diversified firms (conglomerates) may be unable to exploit the benefits enjoyed by the constrained diversifiers and end up behaving more like a collection of isolated single business firms.

In this context, Lubatkin and O'Neill (1987)[31] found that mergers of firms in related sectors are associated with a significant decline in systematic and total risk (although mergers tend to increase idiosyncratic risk).

Heathrow's regulatory framework

A comparatively small proportion of the comparators' business activities are regulated. By contrast, Heathrow's revenues have higher exposure to regulatory discretion.

Between 2019 and 2024, on average 40% of the comparators revenues were subject to regulation. At Heathrow almost all of its revenues are regulated through the single till system. This restricts Heathrow's ability to adapt flexibly to changing commercial conditions by adjusting prices.

Airport Operator	Average Proportion of Regulated Revenues (2019-2024)
HAL	99%
AENA	53%
AdP	43%
Fraport	19%
Zurich	46%

Figure E.ii.b.11 – Table showing proportion of revenues falling under regulation for each operator

If we are willing to make the assumption that the CAA does not differ quantifiably from other airport regulators in terms of the risk it allocates to equity from the regulated business, then--purely mechanically--regulatory risk is more material at Heathrow than comparators.

5. Total market return

In the next sub-sections, we discuss in more detail the CAA's methodological choices in relation to:

- the ex post TMR
- the ex ante TMR

In this section , we present the TMR range that we consider to be most appropriate based on the discussion in the previous sections.

---

[31] See: Merger Strategies and Capital Market Risk on JSTOR.

Ex post total market return

In its H8 Initial Proposals, the CAA reaffirmed its H7 approach and estimated the historical ex post TMR using the arithmetic average of DMS data from the Long-run DMS Database (2025) over 1900–2024. Nominal returns are deflated using annual CPIH inflation rates for each year. On this basis, CAA estimates a CPIH-real ex post TMR of 6.92%.[32]

While we agree with the CAA's approach, we consider that the ex post TMR estimate should be updated to incorporate the additional year of data available in the Long-run DMS Database (2026) , published in March 2026. Accordingly, we propose to estimate the ex post TMR using the arithmetic average of DMS data over 1900–2025, with nominal returns deflated using annual inflation in each year. This update results in a CPIH-real ex post TMR of 7.03%, primarily driven by the exceptionally strong equity returns observed in 2025 in the UK (24.4%). These outcomes are consistent with the broader global environment, which also experienced very high equity returns during 2025 (22.3%).

Ex ante total market return

In its H8 IPs, the CAA reaffirmed its view that both the Fama–French dividend growth model and the DMS decompositional approach provide credible estimates of the ex ante TMR, and that weight should be placed on both. Accordingly, the CAA derives a CPIH-real ex ante TMR range with a lower bound of 6.65%, based on the Fama–French dividend growth model, and an upper bound of 6.79%, based on the DMS decompositional approach. The CAA sets its CPIH-real ex ante TMR at the midpoint of this range, i.e. 6.72%.

We again update these estimates to reflect the latest year of data available from the Long-run DMS Database (2026), published in March 2026. Accordingly, we update both the Fama–French dividend growth model and the DMS decompositional approach to incorporate data spanning 1900–2025, rather than 1900–2024. Doing so results in an ex ante TMR range of 6.69%–6.83%, with a midpoint of 6.76%.

Details of the CAA's base estimates, together with the updated values, are set out in Figures E.ii.b.12 and E.ii.b.13 below.

CPIH-real	Formula	CAA H8 IPs	Updated CAA H8 IPs
Cut-off date		*1 November 2025*	*31 March 2026*
Average dividend yield	[A]	4.41%	4.40%
Average dividend growth rate	[B]	1.72%	1.79%

[32] CAA (2026), '[H8 initial proposals section 3: financial issues and calculating the price cap](#)', March, pp 50–51.

Bias adjustment	[C]	0.53%	0.53%
Historic ex ante return	[D] = [A]+[B]+[C]	6.65%	6.69%

Figure E.ii.b.12 – Table showing CAA ex ante total market return based on the Fama–French dividend growth model. Note: Based on a cut-off date of 31 March 2026. Values may not add up due to rounding. Source: Oxera analysis based on CAA H8 Initial Proposals and DMS data.

