
Global Certificate in Health Technology Assessment (Advanced)

Health Technology Assessment Principles

Health Technology Assessment (HTA) is a multidisciplinary process that systematically evaluates the properties, effects, and impacts of health technologies. The purpose of HTA is to inform decision-makers about the value of a technology, guiding policy, reimbursement, and clinical practice. Mastery of the terminology used in HTA is essential for professionals who must interpret evidence, communicate findings, and apply results in real-world settings. The following exposition defines the core terms and vocabulary that underpin the principles of HTA, illustrating each concept with practical examples, discussing typical applications, and highlighting common challenges encountered in practice.

Technology refers to any intervention that can be used to diagnose, treat, prevent, or monitor disease. This includes drugs, medical devices, diagnostic tests, procedures, and even complex service delivery models such as telemedicine platforms. For instance, a new implantable cardiac defibrillator is a technology, as is a software algorithm that predicts sepsis risk from electronic health record data. Recognising the breadth of what constitutes a technology is crucial because HTA methods differ depending on the type of intervention. A challenge often arises when a technology blurs traditional categories, such as a combination product that integrates a drug with a device, requiring analysts to adapt evaluation frameworks accordingly.

Clinical effectiveness measures the extent to which a technology produces the intended health outcomes under typical conditions of use. It is usually expressed in terms of improvement in survival, symptom relief, or disease progression. Clinical effectiveness is distinguished from efficacy, which reflects outcomes under ideal, controlled trial conditions. An example of clinical effectiveness is the reduction in hospital readmissions observed when a home-based cardiac rehabilitation program is implemented across multiple community health centers. In practice, gathering real-world effectiveness data can be hampered by incomplete records, heterogeneity of patient populations, and variations in adherence, all of which challenge the reliability of the assessment.

Safety encompasses the frequency and severity of adverse events associated with a technology. Safety data are typically derived from clinical trials, post-marketing surveillance, and pharmacovigilance databases. For a new anticoagulant, safety considerations might focus on the incidence of major bleeding events compared with standard therapy. A practical challenge is that rare adverse events may not emerge until after widespread adoption, requiring HTA bodies to update recommendations as new safety information becomes available.

Economic evaluation is the systematic comparison of alternative interventions in terms of both costs and consequences. The most common forms of economic evaluation in HTA are cost-effectiveness analysis (CEA), cost-utility analysis (CUA), and cost-benefit analysis (CBA). Economic evaluation provides a framework for decision-makers to allocate limited resources efficiently.

Cost-effectiveness analysis (CEA) compares the costs and health outcomes of two or more interventions, expressing results as a ratio of incremental cost to incremental effect, often called the incremental cost-effectiveness ratio (ICER). For example, a CEA might compare a novel biologic therapy for rheumatoid arthritis with an existing disease-modifying drug, reporting an ICER of \$45,000 per additional responder. The primary challenge in CEA is selecting an appropriate measure of effect that captures the relevant health benefit, especially when outcomes are multidimensional.

Incremental cost-effectiveness ratio (ICER) is calculated by dividing the difference in costs between two interventions by the difference in their effects. It represents the additional cost required to achieve one additional unit of health benefit. In the earlier rheumatoid arthritis example, the ICER indicates the extra cost per extra patient achieving a defined response. Interpreting the ICER requires reference to willingness-to-pay thresholds, which differ across jurisdictions and can be politically sensitive.

Cost-utility analysis (CUA) is a specific form of CEA that uses quality-adjusted life years (QALYs) as the measure of effect. QALYs combine length of life with health-related quality of life, allowing comparison across disparate health conditions. A CUA of a new vaccine may reveal that it yields 0.08 QALYs per vaccinated individual at an incremental cost of \$1,200, resulting in an ICER of \$15,000 per QALY gained. A frequent challenge in CUA is the derivation of utility values, which may rely on generic instruments such as the EQ-5D or disease-specific measures, each with methodological limitations.

Quality-adjusted life year (QALY) quantifies health outcomes by weighting each year of life lived by a utility factor ranging from 0 (death) to 1 (perfect health). For instance, a year lived with moderate chronic pain might be assigned a utility of 0.6, resulting in 0.6 QALYs for that year. The use of QALYs enables comparison of interventions across disease areas; however, critics argue that QALYs may undervalue interventions for populations with lower baseline utilities, such as the elderly or those with disabilities.

