

Proposed approach for incorporating RCF and QCRA evidence for informing risk funding for Major Projects

Report prepared for NGET

05 March 2026

1 Introduction and purpose of this paper

Frontier has been commissioned to undertake an independent assessment of the approach to risk funding applied to BTNO. While this paper addresses Ofgem's minded-to position for BTNO specifically, the discussion is relevant to the wider pipeline of major projects.

Our findings in this paper relies on evidence and analysis from NGET and Oxford Global Projects (OGP). We have also relied on wider expertise and analysis produced by OGP (e.g. where we refer to industry convention). The views in this paper have been formed by interpreting the evidence and data shared with us as part of this process, and we have not undertaken any independent verification of that information and analysis.

We note that Ofgem seems to have departed from its recent approach of setting the risk allowance at the P70 of the Quantitative Cost Risk Assessment (QCRA) distribution, when its previous approach¹ was more consistent with industry convention. Ofgem has chosen instead to set risk funding for BTNO at the P50 of the QCRA distribution submitted by NGET, stating that the P50 is the "*optimal mix of balancing risk while mitigating the costs from escalating*."²

As we discuss in this paper, the P50 is in fact, not optimal, and may result in systematic underfunding of risk for major projects as it is likely materially lower than the expected cost of these projects. This is because major project risks are typically skewed, meaning the mean exceeds the P50 and better reflects the expected value. NGET has also exercised caution in producing their QCRA to avoid double counting of the protections provided under the ASTI framework. However, this approach may mean that even the mean understates true delivery risks, as some risks are not fully captured. To address this, NGET has set allowances in the P50 - P70 range as a pragmatic approach to reflect additional risks while avoiding double counting protections already included in the ASTI framework.

¹ The approach taken in previous ASTI project assessments

If Ofgem's objective is to set a fair level of funding, the starting point should be the mean of the QCRA distribution (rather than the median). This is discussed in the remainder of our report, which is organised as follows:

- First, we outline Ofgem's minded-to position for assessing risk funding as set out in the BTNO PA consultation.
- Then, we present a framework for setting the risk allowance, led with QCRA evidence but and relying on reference class forecasting evidence as a cross-check. This framework sets out how different sources of evidence can be used to inform the appropriate level of risk allowance.
- Following this, we explain how the framework can be applied to BTNO.
- Finally, we discuss our overall conclusions and further considerations that need to be accounted for to develop a robust approach towards setting risk allowances for future major projects.

2 Ofgem's minded-to approach based on the P50 could result in systematic underfunding

Ofgem's minded-to position on risk funding states that funding at the P50 of the Quantitative Cost Risk Analysis (QCRA) represents an efficient level of confidence. On this basis, Ofgem proposes reducing BTNO's risk funding by c.£8m (24/25 prices) to align with a P50 position (on the QCRA).³

We understand that this represents a departure from Ofgem's existing approach in previous ASTI project assessments (e.g. OPN2, EGL1), where Ofgem allowed risk funding at levels above the median implied by QCRA evidence.⁴ NGET considered the previous approach to be more appropriate for a number of reasons.

First, the P50 of the QCRA does not represent the expected (fair) outcome; this would be represented by the true mean.⁵ In large infrastructure projects, risk distributions can be asymmetrical, with a longer upper tail. In such distributions:

- The outcomes are skewed to the downside; and,
- Even though the P50 represents the most likely outcome, the expected cost is higher than the P50 level.

3 [REDACTED]

4 [REDACTED]

⁵ Specifically, the probability weighted mean.

Under these conditions, setting the risk allowance at P50 will systematically underfund projects over time. While P50 implies that 50% of projects will likely require an allowance at or below this level, the law of large numbers implies that, across projects, outturn costs will converge to the true mean, not P50.

If the regulatory objective is to provide a fair bet and efficient level of allowance, the appropriate benchmark is the expected value of the risk distribution, not the P50. This principle is reflected in established public-sector appraisal guidance (e.g. Department for Transport).⁶

Second, QCRA captures only quantifiable risks - not the full risk exposure. QCRA models reflect the distribution of identified and quantifiable risks at the time of submission. QCRA typically cannot capture:

- Unknown risks;
- Emerging risks; and,
- Risks that are difficult to quantify or cannot be robustly quantified.

