
9. Aggregating and Comparing Costs and Benefits

Once costs and benefits of different regulatory options have been identified and estimated, the next critical step in RIA is their aggregation and comparison. This stage is central to decision-making, as it determines whether a regulatory option produces a **net positive impact** on society, and how that impact is distributed across stakeholders.

This step can be thought of as answering the question:

What is the likely net benefit (benefit minus cost) of each option?

RIA frameworks across jurisdictions emphasise that the purpose of aggregation is not merely arithmetic consolidation, but structured comparison across options, enabling policymakers to identify the alternative that delivers the greatest net benefit. As noted in a policy primer by the Mercatus Center, RIA requires a comparison of alternatives “so that the alternative with the most net benefits can be identified.”

Similarly, guidance from the OECD stresses that policymakers must assess costs and benefits across stakeholders to determine whether regulatory interventions enhance overall welfare and justify their associated burdens.

9.1. Conceptual Foundations of Aggregation

Aggregation in RIA operates at two levels. First, it involves consolidating impacts across different categories—monetary, quantitative, and qualitative. Second, it involves aggregating across stakeholders—such as firms, consumers, government, and broader society.

At a conceptual level, aggregation is grounded in welfare economics, where policy choice is framed as selecting the option that maximises net social benefit (i.e., total benefits minus total costs). It must account not only for total magnitude but also for distributional effects, as costs and benefits may accrue to different stakeholders. UK and EU guidance explicitly require that RIAs consider how impacts are distributed “within and between groups,” highlighting that aggregation is not purely an exercise in summation but also in distributional assessment.

9.1.1. Aggregation at the Stakeholder Level

One approach to aggregation involves first analysing costs and benefits at the level of individual stakeholder groups and then comparing net impacts across these groups.

Under this approach, regulators estimate the net impact for each stakeholder category—for example:

- regulated firms (compliance costs vs efficiency gains)
- consumers (price reductions vs reduced choice, if any)
- government (administrative costs vs improved policy outcomes)

This method is particularly useful where impacts are unevenly distributed or where trade-offs between stakeholders are central to policy choice.

The approach is reflected in the framework used by the Prudential Regulation Authority, which explicitly analyses costs and benefits at multiple levels—firm-level, market-level, and macroeconomic level—before considering their overall effect.

For instance, prudential regulation may impose direct compliance costs on financial institutions but generate benefits for depositors, policyholders, and the broader economy through increased financial stability. The PRA recognises that such impacts may involve transfers between stakeholders, where costs borne by one group may correspond to benefits for another.

This stakeholder-level aggregation is particularly valuable in sectors such as finance, telecoms, and competition law, where regulatory interventions often involve redistribution of surplus rather than purely net creation of welfare.

Telecom Market Intervention (Inspired by UK Smart Data Impact Assessment)

Assume a regulatory option that reduces switching barriers and improves transparency.

TABLE 9.1

Stakeholder	Costs (₹ million/year)	Benefits (₹ million/year)	Net Impact
Telecom Firms	500 (compliance, IT systems)	200 (efficiency gains, reduced churn costs)	-300
Consumers	50 (time spent learning system)	1,200 (price savings, better plans)	+1,150
Small Businesses	30 (transition costs)	400 (reduced overpayment, better contracts)	+370
Government	100 (administrative and enforcement costs)	150 (better market outcomes, tax gains)	+50

9.1.2. Aggregation at the Aggregate (Economy-Wide) Level

An alternative and widely used approach is to aggregate all costs across stakeholders and all benefits across stakeholders to determine a single net impact at the societal level.

This approach is characteristic of standard cost-benefit analysis, where $\text{Net Social Benefit} = \text{sum of Total Benefits on all stakeholders} - \text{sum of Total Costs on all stakeholders}$

Under this framework, a regulatory option is considered desirable if aggregate benefits exceed aggregate costs, even if some stakeholders are worse off.

In the United Kingdom, regulatory proposals are often summarised using the Equivalent Annual Net Cost to Business (EANCB) metric, which converts impacts into an annualised net figure to facilitate comparison across options. Similarly, EU-level RIAs aim to demonstrate that “the benefits resulting from the proposal justify the costs” and that the regulation produces a positive net impact on society as a whole.

This aggregate approach is particularly useful where policymakers need a single decision metric, but it must be complemented by stakeholder-level analysis to ensure that important distributional effects are not obscured.

