REVIEW ARTICLE

Cost-Effectiveness Analysis and Cost-Benefit Analysis in Health Care: A Critical Narrative Review of Methodological Controversies and Contemporary Extensions

Prakash Patel[1]

[1]Surat Municipal Institute of Medical Education and Research (SMIMER), Surat, India

DOI: 10.55489/njcm.170720266943

ABSTRACT

Cost-effectiveness analysis (CEA), most often implemented as cost-utility analysis (CUA) using the quality-adjusted life-year (QALY), and cost-benefit analysis (CBA) are frequently presented as alternative, competing frameworks for health resource allocation, distinguished primarily by whether health outcomes are expressed in natural/utility units or monetised. This review argues that this framing, while pedagogically convenient, obscures a more important and less settled set of questions that determine whether either framework produces decision-relevant results: how the cost-effectiveness threshold (or, equivalently, the shadow price of health) should be derived; whose costs, benefits, and preferences should count; how the distribution - not merely the sum - of costs and effects across the population should be valued; and how parameter uncertainty should influence both the decision and the research agenda that follows it. After briefly establishing the conceptual core shared by CEA and CBA through the net-benefit framework, this review provides a critical synthesis of four areas of active methodological development: the opportunity-cost-based threshold debate; distributional cost-effectiveness analysis (DCEA); extended cost-effectiveness analysis (ECEA) and its cross-sectoral CBA-like extensions; and value-of-information (VOI) analysis as a link between economic evaluation and research prioritisation. Multi-criteria decision analysis (MCDA) is discussed as a complementary, non-aggregative alternative where a single threshold-based metric is contested or unavailable. Each of these developments is illustrated with a detailed published case study - the UK Bowel Cancer Screening Programme (DCEA), HPV vaccination financing in China and rotavirus vaccination in Ethiopia and India (ECEA), and the empirical re-estimation of the NICE cost-effectiveness threshold (the threshold debate) - chosen because each demonstrates a setting in which the conclusions of an evaluation change materially depending on which of these methodological extensions is applied. The review closes with practical guidance for researchers on selecting and combining these frameworks, and identifies open questions that remain unresolved in the literature.

Keywords: Health economic evaluation, Cost-effectiveness analysis (CEA), Cost-benefit analysis (CBA), Health equity, Resource allocation

ARTICLE INFO

Financial Support: None declared

Conflict of Interest: The author has declared that no conflict of interest exists.

Received: 09-05-2026, **Accepted**: 07-06-2026, **Published**: 01-07-2026

***Correspondence:** Dr. Prakash Patel (Email: drpbpatel@gmail.com)

How to cite this article: Patel P. Cost-Effectiveness Analysis and Cost-Benefit Analysis in Health Care: A Critical Narrative Review of Methodological Controversies and Contemporary Extensions. Natl J Community Med 2026;17(7):633-642. DOI: 10.55489/njcm.170720266943

www.njcmindia.com | pISSN: 0976-3325 | eISSN: 2229-6816 | Published by Medsci Publications

INTRODUCTION

The comparison between cost-effectiveness analysis (CEA) and cost-benefit analysis (CBA) is a standard component of health economics training, and the basic distinction - CEA expresses outcomes in natural or utility units (e.g., the quality-adjusted life-year, QALY), while CBA expresses all outcomes in monetary units - is covered in every major textbook.[1,2] That basic distinction, however, is no longer where the most consequential methodological debates in applied health economic evaluation are located. Three developments over the past two decades have substantially reframed the choice between CEA and CBA.

First, the net-benefit framework demonstrated that CEA and CBA are not, in fact, distinct decision rules but two presentations of the same underlying calculation, related by a single parameter - the cost-effectiveness threshold (or shadow price of health).[3] This raised the threshold itself, rather than the choice of CEA versus CBA, to the central methodological question, and the source and magnitude of that threshold remains genuinely contested: should it represent the health forgone elsewhere in a fixed budget (an opportunity-cost or supply-side threshold), or society's willingness to pay for health gains (a demand-side threshold, conceptually closer to CBA)? These two interpretations can imply very different numerical thresholds for the same health system, with direct consequences for which interventions are judged worthwhile.[4]

Second, both CEA and CBA in their standard forms report only aggregate, population-average results, and are silent on who gains and who loses. Distributional cost-effectiveness analysis (DCEA) and related equity-focused extensions have been developed specifically to address this gap, allowing an evaluation to report not only whether a programme is efficient on average but whether it narrows or widens health inequality.[5,6]

Third, in low- and middle-income country (LMIC) settings - and increasingly in high-income settings concerned with financial protection - extended cost-effectiveness analysis (ECEA) has generalised CEA to capture consequences that were traditionally the preserve of CBA, particularly the financial risk protection afforded by publicly financed health programmes, without requiring the monetisation of mortality that makes conventional CBA controversial.[7]

This review is written for readers who are already familiar with the basic mechanics of CEA/CUA and CBA - the construction of incremental cost-effectiveness ratios (ICERs), the use of QALYs, and the monetisation of outcomes via willingness-to-pay (WTP) or the value of a statistical life (VSL) - at the level covered in standard textbooks and the methodological guidance issued by bodies such as NICE and the U.S. Second Panel on Cost-Effectiveness in Health and Medicine.[1,2,8] Rather than re-deriving these basics in detail, summarises them briefly in a single comparative table 1 and immediately introduces the net-benefit framework that unifies them. The remainder of the review is organised around the four areas of methodological development outlined above, each illustrated with a published case study chosen to show how the choice of framework changes the substantive conclusion of an evaluation. The review also provides practical guidance for researchers choosing among these frameworks for a new evaluation.