For large, complex and first-of-a-kind infrastructure projects, these unquantified risks can be material. As a result, the QCRA distribution is likely not a complete representation of true project risk distribution. Interpreting P50 as an “efficient” funding level assumes that the QCRA fully captures risk exposure, which is not the case.

This limitation is one reason why industry practice has traditionally interpreted QCRA outputs with caution and avoided funding at the median.⁷ Across major infrastructure sectors and geographies, it has long been standard practice to fund risk allowance at higher confidence levels, especially for complex projects (e.g. P80).⁸ This reflects a practical understanding that risk distributions are skewed for major and/or complex projects, not all risks are quantifiable and underfunding can increase delivery risk.⁹

⁶ Department for Transport (2025), [TAG Unit A1.2, Scheme costs](#), 3.2.23.

⁷ OGP notes that this is recognised in the Green Book for example, where HMT discusses the systematic tendency for “practitioners to be over-optimistic about key assumptions in appraisal”. HMT notes that capital costs tend to be higher in reality than expected. This reflects the challenge of accounting for unknown risks at the point at which the budget is set. HMT recommends “increasing estimated costs and timeframes” to provide a more realistic view. Please see HMT (2026) The Green Book, UK Government Guidance on Appraisal.

⁸ See for example Department for Transport (2025), TAG Unit A1.2, Scheme costs, para 3.5.3. See also Contingency Guidelines Risk Engineering Society (RES) and Engineers Australia Author: Pedram Danesh-Mand, 2nd Edition, 2019, where the authors state “Where applicable, methods should consider the unique nature of mega-projects and complexity; the overall outcome may be more than the sum of the uncertainties and risks (i.e., non-linear)...With skewed distributions, the mean is generally greater than P50 – P60 being typical. This also highlights that the P50 (median) is not risk neutral, it is aggressive because the median does not address skewed reality. Heathrow and Network Rail take a similar approach, see Heathrow Capital Projects (2016 Procedures), and Network Rail (2024) <https://www.orr.gov.uk/sites/default/files/om/pr08-sdgenh-230508.pdf> .

⁹ Department for Health and Social Care (2022) UK New Hospitals Programme: Approach to Retained Contingency, Managing uncertainty and increasing the confidence of delivering a first of a kind programme within its funding envelope

Notably, in previous ASTI determinations (e.g. OPN2, EGL1), Ofgem did not adopt a strict P50 approach and allowed funding above the median implied by QCRA evidence. The current proposal therefore represents a departure from existing practice.

We note that whilst Ofgem has put in place other risk mitigation measures for TOs, many of these, e.g. access to the ASTI Uncertain Cost Reopener and the Totex Incentive Mechanism, are symmetric (i.e. they can impact TOs equally on the upside and downside). Therefore to the extent the underlying risk distribution is skewed, these regulator mechanisms are unlikely to fully mitigate even if some of them may moderate those most extreme tail end risks. The requirement to set a fair allowance still requires the true mean of the distribution, albeit post regulatory mechanism treatment.

Ofgem has also introduced risk mitigation measures for TOs such as Early Construction Funding, the Cost and Outcome Adjusting Events (COAEs) and Material Scope Change. While these measures do have the potential to reduce extreme tail-end risks, their impact is difficult to quantify. Therefore in NGET's risk quantification within its PA submissions, NGET has used professional judgement alongside ASTI guidance to exclude risks covered by these mechanisms in the QCRA.¹⁰ However, it is unlikely that the COAE and MSC mechanisms can fully mitigate the asymmetry of the true risk distribution even if they cover the most extreme outcomes. The outcomes that are just short of the cut-off point of these mechanisms still cause the risk distribution to be skewed, and because these are not included in NGET's current QCRA (because NGET had not envisaged that Ofgem would propose to fund the risk allowance at P50), the P50, as well as the Pmean of NGET's current QCRA is likely lower than the true expected cost of the project.¹¹

To sum up, setting a risk allowance at the P50 of this distribution therefore does not represent fair allowance that reflects the expected outcome for that individual project, due to the combination of the factors above. To provide a fair allowance, Ofgem's allowance must cover the cost in expectation, and because the current QCRA analysis incorporates an incomplete set of risks, the correct allowance will need to be higher than Ofgem's minded-to position - the P50 of that the QCRA distribution. Overall, We consider that it is most pragmatic to continue to rely on industry best practice of selecting higher confidence levels on the QCRA (as NGET has done so far).