CPIH-real	Formula	CAA H8 IPs	Updated CAA H8 IPs
Cut-off date		*1 November 2025*	*31 March 2026*
Geometric mean dividend yield	[A]	4.55%	4.54%
Growth rate of real dividends	[B]	0.64%	0.69%
Geometric mean ex ante TMR	[C] = [A]+[B]	5.19%	5.23%
Geometric-to-arithmetic conversion	[D]	1.61%	1.60%
Arithmetic ex ante TMR	[E] = [C]+[D]	6.79%	6.83%

Figure E.ii.b.13 - CAA ex ante total market return based on the DMS decompositional approach. Note: Based on a cut-off date of 31 March 2026. Values may not add up due to rounding. Source: Oxera analysis based on CAA H8 Initial Proposals and DMS data.

Total market return estimate

While we broadly agree with the CAA's approach to deriving an ex post TMR under a stable TMR assumption, we do not consider it appropriate to assume that the TMR is fixed.

In particular, the current 'through-the-cycle' and fixed TMR approach can easily understate the returns required by investors in a rising or persistently high interest rate environment. Notably, where gilt yields were last at comparable levels to current yields, regulatory TMR allowances in the water and energy sectors were in the range of 7.5-8.0% (CPIH-real),[33] which is materially above the long-run ex-post estimate of

[33] Ofgem (2012), 'RIIO-GD1: Final Proposals Finance and uncertainty supporting document', 17 December, p. 22; Ofwat (2009), 'Future water and sewerage charges 2010-15: Final determinations', November, p.128. Ofwat's PR09 and Ofgem's RIIO-GD1 and RIIO-T1 decisions were set in RPI-real terms. These figures have been converted into CPIH-real allowances using the long-term RPI–CPIH wedge published by the OBR. We have reflected the changes in the long-term wedges over time.

approximately 7.03% (CPIH-real) derived under a stable TMR assumptions. As such, this warrants including a moderate adjustment to reflect evolving interest rate conditions.

We suggest supplementing the existing Ex Post approach with a TMR estimate based on a stable ERP framework, whereby the ERP is estimated directly, and the RFR is then added back to derive the TMR. This method builds upon the CMA approach taken in its PR24 Determinations, in which the CMA explicitly drew on a stable ERP methodology, alongside ex post and ex ante TMR estimates to inform the upper end of its provisional TMR range. Our approach updates this method with an additional year of returns data, and adjusted for geometric differencing of annual ERP estimates to follow academic convention, such as the work by Dimson, Marsh and Staunton. The geometric differencing ensures that the impact of inflation is accurately captured within the risk-free return and that the final ERP estimates are unaffected by the choice of the inflation index. This approach results in a TMR that varies through the cycle, appropriately reflecting changes in the interest rate environment.

Under a stable ERP approach, the ERP is defined as the average CPIH-real annual excess return on UK equities relative to the risk-free asset, proxied by the DMS UK government index, over 1900–2025, which yields an average CPIH-real ex post ERP estimate of 5.39%. Adding the updated RFR estimate of 2.20% implies a CPIH-real ex post TMR of 7.59%, which is 56 basis points (bps) higher than the TMR estimated under a stable TMR assumption.

Based on these considerations, HAL's updated TMR ranges adopt a lower bound that is aligned with the CAA's approach to ex ante TMR, updated to a cut-off date of 31 March 2026 and incorporating an additional year of DMS data.

HAL's upper bound is updated to reflect the adoption of a stable ERP approach, alongside the inclusion of the additional year of DMS data in deriving the ex post TMR. By contrast, the updated CAA upper bound relies solely on incorporating the additional year of DMS data under a stable TMR assumption.

The resulting TMR estimate ranges are presented in Figure E.ii.b.14 below.

CPIH-real	Updated CAA H8 IPs	HAL
Ex ante TMR	6.76%	6.76%
Ex post TMR	7.03%	7.59%
Midpoint	6.89%	7.18%

Table E.ii.b.14 - TMR estimation. Note: Based on a cut-off date of 31 March 2026. Values may not add up due to rounding. Source: Oxera analysis based on DMS data.

6. Effect of illiquidity on beta estimation

As discussed in Section 2.2, Fraport is an outlier with low liquidity among the comparators.

	AENA	AdP	Fraport	Zurich	Average of FTSE 100 companies	Average of Stoxx 600 companies
Relative (%) Bid-Ask Spread (2020-2026)	0.09%	0.11%	0.68%	0.04%	0.08%	0.12%

Figure E.ii.b.15 – Bid-Ask spread of comparators and constituent companies of major indices. Source: Heathrow analysis of Bloomberg.

Illiquidity has four dimensions, each of which may be cause for concern when estimating the comparators' betas.