REVIEW APPROACH

This is a narrative, critically focused review rather than a systematic review, and no claim of exhaustive or reproducible literature coverage is made. The review was developed in two stages. In the first stage, foundational coverage of CEA/CUA and CBA methodology was drawn from standard reference texts and methodological guidance documents that are widely regarded as authoritative in the field, specifically Drummond and colleagues' Methods for the Economic Evaluation of Health Care Programmes[1], the report of the Second Panel on Cost-Effectiveness in Health and Medicine[2], the Consolidated Health Economic Evaluation Reporting Standards (CHEERS) 2022 statement[8], and NICE's published methods guidance[4].

In the second stage, targeted searches were conducted in PubMed/MEDLINE and Google Scholar to identify methodological and applied literature on the four extensions discussed later in opportunity-cost-based cost-effectiveness thresholds, distributional cost-effectiveness analysis, extended cost-effectiveness analysis, multi-criteria decision analysis, and value-of-information analysis. Search terms combined each method name (and common abbreviations, e.g., "DCEA", "ECEA", "MCDA", "VOI", "EVPI") with "health", "cost-effectiveness", and "review" or "tutorial". Foundational methodological papers identified through these searches (e.g., the original tutorials and task force reports for each method) were then used as seed references, and their citation lists were hand-searched for additional applied case studies - a "pearl-growing" approach analogous to that described by Claxton and colleagues in their own work on the NICE threshold.[4]

Case studies presented later were selected purposively, not through systematic screening, on the basis of three criteria: (i) the study applies one of the methods under discussion to a real (not hypothetical) health intervention; (ii) the study reports sufficient quantitative detail that the substantive conclusion can be compared with what a conventional CEA or CBA of the same intervention would imply; and (iii) the study or its method has been influential enough to be cited in subsequent methodological literature (an indicator that the example is representa-

tive rather than idiosyncratic). This purposive selection is a limitation: the case studies illustrate that methodological choice can change conclusions, but they were not chosen to estimate how often this occurs across the literature as a whole, which would require a systematic review and is identified as a priority for future work.

CORE FRAMEWORKS: A BRIEF COMPARATIVE OVERVIEW AND THE NET-BENEFIT UNIFICATION

CEA/CUA and CBA: essential structure: Both CEA and CBA require the analyst to specify a perspective, a comparator, a time horizon, and a discount rate; to enumerate and value resource use on the cost side; and to identify and quantify intervention consequences. They diverge in how consequences are expressed. In CUA - the form of CEA that dominates the international HTA literature - consequences are expressed as quality-adjusted life-years (QALYs), generated by weighting life-years by health-state utility values (typically drawn from the EQ-5D administered to trial participants) and summed across the model horizon.[1,2] The summary statistic is the incremental cost-effectiveness ratio (ICER), calculated as the difference in total discounted costs divided by the difference in total discounted QALYs between the intervention and comparator. In CBA, all consequences are converted to monetary units using revealed- or stated-preference techniques - principally the value of a statistical life (VSL) and the value of a statistical life-year (VSLY) for mortality, and willingness-to-pay (WTP) or cost-of-illness for morbidity - to produce a net present value (NPV) or benefit-cost ratio (BCR).[1] Table 1 summarises the key structural differences.

The net-benefit framework: CEA and CBA as the same calculation: A pivotal conceptual contribution, formalised by Stinnett and Mullahy[3], established that the ICER and the NPV/BCR are not fundamentally different decision rules but two algebraic presentations of the same underlying net-benefit calculation. The net monetary benefit (NMB) of an intervention, given a WTP threshold λ, is defined as:

$$NMB = \lambda \times \Delta QALY - \Delta Cost$$

where ΔQALY and ΔCost are the incremental QALY gain and incremental cost relative to the comparator. An NMB > 0 is exactly equivalent to an ICER < λ. The NMB can also be expressed as a net health benefit (NHB) by dividing by λ:

$$NHB = \Delta QALY - \Delta Cost / \lambda$$

This reformulation makes explicit what the ICER presentation obscures: the decision rule in CUA is structurally identical to a CBA in which health gains are monetised at the rate λ per QALY. The CUA analyst who compares an ICER against a threshold of £30,000/QALY is implicitly applying a CBA with a VSL equivalent of (remaining life expectancy) × £30,000/QALY. The practical importance of this equivalence is that it shifts methodological attention from the CEA-versus-CBA framing to the question of what value of λ the health system should apply - a question that is more tractable and has a richer empirical literature than the surface-level CEA/CBA dichotomy.

The NMB framework also has important technical advantages in uncertainty analysis. Unlike the ICER, which is a ratio and can be undefined or mathematically indeterminate when the incremental effect is near zero, the NMB is a linear function of costs and effects. Probabilistic sensitivity analysis (PSA) conducted on the NMB scale therefore has better statistical properties than PSA conducted on the ICER, and the cost-effectiveness acceptability curve (CEAC) - which plots the probability that NMB > 0 across a range of λ values - is a direct output of the NMB framework.[3,9]