Cost-benefit analysis (CBA) expresses both costs and benefits in monetary terms, allowing a direct assessment of net monetary gain or loss. In a CBA of a smoking-cessation program, benefits might include reduced healthcare expenditures, increased productivity, and avoided premature mortality, all monetised. While CBA offers a clear decision rule (net benefit > 0), assigning monetary values to health outcomes raises ethical and methodological concerns, especially when valuing human life.

Perspective defines the viewpoint from which costs and outcomes are measured. Common perspectives include the health-care payer, the societal, and the patient perspective. A payer perspective includes only direct medical costs reimbursed by the insurer, whereas a societal perspective adds indirect costs such as lost productivity and caregiver burden. For example, when evaluating a new chemotherapy regimen, a payer perspective might consider drug acquisition and administration costs, while a societal perspective also incorporates patients' time off work and transportation expenses. Selecting the appropriate perspective is critical because it influences the magnitude of costs and benefits, and differing perspectives can lead to divergent conclusions about cost-effectiveness.

Time horizon denotes the period over which costs and outcomes are assessed. Short-term horizons may capture immediate costs but miss long-term benefits or harms, whereas long-term horizons provide a more comprehensive view but increase uncertainty due to extrapolation. In assessing a vaccine that prevents a disease with lifelong sequelae, a lifetime horizon is typically appropriate to capture all future health gains and cost savings. The challenge lies in choosing a horizon that balances relevance with data availability, and in modelling long-term outcomes when empirical data are limited.

Discounting is the process of adjusting future costs and benefits to reflect time preference, acknowledging that a dollar (or health benefit) today is valued more than the same amount in the future. Standard discount rates range from 3 % to 5 % per annum for both costs and outcomes, though some jurisdictions apply differential rates. For a health technology with upfront costs and delayed benefits, discounting reduces the present value of those benefits, potentially affecting the ICER. Debate persists regarding the appropriate discount rate, especially for interventions with long-term health impacts, such as preventive measures.

Sensitivity analysis examines how results change when key assumptions or input parameters are varied. Deterministic (one-way) sensitivity analysis varies one parameter at a time, while probabilistic sensitivity analysis (PSA) assigns probability distributions to inputs and runs simulations to generate a range of possible outcomes. In a CUA of a new insulin analogue, a PSA might reveal a 70 % probability that the ICER falls below the accepted willingness-to-pay threshold. Conducting thorough sensitivity analysis is essential to assess the robustness of conclusions, yet it can be resource-intensive and requires expertise in statistical modelling.

Willingness-to-pay threshold represents the maximum amount a decision-maker is prepared to spend for an additional unit of health benefit, such as a QALY. In the United Kingdom, the National Institute for Health and Care Excellence (NICE) commonly uses a threshold of £20,000-£30,000 per QALY. In the United States, there is no formal threshold, but \$50,000-\$150,000 per QALY is often referenced. Thresholds are not static; they may be adjusted for disease severity, rarity, or equity considerations. The selection of a threshold can be contentious, as it directly influences whether a technology is deemed cost-effective.

Budget impact analysis (BIA) estimates the financial consequences of adopting a new technology within a specific budget context over a defined period, usually 1-5 years. BIA complements economic evaluation by addressing affordability, not just value for money. For example, a BIA for a high-cost oncology drug may show that, despite a favourable ICER, its adoption would consume a substantial portion of the oncology budget, prompting negotiations on pricing or patient selection criteria. A key challenge is that BIA often requires detailed utilisation forecasts and may be sensitive to assumptions about market uptake.

Decision-analytic model is a structured framework that synthesises evidence from multiple sources to estimate costs and outcomes. Common model types include decision trees, Markov models, and discrete-event simulation. Decision trees are suitable for short-term, simple pathways, while Markov models capture chronic conditions with health states and transition probabilities. A discrete-event simulation can model complex interactions and resource constraints, such as patient flow in an intensive care unit. Building

a robust model demands transparent documentation, validation, and calibration, yet model complexity can impede stakeholder understanding and acceptance.

Markov model divides the disease process into mutually exclusive health states (e.g., "healthy," "diseased," "dead") and defines transition probabilities between states for each cycle length (often one year). The model aggregates costs and utilities accrued in each state, producing cumulative outcomes over the chosen horizon. In evaluating a chronic heart-failure therapy, a Markov model might include states such as "stable," "hospitalised," and "post-hospitalisation," with transitions reflecting disease progression and mortality. A limitation of Markov models is the "memoryless" property, which assumes that future transitions depend only on the current state, potentially oversimplifying histories of prior events.