- **Tightness.** Tightness refers to the total cost of executing a 'round-trip' transaction (purchasing and then selling a stock). This is most commonly proxied by the bid-ask spread, which is the difference between the price at which a market maker will buy and sell a share at any given moment. The wider the bid-ask spread (i.e., the less 'tightness') the more the observed price 'bounces' between trades that are made that match the 'bid' prices and those that match the 'ask' prices. A stock with less tightness will 'bounce' more, which in turn reduces the correlation of the stock's most recent trade price with the price of the market portfolio. This causes a downwards bias in the beta of stocks with low tightness.
- **Depth.** Depth refers to the volume of shares that can be bought or sold at any given price. A deep order book has many shares available at a given price and a single transaction is unlikely to exhaust the available shares at said price. Conversely, a stock with a shallow order book has few shares listed with buying and selling orders at any given price. This means that even modest trades can cause significant movements in the mid-price (the mid-point between the best bid and best ask) as stocks at the prevailing 'best price' will quickly run out. A stock with less depth will experience these jumps in listed mid-prices more frequently, which in turn reduces correlation with the market portfolio. This causes a downwards bias in the betas of shallow stocks.
- **Resiliency.** Resiliency refers to the speed with which a stock's price recovers to its true value following a liquidity-driven[34] price disturbance. A stock with low resiliency will exhibit persistent price deviations following a trade, as there are insufficient active participants to quickly arbitrage the price back to the true value. This slow reversion to the pre-shock price introduces further noise into the time series of

[34] A liquidity-driven price disturbance is a temporary deviation of a stock's observed price from its true underlying value that is caused by the mechanics of trading rather than by any new information about the stock's fundamentals.

stock returns used to estimate beta, compounding the biases arising from tightness and depth.

- **Delay**. Delay refers to the speed with which market-wide information is incorporated into a stock's price. This is influenced by two different mechanisms: (i) the frequency of trades; and (ii) the speed of information diffusion. Regarding (i), illiquid stocks trade less frequently, which can lead to a delay between when there is a change in the underlying value of a stock, and when this is reflected in market transaction prices. Regarding (ii), illiquid stocks attract less attention from analysts and active traders whose activity rapidly incorporates new macroeconomic and market-wide information into prices. Illiquid stocks therefore react to market-wide news slowly, as there are fewer participants monitoring and actively trading them. The true underlying value of the stock may change following new information, but it may take time for this to be reflected in the price if trades are slow to happen. Both effects (i) and (ii) lead to a discrepancy between the most recent price data for the illiquid stock and the wider market. This reduces the empirical correlation between the two and leads to a downwards bias in the estimated beta of the stock.

In this section we discuss the consequences of low liquidity for beta estimation:

- Theoretical relationship between illiquidity and beta estimation.
- Literature that empirically investigates the theoretical relationships discussed.

Theoretical relationship between illiquidity and beta estimation

In this section we elaborate on the mechanisms that lead illiquidity across the above four dimensions to affect stock prices, leading them to appear empirically less correlated with market returns and causing a downwards bias in estimated betas.

In each of the following subsections, we set out further details on the four dimensions of illiquidity and the particular mechanism through which they cause bias in the estimation of beta:

a) Low tightness
b) Low depth
c) Low resiliency
d) Greater delay

Tightness and beta bias

At any moment, the market is characterised by two key prices: the lowest available ask (selling) price and the highest available bid (buying) price, together known as the inside quotes or best prices. The gap between them is the bid-ask spread, and serves as a fundamental measure of transaction cost for those trading at the market price. Prices do not move continuously, but in discrete steps (ticks).

In practice, recorded transaction prices bounce between buyer-initiated trades (executed at the ask price) and seller-initiated trades (executed at the bid price). Even when a stock's fundamental value has not changed, its observed price therefore bounces between these two levels, introducing artificial variability into measured returns. This is known as bid-ask bounce.

This bounce is unrelated to movements in the broader market (which, being a diversified portfolio, exhibits very little bounce of its own), therefore, it acts as noise that dilutes the estimated correlation between the stock and the market. The result is a downward bias in estimated beta. This bias is more pronounced the wider the bid-ask spread: a stock with low tightness bounces more, and its measured returns diverge further from the market's true underlying return.

A related but distinct source of bias arises from price discreteness. Since share prices are quoted in discrete increments (such as pennies and cents), prices cannot move continuously, imposing a minimum effective spread even when the quoted bid-ask spread is narrow. While this effect is most material for low-priced stocks (where even a one-penny increment represents a meaningful fraction of total return) it operates through the same mechanism as bid-ask bounce and compounds the downward bias in beta estimation.