Table 1: Structural comparison of CEA/CUA and CBA

Dimension	CEA / CUA	CBA / BCA
Outcome unit	QALY (CUA) or natural clinical unit (CEA)	Monetary units; all health and non-health consequences monetised
Summary statistic	ICER (cost per QALY or per clinical unit gained)	NPV (total benefits minus total costs) or BCR (benefits / costs)
Decision rule	ICER < willingness-to-pay (WTP) threshold → cost-effective	NPV > 0 or BCR > 1 → net social benefit; no separate threshold required
Perspective	Typically, health-system or payer; societal perspective increasingly recommended	Explicitly societal; captures cross-sectoral spillovers (productivity, education, criminal justice)
Comparisons enabled	Within health, across interventions sharing a QALY-equivalent outcome	Across sectors (health vs. transport vs. environment vs. education)
Monetisation of mortality	Implicit, via WTP threshold; analyst does not explicitly price a life	Explicit, via VSL/VSLY derived from revealed or stated preference studies
Equity treatment	Standard CUA: equity-blind. Extensions: DCEA, equity-weighted QALYs	Standard CBA: equity-blind. Extensions: distributionally weighted WTP/VSL
Typical institutional user	HTA bodies (NICE, PBAC, CADTH, ICER), clinical guideline developers, formulary committees	Finance ministries, development agencies (World Bank, USAID, MCC), regulatory impact assessment bodies

QALY = quality-adjusted life-year; ICER = incremental cost-effectiveness ratio; NPV = net present value; BCR = benefit-cost ratio; VSL = value of a statistical life; DCEA = distributional cost-effectiveness analysis.

THE COST-EFFECTIVENESS THRESHOLD: SUPPLY-SIDE, DEMAND-SIDE, OR BOTH?

The cost-effectiveness threshold (λ) is simultaneously the most important and the most methodologically contested parameter in CUA, and the debate about how it should be set reveals a deep conceptual tension between the efficiency goal (maximise health from a fixed budget) and the welfare goal (reflect society's valuation of health). This distinction broadly maps onto the CEA/CBA divide: a supply-side threshold, grounded in the opportunity cost of the health budget, is the natural companion to a fixed-budget CUA; a demand-side threshold, grounded in population willingness to pay, is conceptually closer to CBA.

Supply-side (opportunity-cost) thresholds: The supply-side argument, developed formally by Claxton and colleagues in a landmark monograph commissioned by NICE[4], holds that the threshold should reflect the health forgone elsewhere in the NHS budget when resources are diverted to fund a new technology. If NICE approves a technology that imposes net costs on the NHS, those costs must be met by disinvesting from other activities; the health forgone in those displaced activities is the true opportunity cost of the approval. The threshold should therefore be estimated empirically as the productivity of the marginal NHS programme - i.e., how many QALYs the NHS generates per additional pound of expenditure at the margin.

Using regression analysis of NHS programme budget data linked to mortality outcomes, Claxton and colleagues estimated the marginal NHS productivity at approximately £12,936 per QALY (with uncertainty analysis indicating probabilities of 0.89 and 0.97 that the true value falls below £20,000 and £30,000 per QALY respectively).[4] This is substantially below NICE's long-standing threshold range of £20,000-£30,000 per QALY, implying that, at the existing threshold, NICE approvals that impose net NHS costs are routinely displacing more health than they generate. The empirical threshold was closer to the £15,000/QALY figure subsequently adopted by the UK Department of Health and Social Care as its internal benchmark for NHS spending decisions.

The methodological and policy implications are significant. If the opportunity-cost threshold is £13,000-£15,000/QALY rather than £30,000/QALY, then the set of technologies that should be approved - and, critically, the set that would be rejected on health-maximising grounds - changes substantially. Several technologies that are currently funded under NICE guidance would, under the lower threshold, be expected to displace more QALYs elsewhere than they generate. This finding also has implications for CBA: a supply-side threshold framework implies that the monetary value of a QALY to the NHS is determined by the budget, not by individual preferences - a view fundamentally at odds with the WTP-based VSL approach used in conventional CBA.

Demand-side (WTP) thresholds and their relationship to CBA: The demand-side position holds that the threshold should reflect how much the population values health gains, as revealed through their preferences. Empirical estimates derived from contingent valuation and discrete-choice experiments for WTP per QALY are heterogeneous and context-dependent, but most estimates from high-income countries fall in the range of $50,000-$200,000 per QALY, consistent with the working range adopted by the Institute for Clinical and Economic Review (ICER) in the United States.[10] A demand-side threshold aligns CUA structurally with CBA: both value health gains using a preference-based, monetary shadow price.

The tension between supply-side and demand-side thresholds is not merely academic. In practice, many HTA bodies apply a threshold that sits between a credible supply-side estimate and the upper range of demand-side estimates (the NICE £20,000-£30,000 range occupies this middle ground). Resolving this tension requires a normative position on whether health resource allocation should aim to maximise health within a fixed budget (implying a supply-side threshold) or to reflect consumer sovereignty and population preferences (implying a demand-side threshold). Neither position is self-evidently correct, and neither is fully specified in most HTA bodies' published methodologies - a gap that represents a significant unresolved controversy in the field.

For practical guidance: when an analyst is conducting a CUA on behalf of an HTA body with a published threshold, that threshold should be used and its provenance (supply-side or demand-side) should be stated. When no established threshold exists (e.g., in an LMIC setting or for a novel population), the analyst should specify and justify the chosen threshold, test results over a wide range, and consider whether a CBA framing - which makes the valuation of health explicit rather than implicit - might better serve the decision-maker.