However, if Ofgem decides to deviate from this method and focus on setting an allowance equal to the expected outcome, then it needs a framework that allows this expected outcome to be estimated appropriately, as opposed to simply taking the P50 of the QCRA analysis which is likely to lead to systematic underfunding of costs. We set out a framework that

¹⁰ NGET has excluded risks qualifying for COAEs from its quantitative risk analysis and has not included scope change risks in the underlying risk register.

considers a wider set of information in the following section, but noting there may be other appropriate approaches that could be followed.

3 A framework for setting the risk allowance at the expected outcome

In this section, we set out how Ofgem can rely on a wider range of evidence to inform the appropriate risk allowance based on the expected outcome for major projects. We set out a systematic approach to relying on not just the QCRA, but additional evidence from Reference Class Forecasting (RCF) which considers the actual performance of completed comparable projects.

We still consider the primary benchmark should be the QCRA, as it is targeted to the specific project, and sources of risks can be transparently explained. However, any method will be subject to limitations, including the QCRA, and we discuss how RCF can help provide additional information on the appropriate level of the required risk allowance.

3.1 Primary analysis based on QCRA

QCRA is a Monte Carlo-based methodology designed to estimate the probability distribution of project risks. It can capture:

- Known risks;
- Uncertainty ranges; and,
- Interactions between identified risks.

NGET's approach to date has focused on standard practice of modelling known risks only, in order to remain transparent and easy to follow (e.g. NGET has hitherto opted not to factor in uncertainty ranges and interactions between risks). Additionally, NGET has not so far factored in the cost of mitigating delay risk in its QCRA analysis, which is particularly relevant for major projects such as BTNO.

As set out above, Ofgem's proposed approach to fund risks at the QCRA modelled P50 is likely to be insufficient and thus, does not balance TO and customer interests. This is because of known limitations of the QCRA analysis, especially its inability to capture unknown and unquantifiable risks. Additionally, major projects typically exhibit tail risks and the P50 does not capture this. More generally, the limitations of QCRA imply that even the QCRA mean or P50 will likely be lower than the true mean (but the true mean is unobservable ex ante).

We consider that to mitigate the two issues above, two measures need to be taken:

- First, instead of P50, the mean implied by the QCRA should be used as a starting point (although this is still likely to be insufficient). But starting with the mean would at least ensure that the allowance covers the fair outcome, or expected value, of the QCRA risk distribution, reflecting the tail risks that have been incorporated into the QCRA analysis; and,
- Second, QCRA analyses will need to be more comprehensive such that they attempt to capture the impact of unknown or unquantifiable risks. But, it is difficult for these types of risks to be captured with accuracy, and thus a level of judgement and project-specific expertise must be applied when setting a fair and efficient allowance that serves the objectives of TOs and customers.

The level of this additional element can be cross-checked using reference class forecasting (RCF).

3.2 Reference class forecasting as a cross-check

Reference class forecasting (RCF) is considered a method for improving predictions by comparing a current project or decision to a broader “reference class” of similar past cases. Instead of relying solely on detailed, inside-view assumptions about the specific project, it uses empirical data from comparable, completed projects to estimate likely outcomes on costs and timelines. Because it is based on realised data, RCF inherently captures:

- Base cost and risk underestimation – due to optimism and other bias;
- Unknown risks; and,
- Internal modelling limitations and systematic forecasting error.

The approach is particularly useful in areas like infrastructure planning, policy evaluation, and large capital projects where overruns are common.