Low depth and beta bias

When trading shares, agents can place two main types of orders. Firstly, a market order is a request to buy or sell a specified number of shares immediately at the best available price. A market order prioritises speed of execution over price certainty. Secondly, a limit order specifies the worst price the trader is willing to accept (the limit price would be higher for a buyer and lower for a seller, all else equal). Limit orders may not result in an immediate trade and instead sit in a queue known as the limit order book, waiting to be matched. Buy limit orders are referred to as bids, and sell limit orders as offers or asks.

As market orders arrive, they are matched against limit orders on the opposite side of the book, first by price and then by time of arrival. Because orders vary in size, this matching is not necessarily one-to-one. For example, suppose a market buy order is made for 300 shares and:

- The best ask is 200 shares with an ask of £1; and
- The next best is 300 shares with an ask of £2.

The buy market order for 300 shares would first consume all 200 shares at £1, then take the remaining 100 at £2[35], yielding an average execution price above the quoted ask.[36] This divergence between the quoted price (£1) and effective price (£1.33[37]) is known as the effective spread) reflects the true cost of executing an order of meaningful size or in an illiquid market.

The density of limit orders at each price level is the order depth of the book. A deep order book, with many shares available at or near the best bid and ask prices, allows large trades to be absorbed with minimal price impact. A shallow order book, where few limit orders are stacked close to the prevailing price, has low depth. In such a market, even modest trades can exhaust the available shares at the best price and force subsequent portions of the order to be filled at progressively worse prices. This will move the observed market price materially, even in the absence of any new information about the stock's fundamental value.

It is this last point that gives rise to downward bias in beta. When a stock has low depth, ordinary trading activity generates price movements that reflect the mechanics of order matching rather than genuine changes in underlying value. These movements are largely idiosyncratic. That is to say, they arise from the particular pattern of orders arriving in that stock's order book, and are unrelated to movements in the price of the market portfolio. When beta is estimated by regressing a stock's returns against market returns, this microstructure noise[38] dilutes the apparent relationship, causing beta to appear lower than it truly is.

Low resilience and beta bias

Even after a trade has been executed, its impact on the observed price may persist for some time. Resiliency refers to the speed with which a stock's price recovers to its fundamental value following a liquidity-driven price disturbance (a price movement caused not by new information about the firm, but by the mechanics of order flow and market structure).

In a resilient market, active participants rapidly identify and arbitrage away any such distortions, restoring the observed price to its fair value quickly. In a stock with low Resiliency, however, this correction is slow. Few participants may be monitoring the stock closely enough to act on the temporary mispricing, and so the price deviation

[35] The best price available on the market would then be £2.
[36] A shallow order book will mean that the supply and demand curves for the stock at any instant in time are irregular step-like functions with long flat regions and large jumps. In this example, this jump occurs at 200 shares.
[37] (200/300) * £1 + (100/300) * £2.
[38] Microstructure noise is a commonly used term within the academic literature used to refer to the wedge that exists between the observed price of a stock and its 'efficient', true or fundamental value. This noise arises due to the mechanics of stock trading and inherent frictions such as bid-ask spreads, discrete pricing, and order processing delays. This noise largely cancels out when stocks are highly aggregated, such as in the market portfolio.

persists across multiple observations in the stock returns time series data. The observed price lingers away from fundamental value, not because anything has changed about the firm's prospects, but simply because the market lacks the active participation needed to correct itself promptly.

This slow reversion introduces a further layer of noise into the time series of stock returns used to estimate beta. Just as with tightness and depth, the noise is idiosyncratic and reflects the particular liquidity conditions of that stock rather than any genuine co-movement with the broader market. This idiosyncratic component dilutes the measured relationship with market returns, producing a downward bias in the estimated beta.

Low resiliency therefore compounds the biases arising from tightness and depth. Where those dimensions of illiquidity introduce noise at the moment of trading, low resiliency causes that noise to persist and propagate through subsequent return observations, further weakening the signal that beta estimation relies upon.

Delay and beta bias

Delay refers to the speed with which information associated with market-wide movements is incorporated into a stock's price. Unlike the previous three dimensions of illiquidity, which introduce noise into observed prices at or around the moment of trading, delay operates through a distinct mechanism.