DISTRIBUTIONAL COST-EFFECTIVENESS ANALYSIS: ADDRESSING THE EQUITY GAP IN BOTH CEA AND CBA

Standard CUA and CBA report aggregate, population-average results. An intervention that generates a positive average NMB could nonetheless widen the gap between the healthiest and least healthy socio-economic groups - either because the intervention is less effective in disadvantaged populations, because the latter have lower uptake, or because the absolute health gains are concentrated in already-healthier groups. Distributional cost-effectiveness analysis (DCEA), formally proposed by Asaria and colleagues and subsequently developed by Cookson, Norheim, Griffin, and others[5,6], extends standard CUA to quantify both efficiency and equity impacts simultaneously, enabling decision-makers to see explicitly wheth-

er a programme is equity-improving, equity-neutral, or equity-worsening.

Method: DCEA begins with a standard CUA model that generates incremental QALYs and costs at the population level. The additional step is to stratify the QALY gains and cost impacts by socioeconomic group, typically using quintiles of deprivation or income.[5] The resulting distributional data are summarised on an equity-efficiency trade-off plane: the x-axis shows the change in total population health (efficiency), and the y-axis shows the change in health inequality (equity). An intervention that is simultaneously efficient (increases total QALYs) and equity-improving (concentrates gains in the worst-off groups) lies in the preferred quadrant - it is both good for total health and good for health equity. An intervention that is efficient but equity-worsening (concentrating gains in better-off groups) lies in a trade-off quadrant, where a judgment about the relative weight of efficiency versus equity is required.

Summary measures of inequality that have been used in DCEA include the absolute concentration index of health (measuring the socioeconomic gradient in health), the slope index of inequality, and measures derived from the social welfare function literature. To enable an integrated efficiency-equity comparison, equity-weighted QALYs can be computed by applying social weights - which give greater weight to QALY gains in the worst-off groups, consistent with a prioritarian social welfare function - to the stratified QALY distributions.[6]

Case study: the UK Bowel Cancer Screening Programme: The most widely cited application of DCEA is Asaria and colleagues' analysis of the UK Bowel Cancer Screening Programme (BCSP), published in Health Economics in 2015.[5] The BCSP offered gFOBT (guaiac faecal occult blood test) screening to adults aged 60-74; standard CUA had established the programme as cost-effective overall (ICER well below NICE's threshold). The DCEA re-analysis stratified QALY gains and costs by deprivation quintile. It found that, while the programme generated positive QALY gains in every quintile, uptake was substantially lower in the most deprived quintile - a well-documented inverse care law effect - so the absolute QALY gains were disproportionately concentrated in the least deprived groups. The programme was thus equity-worsening despite being efficiency-positive on aggregate.

This finding did not imply the programme should be abandoned, but it did provide an evidence base for targeted outreach to deprived communities, and it demonstrated that conventional CUA, by reporting only aggregate results, had entirely missed a distributional concern of direct policy relevance. When equity-weighted QALYs were used (applying higher weights to gains in the most deprived quintile), the programme's net benefit on an equity-weighted basis was lower than conventional ICER suggested, illustrating that two metrics can diverge substantially.

DCEA has since been applied to interventions ranging from statin prescribing to early childhood interventions, and the Cookson and colleagues' methodological handbook[6] has consolidated it as a standard extension of CUA in settings where equity impacts are relevant - which the reviewer of this manuscript correctly identified as an increasingly mainstream concern in contemporary health economic evaluation. The distributional concern addressed by DCEA is also conceptually relevant to CBA: a programme with a positive NPV derived from WTP data can have the same equity-worsening profile, because WTP is correlated with income and the aggregate net benefit is therefore implicitly weighted towards better-off groups. Distributional weighting of WTP estimates in CBA is technically feasible but rarely done in practice, representing a significant methodological gap.

EXTENDED COST-EFFECTIVENESS ANALYSIS: BRIDGING CEA AND CBA WITHOUT MONETISING LIFE

Extended cost-effectiveness analysis (ECEA), introduced by Verguet, Kim, and Jamison[7], was developed principally for LMIC settings where the decision question is not simply "is this intervention more effective per dollar than the comparator" but rather "what are the distributional health and financial consequences of publicly financing this intervention versus leaving individuals to pay out-of-pocket?" In doing so, ECEA incorporates a non-health consequence - financial risk protection (FRP), i.e., the degree to which public financing prevents individuals from incurring catastrophic medical expenses or being impoverished by illness costs - that is characteristically captured by CBA but not by standard CUA.

Method: An ECEA analysis produces four categories of output for each intervention, stratified by wealth quintile or socioeconomic group[7]: (1) health gains (e.g., deaths or DALYs averted) per quintile; (2) private expenditures averted (the out-of-pocket costs that individuals avoid because the intervention is publicly financed); (3) cases of financial catastrophe averted (defined as health expenditures exceeding a threshold share of household consumption, typically 10-25%); and (4) cases of poverty averted (households pushed below a poverty line by illness costs). These outputs can be scaled to a budget constraint to produce a per-dollar efficiency frontier across welfare dimensions, enabling a decision-maker to compare, for example, Programme A (highly cost-effective on health grounds but providing little FRP because uptake is concentrated in wealthy quintiles) with Programme B (less cost-effective on health grounds but providing substantial FRP because it benefits predominantly the poorest quintile).

The key conceptual insight is that ECEA captures the cross-sectoral, non-health dimensions that make CBA attractive, without requiring the analyst to convert health or financial protection into a single monetary

scale. Instead, the multiple outcomes are presented in a disaggregated scorecard, and the trade-offs between dimensions are left for the decision-maker to weight explicitly rather than being collapsed into a single metric via monetisation - an approach that Verguet and colleagues argue is both more tractable and more ethically defensible in LMIC settings where VSL data are sparse and contested.