Probabilistic sensitivity analysis (PSA) assigns probability distributions to uncertain parameters (e.g., costs, utilities, transition probabilities) and uses Monte Carlo simulation to propagate uncertainty through the model. PSA generates a scatter plot of incremental costs versus incremental QALYs, facilitating the construction of cost-effectiveness acceptability curves (CEACs). A CEAC shows the probability that an intervention is cost-effective across a range of willingness-to-pay thresholds. PSA is valued for its ability to reflect overall uncertainty, but it requires careful selection of appropriate distributions and sufficient computational resources.

Cost-effectiveness acceptability curve (CEAC) displays the probability that a technology is cost-effective compared with an alternative across different willingness-to-pay thresholds. The curve assists decision-makers in visualising uncertainty and in selecting a threshold that aligns with policy objectives. For a new antiviral, the CEAC might indicate a 55 % probability of cost-effectiveness at a threshold of \$100,000 per QALY, rising to 85 % at \$150,000 per QALY. Interpreting CEACs demands an understanding of the underlying PSA assumptions and the context of the health system's budget constraints.

Health-related quality of life (HRQoL) captures patients' perceived physical, mental, and social well-being as related to health status. Instruments such as the EQ-5D, SF-36, and disease-specific questionnaires are used to derive utility values for QALY calculations. In a trial of a novel migraine prophylactic, HRQoL may be measured using the Migraine-Specific Quality of Life questionnaire, which can be mapped onto utility scores. Challenges include cultural differences in health perception, ceiling effects, and the need for validated translation of instruments.

Incremental analysis isolates the additional costs and benefits of a new technology relative to the next best alternative, rather than comparing each option to a "do-nothing" baseline. Incremental analysis ensures that resources are directed toward the most efficient improvement. For a new surgical robot, incremental analysis would compare it with conventional laparoscopy, not with no surgery. A common pitfall is the "dominated" scenario, where a technology is both more costly and less effective, making it unequivocally inferior.

Dominated describes a technology that is strictly inferior because it incurs higher costs and delivers lower

effectiveness than an alternative. In such cases, the dominated technology is excluded from further consideration. For example, if a new oral antidiabetic agent costs \$2,000 more per patient per year and yields fewer QALYs than an existing generic drug, it is dominated. Identifying dominated options simplifies the decision-making process but requires accurate comparative data.

Extended dominance occurs when a technology is less efficient than a combination of two other options. If Technology A has an ICER of \$10,000 per QALY, Technology B has an ICER of \$30,000 per QALY, and Technology C has an ICER of \$25,000 per QALY, Technology C may be excluded due to extended dominance because a mix of A and B provides a more favourable cost-effectiveness profile. Recognising extended dominance helps avoid sub-optimal choices, yet the concept can be confusing for stakeholders unfamiliar with incremental analysis.

Budget constraint represents the finite resources available to a health system, limiting the number and type of technologies that can be funded. Budget constraints compel decision-makers to prioritise interventions based on cost-effectiveness, disease burden, and equity. A health ministry may allocate \$500 million to oncology drugs, requiring HTA to rank therapies within that envelope. Balancing cost-effectiveness with other policy goals, such as access for rare disease patients, poses a persistent challenge.

Equity concerns the fair distribution of health benefits across population groups, considering factors such as age, socioeconomic status, and geography. HTA frameworks increasingly incorporate equity weighting, adjusting QALY valuations to reflect societal preferences for reducing health inequalities. For instance, a vaccination program targeting low-income neighborhoods may receive an equity boost in the analysis. Operationalising equity is complex, as it requires explicit value judgments and may conflict with efficiency-driven decisions.

Stakeholder refers to any individual or group with an interest in the HTA process, including clinicians, patients, industry representatives, payers, and policymakers. Engaging stakeholders throughout the assessment improves relevance, transparency, and acceptance of recommendations. In practice, stakeholder consultations may involve public hearings, advisory panels, or written submissions. Managing diverse viewpoints can be challenging, especially when commercial interests clash with public health priorities.

Evidence synthesis is the systematic aggregation of data from multiple studies to inform HTA inputs. Techniques include systematic reviews, meta-analysis, network meta-analysis, and indirect treatment comparisons. For a new antihypertensive, evidence synthesis might combine results from several randomized controlled trials to estimate a pooled effect on systolic blood pressure. Conducting rigorous evidence synthesis demands adherence to methodological standards, such as PRISMA, and the ability to assess heterogeneity, publication bias, and study quality.