Delay creates a temporal disconnect between the information reflected in an illiquid stock's price and that reflected in the broader market, attenuating their measured co-movement and biasing estimated beta downward. For example, delay may mean that movements in the wider market for a given period are only reflected in movements in the illiquid stock’s price in the following period, rather than occurring within the same period (as could be done for a liquid stock). This may lead to a weaker *estimated* relationship between the stock’s returns and market returns, which are conventionally compared within the same period, than the *true* relationship, which occurs with a lag. This source of bias is distinct from the noise associated with resiliency, which relates to stock-specific price disturbances, and the speed of their correction.[39]

This delay arises through two related but conceptually distinct channels. The first channel is the infrequency of trades, sometimes referred to as non-synchronous trading. Beta is typically estimated using the most recent recorded price of the stock and the most recent recorded price of the market portfolio. For liquid stocks, these prices are (closely) contemporaneous. For illiquid stocks, however, there may be significant time gaps between transactions, spanning periods within a trading day or even across trading days. As a result, the most recent recorded price of an illiquid stock

[39] We note that in practice, the extent to which a stock is illiquid with respect to Resiliency may be correlated with the extent to which it is illiquid with respect to Delay.

may reflect market conditions and information from a materially earlier point in time than the corresponding market price. These two time series are therefore not directly comparable: the stock's return may be stale, or may reflect a delayed catch-up to trends that have already been incorporated into market prices. This discrepancy mechanically reduces the measured correlation between the stock and the market, producing a downward bias in estimated beta.

The second channel is the slow diffusion of information. Liquid stocks attract sustained attention from analysts and active traders, whose monitoring and trading activity rapidly incorporates new macroeconomic and market-wide developments into prices. Illiquid stocks, by contrast, attract fewer such participants. Even when new information arrives that is relevant to the stock's fundamental value, it may take time before sufficient trading activity occurs to fully reflect that information in the observed price. The stock therefore responds to market-wide news with a lag, further widening the temporal gap between its return series and that of the market.

Both channels produce the same result: an asynchronicity between the illiquid stock's returns and contemporaneous market returns that weakens their measured relationship with the market portfolio and depresses estimated beta. It is worth noting, however, that delay does not operate entirely independently of the other dimensions of illiquidity discussed above. In practice, stocks with wide bid-ask spreads tend also to trade infrequently, meaning tightness and delay are empirically correlated even where they are conceptually distinct and operate through different mechanisms. Similarly, a stock with low resiliency, whose price corrects slowly following a liquidity disturbance, may also be slow to incorporate new information, creating overlap between delay and resiliency. These correlations mean that in practice the four dimensions of illiquidity tend to co-occur and their biasing effects on beta estimation are likely to compound one another.

Empirical studies on tightness and beta

The literature on the effect of tightness on stock price noise generally seeks first to understand the causes and dynamics of volatility in asset pricing. Within this, the papers often consider a wide range of measures of illiquidity, including those that relate to tightness.

For example, Aït Sahalia and Yu (2009)[40] seek to better understand the nature of the noise contained in high frequency price data that is averaged out when data is aggregated, either over time or across stocks. They further seek to relate stock price noise to observable financial characteristics of the underlying assets and in particular to different financial measures of liquidity. They hypothesise that more liquid assets

[40] 'High frequency market microstructure noise estimates and liquidity measures'. Y. Aït Sahalia; J. Yu. Annals of Applied Statistics (2009).

tend to generate less microstructure noise. In testing this hypothesis, they assess how a wide range of liquidity measures affect measures of stock return volatility. This represents an extension of previous analysis, for example, Roll (1984)[41], who assumes that volatility is driven exclusively by the bid-ask spread.

Aït Sahalia and Yu conduct their analysis on a large sample of high frequency data on all transactions recorded from New York Stock Exchange (NYSE) common stocks between 1 June 1995 and 31 December 2005. The specific measures of liquidity they utilise include bid-ask spread, transaction size, effective cost of trading and Amihud's ILLIQ ratio. They then relate these to the high frequency estimates of market microstructure noise and find that the microstructure noise is positively correlated with their measures of illiquidity. In particular, they find that the intraday bid-ask spread explains most (63%) of the variation in noise.

Bandi and Russel (2006)[42] are similarly motivated to Aït Sahalia and Yu and seek to establish a procedure to purge high frequency return data of microstructure components. They do so utilising data from stocks in the S&P100 index, on a shorter data set focused on February 2002. They decompose the effects of microstructure noise components and attribute them to the extent of illiquidity of a given stock. In doing so, they demonstrate that microstructure components create significant price volatility. Specifically, they find that a 1% increase in quoted bid-ask spreads translates into a 1% increase in the standard deviation of the noise.