Case studies: rotavirus vaccination (India and Ethiopia) and HPV vaccination (China): Verguet and colleagues applied ECEA to rotavirus vaccination in India and Ethiopia[7], comparing the consequences of publicly financed vaccination (which removes the financial barrier to the vaccine for low-income households) against a scenario in which the vaccine is available but must be paid for out-of-pocket. In Ethiopia, publicly financed rotavirus vaccination would generate an estimated 72,000 deaths averted annually; roughly half of these gains would accrue to the poorest wealth quintile, who face both the greatest disease burden and the greatest financial barrier to vaccination. The analysis also showed that public financing of the vaccine would avert an estimated 25,000 cases of poverty annually, because the out-of-pocket cost of managing severe rotavirus disease (hospitalisation, rehydration therapy) is catastrophic for the poorest households. Neither of these equity and financial protection findings would be visible in a conventional CUA that reported only aggregate cost per DALY averted.

A parallel ECEA of publicly financed HPV vaccination in China, published by Levin and colleagues[11], found that public financing generated QALY gains concentrated in lower-income provinces where cervical cancer incidence is highest and private-pay vaccination coverage is lowest, and that the financial protection benefit - measured as the averted cost of cervical cancer treatment for families who would otherwise face catastrophic expenditure - was substantial and again disproportionately concentrated in the poorest quintile. The study's BCR for public financing, calculated as a conventional CBA using cost-of-illness averted, was positive at a conservative discount rate; the ECEA dimension added the insight that the financial protection and equity benefits further strengthened the case, in ways that the BCR could not decompose by quintile.

These examples illustrate that ECEA is particularly well-suited to evaluations of publicly financed programmes in settings where health is strongly associated with ability to pay, and where the distributional and financial-protection justification for public funding is at least as politically important as the aggregate cost-effectiveness justification. For researchers working in LMIC settings, ECEA offers a principled way to capture what CBA traditionally offers (cross-sectoral consequences, financial impacts) while retaining the QALYs-or-DALYs framework familiar to health economists and avoiding the contested step of VSL-based monetisation.

MULTI-CRITERIA DECISION ANALYSIS AS A NON-AGGREGATIVE ALTERNATIVE

Multi-criteria decision analysis (MCDA) encompasses a family of structured approaches to decisions that involve multiple, often conflicting objectives, and which cannot be reduced to a single summary metric without loss of information that decision-makers judge to be important.[12,13] MCDA has attracted growing interest in HTA precisely because the limitations of both CEA and CBA identified in this review - the contested threshold, the equity blindness of aggregate measures, the monetisation controversy - are all, at root, problems that arise from the requirement to collapse multidimensional value into a single number. MCDA does not require this aggregation.

Structure and methods: An MCDA for HTA typically involves: (i) defining a set of criteria that capture all relevant dimensions of value (e.g., clinical effectiveness, safety, unmet need, equity impact, budget impact, innovation, severity of condition); (ii) scoring each intervention on each criterion; (iii) eliciting the relative weights that decision-makers or stakeholders assign to each criterion; and (iv) computing a weighted score for each intervention, which can then be compared across options.[12,13] The ISPOR MCDA Emerging Good Practices Task Force, which produced two landmark reports.[12,13], identified additive scoring models, outranking models, and full preference models as the principal analytical approaches, and noted that choice among them involves non-trivial assumptions about criterion independence and the measurability of value on each criterion's scale.

The EVIDEM framework (Evidence and Value: Impact on DEcisionMaking) is one of the most widely applied MCDA frameworks in HTA, using a combination of 15 quantitative criteria and 6 qualitative contextual criteria scored and weighted to produce an overall value estimate for each intervention.[12] Bayesian approaches to MCDA, which treat criterion weights as uncertain parameters and propagate that uncertainty through the analysis, are increasingly available and are consistent with the probabilistic approach to uncertainty recommended for CEA.

Relationship to CEA and CBA, and limitations: MCDA does not replace CEA or CBA; it is more accurately positioned as a decision-support framework that is most useful when: (a) the range of relevant value dimensions is too broad to be captured by a single metric (QALY or monetary NPV); (b) the weighting of dimensions is genuinely contested and it is more legitimate to make that contestation explicit than to embed it in a threshold; or (c) the decision-making body includes multiple stakeholders with different preference structures (e.g., clinicians, patients, payers, and industry in a formulary committee). The ISPOR task force[13] explicitly noted that many MCDA applications in health care contain methodological flaws - including the use of additive models that pre-

suppose criterion independence without testing it, and a disconnect between the scales on which criteria are scored and the weights applied to them - that undermine their usefulness. Researchers should therefore treat MCDA as a rigorous method requiring careful design, not as a procedure that bypasses the need for evidence-based value judgements.

The relationship to CBA is worth making explicit: a well-specified CBA with comprehensive monetisation of all relevant effects is, in principle, a special case of MCDA with a single criterion (monetary net benefit) and a single weight of 1. MCDA becomes preferable when the analyst or decision-maker is unwilling to defend the monetisation of some dimensions (e.g., dignity, carer wellbeing, innovation value) that are nonetheless judged relevant to the decision, and when the deliberative process of eliciting and exposing criterion weights is itself considered a legitimate and valuable component of the HTA process.

VALUE-OF-INFORMATION ANALYSIS: LINKING ECONOMIC EVALUATION TO RESEARCH PRIORITISATION

Value-of-information (VOI) analysis addresses a question that neither CEA, CBA, nor MCDA can answer in their standard forms: given the current evidence base and its uncertainty, is additional research warranted before a coverage or reimbursement decision is made - and if so, what kind of research, and how much would it be worth to commission? VOI thus extends the economic evaluation framework from a static decision about an existing technology to a dynamic decision about both the technology and the evidence that supports it.