However, RCF’s limitation is comparability. While RCF provides a range of potential outcomes from a large sample of comparable – albeit not identical – projects, no past project is identical to the target project. Given this, RCF is unable to predict exactly where a project will outturn, but provides a benchmark perspective of the potential impact of materialised risks for the target project.¹²

When deploying RCF, a sufficiently large sample will be required to mitigate biases from idiosyncratic factors within individual projects in the historical RCF database. But even with a large sample size, if the target project has some unique characteristics which is not reflected in the reference class then the RCF would not provide an accurate forecast. In other words, if sole reliance was placed on RCF to define a risk allowance for a particular project, the

¹² Organisations using RCF recognise that there is a much wider range of outcomes than with QCRA.

estimation error of this risk allowance could be either positive or negative (can be an over- or under-estimate of the true (unobservable) value). The RCF sample may include projects which are less risky than the target project; in this case, the risk allowance implied by RCF would likely be lower than what the target project requires. The opposite case could also be true. This is an important consideration to have in mind when interpreting RCF outputs.

3.3 Proposed approach to cross-check a QCRA-led risk allowance

In this section we proposed to cross-check a QCRA-led risk allowance with RCF outputs. The cross-check consists of broadly two steps:

- The mean of the QCRA is the starting point, noting that it is more likely for the QCRA mean to be lower than the true mean of the distribution of the project's risks.¹³ To account for this, NGET has applied a professional judgement when determining the QCRA-led risk allowance for the project in question.
- Then, RCF evidence for the target project is considered as a cross-check, which provides additional comfort on the professional judgement applied in the previous step.

We describe this process in more detail in the following sub-sections.

Step 1: Identify the QCRA-led risk allowance using the mean as a starting point

For all the reasons set out in Section 3.1 above, the mean of the QCRA should be considered a plausible lower bound.

At this stage, it is also important to identify elements that the QCRA included or excluded, to understand the level of confidence that can be placed in the QCRA. Having regard to project-specific risk characteristics captured in the QCRA, we form a qualitative view on the appropriate uplift needed over the mean.

Step 2: Incorporating RCF evidence as a cross-check

Given RCF's strengths and limitations it can be considered as a cross-check to the QCRA-led allowance. As set out in Section 3.2, project-specific RCFs are established by sampling similar projects and examining the range of cost performance realised by these projects.

As we set out in Section 3.2 above, past projects will not be identical to the target project. The RCF outputs therefore represent a range of outcomes from similar projects, although in

¹³ As set out in Sections 2 and 3.1, NGET has exercised caution in producing their QCRA to date, to avoid double counting of the protections provided under the ASTI framework. This approach implies that even the mean understates true delivery risks, as some risks are not fully captured. The QCRA mean can be brought closer to the true mean by making QCRA analyses more comprehensive; but it is inherently not possible to fully account for unknown or unquantifiable risks. In general, the mean of the QCRA is likely to be lower than the true mean.

principle the mean of this range should be the applicable figure for operationalising this cross-check.

That said, specific circumstances of the target project must always be considered. The mean of a robust QCRA may ultimately be more representative of the target project than the range of uplifts implied by the RCF, which is based on outcomes of a of similar projects delivered in different environments (to the target project).

We would generally expect the mean value derived from the QCRA to be lower than the mean value implied by RCF, as RCF reflects realised risks including unknown risks which cannot be identified in the QCRA. However, as we discussed in Section 3.2, RCF can either over- or under-estimate the level of risk allowance, depending on the nature of the underlying sample of completed projects with differences in scope, complexity, delivery environment or risk profile. Therefore, if the risk estimate implied by the RCF's mean value is lower than the mean value of the QCRA, then this cross-check does not provide any additional information on the appropriate uplift over the QCRA mean. We cannot, however, in this case conclude that the QCRA mean is too high.

In summary, we propose that the RCF evidence is considered as a cross-check to robust, QCRA-led analysis, with a plausible lower bound for the risk allowance set at the QCRA mean. The RCF outputs can be used to cross-check any uplift deemed necessary above the QCRA mean by expert judgement. If the RCF outputs support the level of uplift, then we can be more confident in the expert judgement. If the RCF outputs do not support the uplift, however, this does not mean that the uplift is automatically wrong, rather that it is not supported by RCF analysis (which may not always perfectly match the target project).