Kim, Oh and Wang (2024)[43] conduct a similar decomposition exercise and demonstrate bias specifically in beta estimates of stocks with a high bid-ask bounce. They find that this bias follows a predictable pattern as the number of observations (daily trades) increases, but that it does not disappear as the sample size increases. While this points to a bias that does not fully disappear as liquidity increases, it indicates that bias in beta due to bid-ask bounce is greater for more illiquid assets. Importantly, they find that this bias is predictable as it exhibits autocorrelation over time (i.e., the bias is correlated with its own value in previous periods), and construct a new estimator of beta which accounts for the bias caused by bid-ask bounce.

Empirical studies on depth and beta

Assessments of how market depth affects price volatility generally examine how steep or unevenly distributed order books influence price movements relative to the effect of large individual trades. They find that these order book effects have a greater role in

---

[41] 'A simple model of the implicit bid–ask spread in an efficient market'. R. Roll. Journal of Finance (1984).
[42] 'Separating microstructure noise from volatility'. F. Bandi; J. Russell. Journal of Financial Economics (2006).
[43] 'Robust Realized Integrated Beta Estimator with Application to Dynamic Analysis of Integrated Beta'. D. Kim; M. Oh; Y. Wang. Journal of Econometrics (2024).

driving price volatility and, as such, the shallow order books of illiquid stocks cause a significant increase in noise.

Farmer et al. (2004)[44] study the cause of large fluctuations in prices on the London Stock Exchange. Their dataset includes a sample of 16 high volume stocks traded over the same 4-year period from 1999-2002. They find that large price fluctuations are driven by liquidity fluctuations (variations in the market's ability to absorb new orders) more than sudden spikes in demand or supply. This contrasts with the earlier findings of Gabaix et al. (2003)[45] that large price movements are caused by high volume orders. The rationale for Gabaix et al.'s finding, that large price changes are caused by large market orders, is intuitive; a very large market order will dig deeply into the limit order book, causing transactions at many price levels, increasing the spread, and changing the mid-price.

However, Farmer et al. find that even when a stock is otherwise very liquid, there can be substantial gaps in the order book, corresponding to a block of adjacent price levels containing no quotes. When such a gap exists next to the best price, a new order can buy or sell all the shares at the best price, triggering a large change in the mid-point price.

Empirically, Farmer et al. find that the distribution of large price changes closely reflects the distribution of gaps in the limit order book. In a market where participants placed many small orders uniformly across prices, such large price fluctuations would not happen. Unsurprisingly, they find that more lightly traded stocks are more strongly affected by a shallow order book (i.e., their prices exhibit more noise) but critically note that the pattern holds across all stocks that lack depth in their order book.

Webber and Rosenow (2006)[46] analyse large stock price changes[47] for TAQ data for 1997 and order book data from Island ECN for 2002. Much like Farmer et al. they argue that that these large price changes are not due to large trading volumes but rather due to the shallowness of the relevant stock's order book.

Their results are corroborated further by Duong and Kalev (2008)[48], who investigate the informativeness of the slope of the order book in determining price volatility for the constituent stocks of the S&P/ASX 100 index on the Australian Stock Exchange. They hypothesise:

---

[44] 'What really causes large price changes?'. J. Farmer; L. Guillemot; F. Lillo; S. Mike; A. Sen. Quantitative Finance (2004). As identified in our initial literature review in our document titled "Analysis of HAL Risk Relative to CAA Comparator Operators".
[45] 'A theory of power-law distributions in financial market fluctuations'. X. Gabaix; P. Gopikrishnan; V. Plerou; H. Stanley. Nature (2003).
[46] 'Large stock price changes: volume or liquidity?'. P. Webber; B. Rosenow. University of Koln (2004).
[47] Large is defined as a price change of more than five standard deviations.
[48] 'Order Book Slope and Price Volatility'. H. Duong; P. Kalev. Monash University (2008).

- Firstly, that the limit order book slope is negatively related to future volatility. This is based on the empirical findings of Ahn et al. (2001)[49] and Pascual and Veredas (2006)[50], which support this relationship.
- Secondly, that the limit order book slope is negatively related to long term stock price volatility[51]. This is again based on an empirical precedent, set by Andersen et al. (1997)[52] and Muller et al. (1997)[53].
- Thirdly, that the limit order book slope of the demand (buy) side is more informative than the limit order book slope of the supply (sell) side with regards to long term stock price volatility. This is, in part, motivated by the findings of Griffiths et al. (2000)[54] who provide evidence that aggressive buy orders on the Toronto Stock Exchange are more informative (as to price volatility) than aggressive sell orders.