9.1.3. Matching Types of Costs and Benefits

A critical methodological principle in aggregation is that like must be compared with like. This means that:

- Monetary costs should be compared with monetary benefits
- Quantified (non-monetary) costs should be compared with quantified benefits
- Qualitative costs should be compared with qualitative benefits

This principle ensures analytical consistency and avoids misleading comparisons.

RIA guidance across jurisdictions emphasises that impacts should be monetised wherever feasible, but where monetisation is not possible, they should at least be quantified or qualitatively described.

In practice, this often leads to a layered presentation of results, where:

- monetised impacts are aggregated into net present value or annualised figures
- quantified impacts (e.g., number of consumers benefiting) are presented alongside
- qualitative impacts (e.g., trust, fairness, innovation) are discussed narratively

This layered approach ensures that decision-making is informed by the full spectrum of impacts, rather than being constrained by what can be monetised.

In practice, RIAs across jurisdictions adopt a combination of these approaches.

In the United Kingdom, Impact Assessments routinely present a summary table of costs and benefits, including monetised impacts (e.g., compliance costs, cost savings), alongside qualitative impacts (e.g., improved consumer confidence). The use of EANCB provides a standardised metric for comparing options.

In EU practice, RIAs emphasise demonstrating that benefits outweigh costs, while also identifying who gains and who loses, thereby combining aggregate and distributional analysis.

The PRA's cost-benefit framework provides a particularly sophisticated example, tracing impacts from firm-level effects to market-level and macroeconomic outcomes, and recognising that some impacts may offset each other across stakeholders.

These examples highlight that aggregation is not a purely mechanical exercise but involves judgment, transparency, and careful interpretation of evidence.

9.2. Key Analytical Considerations and Challenges

Several challenges arise in aggregating costs and benefits.

First, there is the risk of double counting, particularly where the same impact is captured through multiple channels (e.g., time savings and productivity gains). Care must be taken to ensure that each impact is counted only once.

Second, aggregation may obscure distributional impacts, particularly where aggregate net benefits mask significant losses for specific groups. This underscores the importance of presenting both aggregate and stakeholder-level results.

Third, some impacts represent transfers rather than net welfare gains. For example, a regulation that shifts surplus from firms to consumers may produce distributional effects without changing total welfare. Regulators must distinguish such transfers from genuine efficiency gains.

Finally, uncertainty in estimates must be transparently addressed, often through sensitivity analysis or scenario-based aggregation.

In synthesis, aggregation of costs and benefits is the culmination of the RIA process, translating detailed analysis into decision-relevant conclusions. It requires both technical rigour and analytical judgment, combining monetised, quantified, and qualitative impacts across stakeholders.

A robust approach to aggregation typically involves:

- analysing impacts at the stakeholder level
- aggregating impacts at the societal level
- ensuring consistency across types of impacts

When undertaken carefully, aggregation enables policymakers to determine not only whether a regulatory option is beneficial in aggregate, but also how those benefits and costs are distributed, thereby supporting more transparent, accountable, and evidence-based decision-making.

9.3. Maintaining proportionate boundaries to RIA analysis

RIA should be sufficiently comprehensive to capture the material consequences of the regulatory options under consideration, but it should not become an open-ended exercise in tracing every conceivable consequence of an intervention. A regulatory measure may produce a first-order effect on the parties directly subject to it, followed by second-order effects on suppliers, consumers, competitors and workers, and then further effects on investment, prices, employment, innovation, household expenditure or other markets. In principle, this chain can continue almost indefinitely. If every possible downstream consequence is required to be identified, quantified and monetised, the RIA may become unending, while the additional analysis may contribute little to the actual regulatory decision.

This is one reason why proportionality is a fundamental principle of RIA. Proportionality does not mean ignoring indirect or unintended consequences. Rather, it requires the regulator to distinguish between effects that are material to the decision and effects that are sufficiently remote, uncertain or insignificant that analysing them in greater detail would not materially change the assessment of the available options.

The European Commission recognises the need to identify the relevant impacts without attempting to analyse every possible effect with the same degree of detail. This is particularly important because regulatory proposals can have economic, social and environmental consequences extending across different groups and over different time horizons. The purpose of proportionality is therefore to ensure that analytical resources are concentrated on those consequences that are most relevant to the decision.

The regulator should consider the strength of the causal relationship, the likelihood and magnitude of the effect, the availability and reliability of evidence, and whether including the effect is likely to alter the comparison between regulatory options.