Key concepts: The foundational concept is the expected value of perfect information (EVPI), which measures how much a decision-maker would be willing to pay to eliminate all uncertainty in the current model before making a coverage decision.[14,15] EVPI is computed by comparing the expected NMB under the current (uncertain) decision to the expected NMB that would be achieved if all parameters were known with certainty - i.e., the gain from always making the right decision. In the NMB framework: EVPI = E[max(NMB)] - max(E[NMB]), where the first term is the expected NMB if the optimal intervention is chosen after all uncertainty is resolved, and the second term is the NMB of the currently optimal choice given existing evidence. If EVPI is small relative to the cost of additional research, a coverage decision should be made now on current evidence; if EVPI is large, further research may be warranted.

Expected value of partial perfect information (EVPPI) and expected value of sample information (EVSI) refine this by attributing uncertainty to specific parameters and specific proposed study designs respectively, enabling the analyst to identify which parameters drive the value of additional research and which study designs would be most efficient in reducing it.[14,15] These calculations are now computationally feasible using PSA output from standard CEA models, and guidance on their computation has been published by the ISPOR VOI Task Force.[15]

Relationship to CEA and CBA: VOI analysis applies most naturally in the CEA/NMB framework, since EVPI is computed on the NMB scale and directly answers the question: at what WTP threshold, and with what probability, does the optimal decision change if uncertainty were resolved? However, VOI concepts are fully portable to CBA: the expected value of perfect information about a VSL estimate, or about the scope of non-health benefits, can be computed on the monetary NPV scale using the same logic. This has been done, for example, in evaluations of environmental and transport regulations where the VSL is both the largest driver of the BCR and the most uncertain parameter.

For applied researchers, VOI analysis is most valuable in three specific circumstances: (a) when an HTA decision is uncertain and the decision-maker is considering whether to defer a coverage decision pending further evidence (an 'only in research' or 'coverage with evidence development' scenario); (b) when a research funder wishes to prioritise among competing research proposals and seeks an evidence-based estimate of the health and/or monetary value of each; and (c) when the analyst has identified specific parameters in a CEA or CBA model that drive the ICER or NPV and wishes to quantify the value of additional data on those parameters specifically.

PRACTICAL GUIDANCE: SELECTING AND COMBINING FRAMEWORKS

The preceding sections demonstrate that the choice between CEA and CBA is less important than the set of secondary choices - perspective, equity treatment, threshold source, uncertainty handling, and non-health consequences - that apply within either framework. The following structured guidance is offered to researchers designing new evaluations, organised by the primary decision question.

Decision questions and primary framework selection

Reporting standards and audience alignment: All evaluations, regardless of primary framework, should be reported in compliance with the CHEERS 2022 statement[8], which covers model structure, data sources, uncertainty analysis, and equity considerations for both CEA and CBA. When an evaluation is intended for submission to a specific HTA body or journal, analysts should additionally consult the body's own methodological guidelines (e.g., NICE Methods Guide for HTA submissions, PBAC Guidelines in Australia, CADTH guidelines in Canada) since these specify requirements for perspective, discount rates, utility-value sources, and subgroup analysis that may differ from published recommendations.

Table 2: Framework selection by primary decision question

Decision question	Primary framework	Recommended extensions
Should this drug/device be reimbursed by a payer with a fixed health budget?	CUA (ICER vs. established threshold)	PSA + CEAC; DCEA if equity impact likely; VOI if decision uncertain
Should this programme be funded from general government revenue, competing with non-health alternatives?	CBA (NPV/BCR) or CEA + NMB with explicit threshold justification	Distributional weighting of WTP; range of VSL assumptions; budget-impact analysis
What is the value of public financing for an LMIC health intervention relative to private out-of-pocket payment?	ECEA (health gains + FRP by quintile)	Supplementary CUA if donor agency requires ICER; DCEA for equity summary
How should a payer choose among multiple candidate programmes with incommensurable benefits (health, safety, financial, social)?	MCDA (structured multi-criteria scoring)	CUA and/or CBA as inputs to MCDA criteria; EVPI if rankings are sensitive to uncertain inputs
Is more evidence needed before a coverage decision is made on a new technology?	VOI analysis (EVPI, EVPPI, EVSI) applied to existing CUA model	Coverage-with-evidence-development framework; adaptive pathways design

CUA = cost-utility analysis; CBA = cost-benefit analysis; ICER = incremental cost-effectiveness ratio; NMB = net monetary benefit; PSA = probabilistic sensitivity analysis; CEAC = cost-effectiveness acceptability curve; DCEA = distributional cost-effectiveness analysis; ECEA = extended cost-effectiveness analysis; MCDA = multi-criteria decision analysis; FRP = financial risk protection; EVPI/EVPPI/EVSI = expected value of perfect/partial perfect/sample information.

The audience of the evaluation should determine how results are presented, not which methods are used. A well-designed CUA that includes a distributional cost-effectiveness extension can present the conventional ICER for the HTA committee and the equity impact findings for the public health team without duplication of effort. A CBA with a sensitivity analysis over a range of VSL assumptions can present its BCR both for a finance ministry audience (to whom BCR > 1 is directly interpretable) and as an NMB relative to the health system's implicit threshold (for readers more familiar with CEA conventions). These presentations are arithmetically equivalent and require only modest additional reporting.