These three hypotheses speak to the same argument – that a steeper order book is associated with an increase in stock price volatility. Their empirical findings support these hypotheses: they find that order book slope predicts future price volatility in the majority of the constituent stocks of the S&P/ASX 100 index over the period between 1 July 2005 and 30 June 2006. In their supplementary analysis, they additionally find that the limit order book slope is more informative than the bid-ask spread and market depth for predicting long term volatility.

Empirical studies on resiliency and beta

The existing literature relating to the impact of resiliency is sparse relative to the other measures of liquidity. Notably, it does not directly estimate the extent of resiliency and link this to an estimated bias on beta, nor does it directly estimate the extent to which low resiliency contributes towards noise in practice. Two useful findings can nevertheless be drawn.

- Firstly, the literature establishes that there are varying degrees of resiliency between stocks, material enough to have a substantial impact on their average returns.[55] As there is a clear mechanism through which low resiliency may contribute towards

[49] '[Limit orders, depth, and volatility: evidence from the Stock Exchange of Hong Kong]'. H-J. Ahn; K-H Bae; K. Chan. Journal of Finance (2001).
[50] '[Does the Open Limit Order Book Matter in Explaining Long Run Volatility?]'. R. Pascual; D. Veredas. University of Warwick (2006).
[51] Which they refer to as the future permanent component of volatility.
[52] '[Heterogeneous Information Arrivals and Return Volatility Dynamics: Uncovering the Long-Run in High Frequency Returns]'. T. Andersen; T. Bollerslev. Journal of Finance (2012).
[53] '[Volatilities of different time resolutions — Analyzing the dynamics of market components]'. U. Muller; M. Dacorogna; R. Dave; R. Olsen; O. Pictet; J. Weizacker. Journal of Empirical Finance (1997).
[54] '[The costs and determinants of order aggressiveness]'. M. Griffiths; B. Smith; A. Turnball; R. White. Journal of Financial Economics (2000).
[55] As outlined below, the literature demonstrates the presence of a 'resiliency returns premium', in which less resilient stocks earn a greater average return.

bias, the potential for the comparators to have low resiliency in itself is indicative of the potential for bias through this form of illiquidity.[56]

- Secondly, the literature establishes that the least resilient companies earn a material average returns premium. One possible explanation for this premium is that there is greater volatility in stock returns for un-resilient stocks (which investors must be compensated for). This volatility may, at least in part, arise from increased noise. If so, this would act as a mechanism through which lower resiliency could be associated with higher noise, potentially leading to bias in the estimation of equity beta.

For example, Hua et al. (2020)[57] demonstrate the presence of varying degrees of resiliency between stocks. They measure resilience as the covariance between a stock's 9:30-10:00 opening return, and its 10:00-16:00 rest of the day return, standardised by the daily return variance. They estimate a significant non-resiliency returns premium that ranges from 33 to 57 basis points per month for companies with the lowest decile of resiliency relative to those with the highest decile.

J Kim and Y Kim (2019)[58] measure resiliency by adopting a model of stock prices in which stock prices are composed of permanent and transitory components (in line with Hasbrouck (1993)[59], Boehmer and Kelley (2009)[60], and Bao, Pan, and Wang (2011)[61]). Under this model framework, they apply the trend-cycle decomposition methodology introduced by Beveridge and Nelson (1981)[62] to estimate the transitory components from the daily stock prices[63]. They then transform the estimated series of transitory prices into spectral functional forms[64] at different frequencies. This transformation enables them to explicitly derive the time and distance components of reversion cycles, which are then used to calculate the speed of a transitory price reverting to its fundamental value as their measure of resiliency. Similar to Hua et al. (2020), they find

[56] As outlined, there is a clear theoretical link between the degree of Resiliency of a stock, price volatility noise, and the extent of bias in the estimation of equity beta.
[57] 'Resiliency and Stock Returns'. J. Hua; L. Peng; R.A. Schwartz; N.S. Alan. The Review of Financial Studies (2020).
[58] 'Transitory prices, resiliency, and the cross-section of stock returns'. J. Kim; Y. Kim. International Review of Financial Analysis (2019).
[59] 'Assessing the Quality of a Security Market: A New Approach to Transaction-Cost Measurement'. J. Hasbrouck. The Review of Financial Studies (1993).
[60] 'Institutional Investors and the Informational Efficiency of Prices'. E. Boehmer; E.K. Kelley. The Review of Financial Studies (2009).
[61] 'The Illiquidity of Corporate Bonds'. J. Bao; J. Pan; J. Wang. The Journal of Finance (2011).
[62] 'A new approach to decomposition of economic time series into permanent and transitory components with particular attention to measurement of the 'business cycle''. S. Beveridge; C.R. Nelson. Journal of Monetary Economics (1981).
[63] It is this transitional part of stock prices that would be determined by the stock's Resiliency.
[64] Rather than examining a time series purely over time, spectral analysis decomposes its variation across different frequencies (e.g., short-term versus long-term cycles), facilitating the identification of transitory components.