First-order effects will ordinarily warrant detailed analysis. These may include compliance costs imposed directly on regulated entities, price changes experienced by consumers, changes in access to a service, or administrative costs incurred by the government. Second-order effects may also require analysis where there is a credible causal pathway and a reasonable likelihood that their magnitude could materially affect the decision. These could include scenarios where compliance costs are likely to be passed through into consumer prices or where a regulatory requirement could materially affect market entry. By contrast, effects that depend on several uncertain behavioural or economic assumptions and are unlikely to affect the relative ranking of the options may appropriately be recorded qualitatively rather than subjected to increasingly elaborate modelling.

This does not mean that regulators should ignore longer causal chains. Material indirect and unintended consequences should remain within the scope of RIA, particularly where they could change the preferred option. The appropriate response to uncertainty is often to record the effect, explain the causal mechanism, assess its likely direction and significance, and, where feasible, test it through sensitivity or scenario analysis. There is little value, however, in attempting to attach increasingly precise monetary estimates to an effect when the underlying causal relationship itself is highly uncertain. False precision can make an RIA appear more rigorous while actually making the decision-making process less transparent.

9.4. Transparency and justification of assumptions and parameters

The conclusions of an RIA are influenced by the assumptions and analytical parameters used to translate that evidence into estimates of costs, benefits and overall impacts. Two assessments of the same regulatory proposal can therefore produce materially different results if they make different assumptions about, for example, the number of affected stakeholders, the extent to which firms pass compliance costs on to consumers, the behavioural response of consumers, the probability of an adverse outcome, the time period over which costs and benefits accrue, or the discount rate applied to future impacts. This does not necessarily mean that one assessment is erroneous. It means that regulatory analysis involves judgements that need to be made explicit and open to scrutiny.

This is particularly important because RIA is undertaken before the full effects of a regulation can be observed. Forecasts of behavioural responses, compliance costs and benefits necessarily involve assumptions about how markets and stakeholders will respond to the proposed intervention.

For instance, when a regulation generates costs immediately but benefits over a long period, the choice of discount rate can materially affect the present value of those benefits and, consequently, the net present value of the regulatory option. A higher discount rate places less weight on future benefits; a lower rate places relatively greater weight on them. The choice of rate is therefore not a technical detail that can simply be buried in a spreadsheet. It should be stated explicitly, justified by reference to the applicable appraisal guidance, and, where the conclusion is sensitive to the choice, subjected to sensitivity analysis.

The US Office of Management and Budget require regulatory analysis to report the assumptions and analytical methods used and, for discounting, provides recommended discount rates while recognising that different regulatory effects may require particular treatment. The implication for RIA is that the regulator should not simply report a single NPV as though it were an objective fact, it should enable the reader to understand how that result was generated and how robust it is to reasonable alternative assumptions.

A similar issue arises in estimating behavioural responses. Suppose a proposed disclosure requirement is expected to improve consumer decision-making. The estimated benefit may depend on an assumption about the proportion of consumers who will read the information, understand it and act upon it. If the RIA assumes that 80 per cent of consumers respond to the information, the estimated benefit could be substantially larger than if the regulator assumes a 20 per cent response rate. Neither assumption should be presented without explanation. The regulator should provide the empirical basis for the assumption, indicate whether it is derived from historical evidence, experimental evidence, stakeholder evidence or expert judgement, and, where appropriate, show how the results change under alternative response rates.

The same principle becomes particularly important when monetising impacts that cannot be directly observed in markets. Estimates of the value of time, willingness to pay, consumer

surplus, health and wellbeing effects, or reductions in risk may depend on methodological choices and parameters. For example, a regulator may use survey evidence to estimate consumers' willingness to pay for an improvement in service quality. The resulting benefit estimate will depend not only on the responses received but also on the survey design, the valuation method, the population to which the results are extrapolated and the assumptions used to translate stated preferences into aggregate benefits. Such methodological choices should therefore be disclosed rather than presented as though the resulting monetary value were directly observable.

This suggests that an RIA should distinguish between evidence-based parameters and normative choices. Some parameters may be relatively straightforward empirical estimates, for example, the number of regulated entities or the observed average cost of compliance. Others necessarily involve judgement, for example, the assumed probability of a behavioural response, the valuation of an avoided harm, the weighting assigned to distributional impacts, or the appropriate time horizon. Both can legitimately be used in RIA, but they should not be presented as having the same evidentiary status. Making this distinction explicit enables decision-makers to understand where the assessment rests on relatively strong empirical evidence and where it depends more heavily on judgement.