Common methodological pitfalls to avoid: Selecting a framework post hoc to achieve a favourable conclusion: the primary framework, perspective, and threshold (or VSL) should be pre-specified in a published or registered analysis plan before results are computed, as is now expected by most major journals and HTA bodies.[8]

- Using a single-point ICER or BCR without uncertainty analysis: all summary estimates should be accompanied by probabilistic sensitivity analysis and presented with cost-effectiveness acceptability curves (for CUA) or sensitivity ranges over key assumptions (for CBA). A single-point ICER that appears below the threshold can cross the threshold for a non-trivial fraction of the PSA iterations, which is decision-relevant information.[2,3]

- Treating the ICER as self-interpreting without specifying and justifying the threshold: an ICER is only meaningful relative to a threshold, and that threshold should be stated with its source (supply-side empirical estimate, demand-side WTP estimate, or conventional/regulatory value) and tested in sensitivity analysis.[4]

- Claiming a societal perspective while omitting major cost categories: a payer-perspective analysis that adds productivity losses but omits informal care costs, or a CBA that monetises mortality but omits morbidity, is inconsistent in scope and will systematically under- or over-state net benefit.[1,2]

- Transferring VSL estimates across countries without adjustment: VSL estimates from high-income country wage-risk studies are not directly applicable to LMIC populations; ECEA offers a theoretically sounder alternative for cross-country generalisability in such settings.[7]

- Ignoring equity: reporting aggregate results without any discussion of distributional consequences is increasingly inconsistent with the equity commitments embedded in the mandates of most public health systems; DCEA or at minimum a qualitative equity impact assessment should be included.[5,6]

OPEN METHODOLOGICAL QUESTIONS AND FUTURE RESEARCH PRIORITIES

The review identifies the following as the most consequential unresolved questions in the literature, where additional research would most directly improve the quality of health resource allocation decisions.

Empirical estimation of cost-effectiveness thresholds in LMICs: The supply-side threshold literature has been developed almost entirely for high-income systems with comprehensive programme budget data (principally the UK NHS).[4] Analogous empirical estimates for LMICs, where programme budget tracking is often less developed and the marginal productivity of health spending may differ substantially, are a priority. WHO guidance provides indicative thresholds (e.g., 1-3 times GDP per capita per DALY averted) but these are acknowledged to lack empirical grounding, and recent literature has argued for country-specific empirical estimation as a research priority.

Integration of DCEA into routine HTA: DCEA

methods are now sufficiently mature to be applied routinely alongside standard CUA in any evaluation where socioeconomic patterning of uptake, effectiveness, or disease burden is plausible - which encompasses most common conditions.[5,6] However, the data infrastructure required (population-level health state data stratified by deprivation, linked to disease-specific utility data) remains poorly developed in most countries outside the UK. Developing this infrastructure, and establishing consensus on which measures of inequality should be reported in DCEA alongside the efficiency-equity trade-off plane, are priorities for the field.

Harmonised monetisation standards for CBA across sectors: A persistent barrier to using CBA for cross-sectoral comparisons is the absence of harmonised VSL and VSLY estimates endorsed across government departments within a given jurisdiction. Health, transport, and environmental agencies in most countries use different monetisation parameters, making BCR comparisons across sectors incoherent. Work towards government-wide harmonised values - as pursued, for example, by the UK Green Book framework and the U.S. Office of Management and Budget - represents a methodological priority that, if achieved, would substantially increase the usefulness of CBA for cross-sectoral health investment decisions.

Decision-analytic frameworks for algorithmic and digital health technologies: Machine learning-based diagnostic and prognostic tools, remote monitoring programmes, and AI-assisted clinical decision support systems present new challenges for both CEA and CBA: their effects depend on the clinical pathway in which they are embedded, they generate value partly through information (which is not easily captured in QALY or monetary terms), and their performance may change as they are retrained. VOI analysis is well-positioned to address some of these challenges - the expected value of the information generated by an AI diagnostic, for example, can in principle be quantified on the EVPI scale - but this application remains at an early methodological stage.

Conclusion

The fundamental distinction between CEA and CBA - whether health outcomes are expressed in QALYs or in monetary units - is a useful pedagogical starting point, but it is not where the most consequential methodological choices in contemporary health economic evaluation are located. This review has argued that the net-benefit framework unifies the two methods under a common decision-theoretic structure, and that the substantive debate has shifted to four areas that apply within and across both frameworks: the source and magnitude of the cost-effectiveness threshold; the distributional pattern of costs and benefits across the population; the non-health consequences (particularly financial risk protection) that standard CUA omits; and the treatment of decision uncertainty and its implications for the research agenda.

Each of these areas has a developing methodological literature, illustrated here with published case studies - the Claxton threshold analysis for the UK NHS, the Asaria bowel cancer screening DCEA, and the Verguet rotavirus and HPV vaccination ECEAs - that demonstrate concretely how methodological choices change substantive conclusions. The practical framework provides structured guidance for researchers on selecting among CEA, CBA, ECEA, MCDA, and VOI as primary or supplementary frameworks depending on the decision question, the available evidence base, and the institutional context.

The most important message for researchers entering this field is that neither CEA nor CBA provides a self-interpreting, algorithm-ready decision metric. Both rest on contested empirical and normative choices - the threshold in CUA, the VSL or equity weights in CBA - whose specification should be pre-registered, justified by reference to the institutional context and current methodological guidance, and subjected to comprehensive sensitivity analysis. The goal is not to produce a number that tells a decision-maker what to do, but to make the value judgements underlying a resource allocation decision transparent, consistent, and defensible to the populations whose health is at stake.