varying degrees of resiliency between firms, and that portfolios with lower resiliency have higher average returns than those with lower resiliency.[65]

Empirical studies on delay and beta

As set out above, there are two elements to illiquidity related to delay, both of which have the potential to bias estimates of beta. Firstly, the infrequency of trades and secondly the slow diffusion of information. Given that these concepts are related, the literature does not always treat them separately.

The academic literature has long recognised the issue of infrequent trades:[66]

- Scholes (1977)[67] shows that non-synchronous trading of securities introduces an errors-in-variables problem in market model regressions, which can bias beta estimates downwards. The authors describe how "*[non-synchronous] trading of securities introduces … a potentially serious econometric problem of errors in variables*".
- Dimson (1979)[68] extends this by demonstrating that thin trading can lead to downward-biased beta estimates, because stock prices adjust to market information with a lag when trading is infrequent.
- Cohen et al. (1983)[69] find that a "*substantial correction is needed to get "true" beta estimates*" when dealing with asynchronous data.

When infrequent trading extends to the point at which trades do not happen for multiple days, Auret et al. (2014)[70] demonstrate a monotonic downward bias in OLS-estimated betas for every additional day that a stock goes without trading and propose a series of alternative estimators to correct for infrequent trading.

Most recently, Wiedemann (2025)[71] conducts a comprehensive assessment of the bias caused to beta estimates by both asynchronous trading and slow information diffusion. While the issue of infrequent trading has been long recognised (as seen above), the

---

[65] Specifically, they find that the monthly return of portfolios with higher Resiliency, those within decile one of measured Resiliency, are higher than those in decile 10, which are 0.934% per month and 0.454% per month respectively. This corresponds to a material illiquidity premium of 0.48%, similar to that found by Hua et al. (February 2020).

[66] As identified in our initial literature review in our document titled "Analysis of HAL Risk Relative to CAA Comparator Operators".

[67] 'Estimating betas from nonsynchronous data'. M. Scholes; J. Williams. Journal of Financial Economics (1977).

[68] 'Risk measurement when shares are subject to infrequent trading'. E. Dimson. Journal of Financial Economics (1979).

[69] 'Estimating and Adjusting for the Intervalling-Effect Bias in Beta'. K. Cohen; G. Hawawini; S. Maier; R. Schwartz; D. Whitcomb. Management Science (1983).

[70] 'Thin-Trading and Beta Estimation: Results From a Simulated Environment'. C. Auret; D. McClelland; T. Wright. Studies in Economics and Econometrics (2014). As identified in our initial literature review in our document titled "Analysis of HAL Risk Relative to CAA Comparator Operators".

[71] 'Quantifying the Beta Estimation Bias and its Implications for Empirical Asset Pricing'. T. Wiedemann (2025).

contribution of slow information diffusion to delay-related bias has received considerably less attention and has proven difficult to measure.[72] This paper seeks to that gap by exploiting intraday transaction-level data from the TAQ database, which records the exact time of each stock's last trade on a given day, to construct a precise empirical measure called the “time-to-close” (or TTC) that captures the degree of asynchronicity between an individual asset and the market close. This allows the author to disentangle the two sources of bias and quantify each separately.

The paper’s core finding is that each of these can be material sources of bias, particularly slow information diffusion, which remains relevant even for more frequently traded stocks where the impact of asynchronous trading times is minimal.

Specifically, Wiedemann finds that:

- Historically, where stocks were traded less frequently,[73] these sources of bias can lead beta to be underestimated by up to 50%, with roughly three-quarters of this bias attributable to slow information diffusion.
- Over more recent years, where stocks have been traded more frequently (and bias associated with asynchronous trading is less of an issue), beta may still be underestimated by over 30% for the 20% least traded stocks due to the impact of slow information diffusion alone.[74]

[72] This is because both mechanisms lead to a similar effect in the observed data, and a delay in movements in a stock’s replace relative to movements in prices for the wider market
[73] These historical estimates on the extent of bias are based on July 1963 to December 1994.
[74] Wiedemann estimates a bias of 32.7% based on January 1995 through to 2020.