3.4 Concluding remarks on the proposed approach

The proposed approach still focusses on the QCRA as the primary source of evidence for setting the risk allowance, but focuses on the mean of the distribution as a lower bound. The Pmean and a range from a project-specific RCF analysis can be used as a cross-check to the primary source of evidence. RCF may provide an additional level of comfort to the QCRA-led estimate.

This approach attempts to rely on a different sources of information where possible to identify the appropriate range for funding for NGET projects at expected level of costs, especially since major projects are likely to be of high complexity or have characteristics that are 'first of a kind'. However, it is worth noting that industry best practice typically suggests setting contingency

funding above the mean level, especially for complex, large projects with material schedule pressure.¹⁴

Setting the project cost allowance below the true expected cost of the project will systematically underfund the project, and is not in line with Ofgem's financing duty. Furthermore, the asymmetry of the risk faced by these large projects are not compensated for elsewhere in the price control package. In the discussion on whether aiming up on allowed equity return is appropriate due to asymmetric risks, Ofgem has made it clear that any risk asymmetry should be fixed at source. In this case, the allowance at the expected cost of large projects with skewed risk distribution is one of the sources of such asymmetries. We therefore consider that Ofgem must ensure that the allowance setting methodology does not lead to asymmetries.

In the following section, we apply the insights developed to BTNO.

4 Applying our framework to estimate the risk allowance for BTNO

In this section, apply the cross-check framework to BTNO. As a first step, we set out the relevant QCRA values that were submitted at PA, in £ 2025/25 in the table below.

Table 1 BTNO QCRA results

Percentile	£ 24/25	As % of pre-risk costs submitted
P50		9.0%
P70		10.2%
P80		11.0%
Mean		9.2%

Source: NGET analysis, BTNO PA submission

Note: Pre-risk costs submitted at PA is in 24/25 prices

It is also important to note that a parsimonious approach was applied to the BTNO QCRA and it excluded consideration of¹⁵

¹⁴ See for example Contingency Guidelines Risk Engineering Society (RES) and Engineers Australia Author: Pedram Danesh-Mand, 2nd Edition, 2019. Heathrow, the Department for Transport, Network Rail, Highways England and Crossrail appear to have adopted similar principles.

¹⁵

- Estimation uncertainty
- Correlation between risks¹⁶
- Joint confidence levels
- Prolongation¹⁷

NGET has taken such an approach because it is in line with previous submissions, and in doing so the QCRA remains transparent and tractable. The implication is a higher percentile on the QCRA distribution needs to be selected, to account for unknown or unquantifiable risks. The starting point is the QCRA mean of [REDACTED] (corresponding to a 9.2% uplift as a percentage of pre-risk costs). In line with industry convention, NGET has requested funding at a level above the median and selected a confidence level at the 70th percentile of its submitted QCRA.¹⁸

RCF evidence can be considered to provide some comfort regarding the uplift of risk allowance above the mean value of [REDACTED] (which is the theoretical minimum level of risk funding). BTNO's project-specific RCF implies mean uplift implied of 8% above pre-risk costs.

In this case, the project-specific RCF does not provide additional guidance on how to set the risk allowance above the minimum (a 9.2% uplift on pre-risk project costs). The Mean value of the QCRA of [REDACTED] remains a lower bound, but is still likely to be insufficient, given the factors excluded from the QCRA analysis.¹⁹ As the RCF output can either underestimate or overestimate the true value for risk funding required, RCF should not substitute the QCRA-led estimate, where the QCRA mean is the lower bound.

Taking all of this together, it may be prudent for Ofgem to adhere to its precedent of setting risk allowances at the 70th percentile of the QCRA distribution. This level of risk funding has been provided for previous projects, where risk allowances have been appraised using a similar approach and evidence.

If Ofgem intends to move to funding the Pmean of the QCRA distribution in future, the corresponding analysis would need to be more comprehensive, and include as much known and unknown risks including inter-project correlation within a wider portfolio, among other things. For BTNO, this level of rigour has not been sought in the QCRA analysis as Ofgem had not signalled a change to its approach prior to NGET making its PA submissions.