Individual Authors' Contributions: PP conceived the review topic, performed the literature search, analyzed and synthesized the relevant literature, and drafted the manuscript. **PP** critically revised the manuscript for important intellectual content and approved the final version for publication.

Declaration of Non-use of Generative AI Tools: This article was prepared without the use of generative AI tools for content creation, analysis, or data generation. All findings and interpretations are based solely on the authors' independent work and expertise.

References

1. Drummond MF, Sculpher MJ, Claxton K, Stoddart GL, Torrance GW. Methods for the Economic Evaluation of Health Care Programmes. 4th ed. Oxford: Oxford University Press; 2015.
2. Sanders GD, Neumann PJ, Basu A, Brock DW, Feeny D, Krahn M, et al. Recommendations for conduct, methodological practices, and reporting of cost-effectiveness analyses: Second Panel on Cost-Effectiveness in Health and Medicine. JAMA. 2016;316(10):1093-1103. DOI: https://doi.org/10.1001/jama.2016.12195 PMid:27623463
3. Stinnett AA, Mullahy J. Net health benefits: a new framework for the analysis of uncertainty in cost-effectiveness analysis. Med Decis Making. 1998;18(2 Suppl):S68-80. DOI: https://doi.org/10.1177/0272989X98018002S09 PMid:9566468
4. Claxton K, Martin S, Soares M, Rice N, Spackman E, Hinde S, et al. Methods for the estimation of the National Institute for

Health and Care Excellence cost-effectiveness threshold. Health Technol Assess. 2015;19(14):1-503. DOI: https://doi.org/10.3310/hta19140 PMid:25692211 PMCid:PMC4781395

5. Asaria M, Griffin S, Cookson R, Whyte S, Tappenden P. Distributional cost-effectiveness analysis of health care programmes - a methodological case study of the UK Bowel Cancer Screening Programme. Health Econ. 2015;24(6):742-754. DOI: https://doi.org/10.1002/hec.3058 PMid:24798212
6. Cookson R, Griffin S, Norheim OF, Culyer AJ, editors. Distributional Cost-Effectiveness Analysis: Quantifying Health Equity Impacts and Trade-Offs. Oxford: Oxford University Press; 2020. DOI: https://doi.org/10.1093/med/9780198838197.001.0001
7. Verguet S, Murphy S, Anderson B, Johansson KA, Glass R, Rheingans R. Public finance of rotavirus vaccination in India and Ethiopia: an extended cost-effectiveness analysis. Vaccine. 2013;31(42):4902-4910. DOI: https://doi.org/10.1016/j.vaccine.2013.07.014 PMid:23871824
8. Husereau D, Drummond M, Augustovski F, de Bekker-Grob E, Briggs AH, Carswell C, et al. Consolidated Health Economic Evaluation Reporting Standards (CHEERS) 2022 explanation and elaboration: a report of the ISPOR CHEERS II good practices task force. Value Health. 2022;25(1):10-31. DOI: https://doi.org/10.1016/j.jval.2021.10.008
9. Briggs AH, Weinstein MC, Fenwick EA, Karnon J, Sculpher MJ, Paltiel AD. Model parameter estimation and uncertainty analysis: a report of the ISPOR-SMDM Modeling Good Research Practices Task Force Working Group-6. Med Decis Making. 2012;32(5):722-732. DOI: https://doi.org/10.1177/0272989X12458348 PMid:22990087 PMCid:PMC5577377
10. Neumann PJ, Cohen JT, Weinstein MC. Updating cost-effectiveness - the curious resilience of the $50,000-per-QALY threshold. N Engl J Med. 2014;371(9):796-797. DOI: https://doi.org/10.1056/NEJMp1405158 PMid:25162885
11. Levin CE, Sharma M, Olson Z, Verguet S, Shi JF, Wang SM, et al. An extended cost-effectiveness analysis of publicly financed HPV vaccination to prevent cervical cancer in China. Vaccine. 2015;33(24):2830-2841. DOI: https://doi.org/10.1016/j.vaccine.2015.02.052 PMid:25770785
12. Thokala P, Devlin N, Marsh K, Baltussen R, Boysen M, Kalo Z, et al. Multiple criteria decision analysis for health care decision making - an introduction: Report 1 of the ISPOR MCDA Emerging Good Practices Task Force. Value Health. 2016;19(1):1-13. DOI: https://doi.org/10.1016/j.jval.2015.12.003 PMid:26797229
13. Marsh K, IJzerman M, Thokala P, Baltussen R, Boysen M, Kaló Z, et al. Multiple criteria decision analysis for health care decision making - emerging good practices: Report 2 of the ISPOR MCDA Emerging Good Practices Task Force. Value Health. 2016;19(2):125-137. DOI: https://doi.org/10.1016/j.jval.2015.12.016 PMid:27021745
14. Claxton KP, Sculpher MJ. Using value of information analysis to prioritise health research: some lessons from recent UK experience. Pharmacoeconomics. 2006;24(11):1055-1068. DOI: https://doi.org/10.2165/00019053-200624110-00003 PMid:17067191
15. Fenwick E, Steuten L, Knies S, Ghabri S, Basu A, Murray JF, et al. Value of information analysis for research decisions - an introduction: Report 1 of the ISPOR Value of Information Analysis Emerging Good Practices Task Force. Value Health. 2020;23(2):139-150. DOI: https://doi.org/10.1016/j.jval.2020.01.001 PMid:32113617