The appropriate response is not to eliminate assumptions, this would generally be impossible in an ex ante assessment, but to make them visible, justify them and test their importance. A well-designed RIA should therefore identify the principal assumptions used in estimating costs and benefits, explain the source and rationale for each important parameter, distinguish observed data from assumptions and expert judgements, and identify which assumptions are most influential to the final conclusion.

9.5. Different levels and forms of regulatory analysis

RIA need not always involve the same degree of analytical sophistication. A proportionate RIA framework recognises that the resources devoted to analysing a regulatory proposal should reflect the scale, significance, complexity and uncertainty of its likely impacts. A major regulation affecting millions of consumers or imposing substantial compliance costs may warrant detailed cost-benefit analysis, extensive modelling and sensitivity analysis. A relatively minor or routine regulatory amendment may not justify the same degree of analysis. The objective is therefore not to maximise the quantity of analysis undertaken, but to ensure that the analysis is sufficient to support a sound decision.

One approach used in some developing-country regulatory systems is “RIA Lite”. The term is not a universally defined methodology, but generally refers to a proportionate assessment that retains the core logic of RIA—identifying the problem, considering alternatives, identifying affected stakeholders and assessing likely costs and benefits—while relying more heavily on existing evidence, qualitative analysis, targeted stakeholder interviews and readily available quantitative information rather than undertaking a full-scale monetised cost-benefit analysis. This can be particularly useful where the proposed regulatory change is relatively limited, where the likely impacts are modest, or where institutional and data constraints make a full RIA disproportionately costly.

RIA Lite should not, however, be interpreted as permitting important regulatory consequences to be ignored. The key distinction is depth rather than analytical principle. A Lite assessment should still ask what problem is being addressed, whether government intervention is warranted, what alternatives are available, who will be affected, what the principal costs and benefits are, and whether the preferred option is proportionate. If the preliminary assessment reveals substantial, uncertain or widely distributed impacts, the regulator should escalate the exercise to a more comprehensive RIA rather than continuing to rely on a simplified methodology.

A second useful approach is break-even analysis. This is particularly valuable where the regulator can estimate the costs of an intervention reasonably well but cannot credibly monetise some or all of its benefits. Rather than assigning a potentially arbitrary monetary value to the benefits, break-even analysis asks what level of benefit would be required for the intervention's benefits to equal its costs. The result can be expressed in terms that are meaningful for the particular intervention, for example, the number of accidents that would need to be prevented, the number of consumers who would need to switch providers, the number of illnesses avoided, or the amount of additional investment generated for the intervention to break even.

For instance, in an impact assessment concerning data portability and consumer switching, when it is difficult to monetise all of the benefits of the policy, a break-even analysis can help estimate the number of consumers who would need to switch suppliers for the resulting consumer welfare gains to offset the costs of the policy.

Break-even analysis is also useful where the regulator has a reasonably good estimate of the cost side but faces substantial uncertainty about the probability or magnitude of harm avoided. Suppose a cybersecurity regulation costs regulated firms ₹100 crore annually, while the regulator cannot reliably estimate the monetary value of avoided cyberattacks. Instead of producing a speculative estimate of that value, the regulator could calculate how many attacks, or how much aggregate loss, would need to be prevented for the regulation to break even. The resulting threshold can then be compared with historical attack frequencies and available evidence on losses. The analysis does not prove that the regulation will generate positive net benefits; rather, it allows the decision-maker to assess whether the required level of benefit appears plausible.

References

1. Mercatus Center, *A Primer on Regulatory Impact Analysis*, <https://www.mercatus.org/research/policy-briefs/primer-regulatory-impact-analysis>
2. OECD (2020), *Regulatory Impact Assessment: OECD Best Practice Principles*, <https://doi.org/10.1787/7a9638cb-en>
3. OECD, *Better Regulation Practices across the EU*, https://www.oecd.org/en/publications/better-regulation-practices-across-the-european-union-2022_6e4b095d-en/full-report/component-9.html
4. Scottish Government, *BRIA Toolkit*, <https://www.gov.scot/publications/business-regulatory-impact-assessment-bria-toolkit/pages/7/>
5. UK Government, *Validated costs and benefits of regulation*, <https://www.gov.uk/government/publications/validated-costs-and-benefits-to-business-of-regulations>
6. Bank of England / PRA, *Approach to Cost Benefit Analysis*, <https://www.bankofengland.co.uk/prudential-regulation/publication/2026/april/pr-a-approach-to-cost-benefit-analysis-statement-of-policy>