¹⁶ [REDACTED]

¹⁷ [REDACTED]

¹⁸ [REDACTED]

¹⁹ The approach to the BTNO QCRA is in line with past submissions which Ofgem had previously accepted. However this was accompanied by a decision to select a higher confidence level than the P50 on the QCRA distribution.

We note in the case of BTNO, the RCF cross-check does not provide additional information on the required uplift over the QCRA mean. However, we consider that relying on a wider base of evidence – including RCF - can only be beneficial for arriving at a more informed view of the appropriate risk allowance going forward, as long as we remain cognisant of the strengths and limitations of different kinds of evidence.

5 Further considerations and overall conclusions

This paper has shown that setting a risk allowance equal the P50 of the QCRA is likely to result in systematic underfunding, and a fair level of funding should be set at the expected outcome of the true risk distribution, however this is not directly observable and not straightforward to estimate.

Firstly, relying on a P50 of the QCRA is likely to result in systematic underfunding of risk. Risks on major projects are unlikely to be symmetric, which means the mean of the distribution is likely to exceed the P50. As the P50 does not reflect the skewedness of the risk distribution, it does not reflect the expected value of the risk distribution. Therefore, the appropriate the starting point for setting the risk allowance should be the mean of the QCRA distribution submitted.

However, NGET has historically adopted a relatively parsimonious approach to the QCRA, consistent with past Project Assessments. As a result, the P50 – and even the calculated mean – may understate the true scale of delivery risks, as certain risks are not captured in the QCRA analysis. In recognition of this, NGET has focused on the P50 - P70 range when setting allowances, seeking to reflect some missing factors while avoiding double counting protections already embedded in the ASTI framework (e.g. COAEs).

Additionally, considering a broader evidence base (having regard to the strengths and limitations of each source of evidence) can help Ofgem arrive at a more informed view of the appropriate risk allowance. As such, this paper discusses a cross-check of the QCRA-led risk allowance using RCF outputs, as an additional point of reference to provide comfort on required uplifts above the mean. While the specific RCF analysis for BTNO does not support the uplift proposed in the QCRA analysis, we consider that for other major projects in the pipeline this cross-check may well provide additional support.

If Ofgem intends to switch its methodology for setting the allowances to the expected outcome, then the QCRA (and the quantification of their underlying risk registers) will have to be made more comprehensive. NGET's previous approach (e.g. for EGL1 and OPN2, and indeed BTNO) has drawn pragmatically on wider industry expertise and best practice; but a more targeted assessment of major project risk allowances will necessarily require incremental time, resource and expertise to carry out properly and therefore will need to be signalled well in advance of the PA submission stage.

While not the focus of this paper, a further and highly relevant consideration is portfolio risk. Currently, Ofgem assesses each project independently at PA stage for risk, without fully accounting for interdependencies across major projects. Common risk drivers – such as supply chain disruption – may affect multiple projects simultaneously. These interdependencies have the effect of increasing overall variance (risk) of the major projects portfolio, because it reduces the frequency where with a good outcome in one area is offset by a bad outcome elsewhere, or vice versa.

We understand there is evidence to suggest that funding risk at each project's median QCRA is insufficient; but, even if Ofgem could identify the 'true' mean risk funding for each project independently, funding each project at the 'true mean' risk level could still result in underfunding overall.²⁰

However, if Ofgem decides to deviate from industry standard approach of setting the allowance at approximately at a level closer to the P70 of the QCRA distribution, and to focus the allowance on a measurement of the expected outcome, then a lot more work needs to be done to make the QCRA analysis more robust and fit for purpose. Ofgem has not done this for the BTNO project, and therefore cannot expect the P50 or indeed even the Pmean of the QCRA of the BTNO PA submission to be a fair level of allowance because they are not reflective of the expected cost of the project.

²⁰ Evidence from OGP shows this principle is recognised by a number of other infrastructure / major project delivery bodies such as Highways England (2024), the Defence sector, Network Rail. Heathrow, Department for transport.