Systematic review follows a predefined protocol to identify, appraise, and summarise all relevant literature on a specific question, aiming to minimise bias. The review process includes a comprehensive search strategy, eligibility criteria, data extraction, and risk-of-bias assessment. A systematic review of the efficacy

of a new biologic for psoriasis would provide the foundation for subsequent economic modelling. A common obstacle is the scarcity of high-quality comparative studies, leading to reliance on indirect evidence.

Meta-analysis statistically combines results from comparable studies to produce a pooled estimate of effect size, increasing precision. Fixed-effect and random-effects models address different assumptions about between-study variability. For example, a meta-analysis of five trials of a novel vaccine might yield a pooled relative risk reduction of 60% for infection. Meta-analysis can be limited by heterogeneity, selective reporting, and the quality of the included studies.

Network meta-analysis (NMA) extends conventional meta-analysis by allowing simultaneous comparison of multiple interventions, even when head-to-head trials are absent. NMA constructs a network of evidence, estimating relative effects between all pairs of treatments. In evaluating treatments for chronic obstructive pulmonary disease, an NMA could rank inhaled therapies, oral agents, and combination regimens. Implementing NMA requires sophisticated statistical software and expertise, and results may be sensitive to assumptions about consistency and transitivity.

Indirect treatment comparison (ITC) derives comparative effectiveness between two interventions using a common comparator, often employed when direct head-to-head trials are lacking. A simple ITC might compare Drug X to Drug Y by referencing their separate trials against a placebo. While useful, indirect comparisons can introduce bias if the underlying studies differ in patient characteristics, outcome definitions, or follow-up duration.

Risk of bias assesses the potential for systematic errors in study design, conduct, or reporting that could distort effect estimates. Tools such as the Cochrane Risk of Bias tool evaluate domains including randomisation, blinding, and outcome reporting. A high risk of bias in a pivotal trial of a new device may diminish confidence in its reported effectiveness, prompting HTA reviewers to downgrade the evidence quality.

GRADE (Grading of Recommendations Assessment, Development and Evaluation) is a systematic approach to rating the certainty of evidence and strength of recommendations. GRADE categorises evidence as high, moderate, low, or very low certainty based on factors like risk of bias, inconsistency, indirectness, imprecision, and publication bias. In HTA, GRADE informs the confidence placed in clinical effectiveness inputs and guides the interpretation of cost-effectiveness results. Applying GRADE consistently can be challenging, especially when evidence is sparse or derived from observational studies.

Uncertainty in HTA arises from variability in data, methodological choices, and model assumptions. Types of uncertainty include parameter uncertainty (variability in input values), structural uncertainty (choice of model form), and methodological uncertainty (differences in discount rates or perspective). Transparent reporting of uncertainty, through sensitivity analysis and scenario analysis, helps decision-makers understand the robustness of conclusions. However, communicating uncertainty to non-technical audiences

remains a persistent difficulty.

Scenario analysis explores alternative sets of assumptions or structural choices to examine how results change under different plausible conditions. For a vaccine, a scenario analysis might compare outcomes under varying coverage rates, waning immunity, or alternative pricing arrangements. Scenario analysis complements probabilistic sensitivity analysis by highlighting the impact of key structural decisions that cannot be captured by parameter variation alone. Selecting relevant scenarios requires stakeholder input and expert judgment.

Threshold analysis identifies the value of a parameter at which a technology switches from being cost-effective to not cost-effective. For example, a threshold analysis might reveal that the price of a new oncology drug must fall below \$10,000 per treatment course to achieve an ICER below \$50,000 per QALY. Threshold analysis is valuable for price negotiations and for informing price-setting policies, yet it depends on the validity of the underlying model.

Value-of-information analysis (VOI) quantifies the benefit of obtaining additional evidence to reduce decision uncertainty. The expected value of perfect information (EVPI) estimates the maximum worth of eliminating all uncertainty, while the expected value of partial perfect information (EVPPi) focuses on specific parameters. In a HTA of a novel gene-therapy, a high EVPI might justify commissioning a post-approval registry to gather long-term safety data. Conducting VOI analyses requires advanced modelling techniques and may be perceived as theoretical rather than practical by some stakeholders.

Real-world evidence (RWE) comprises data collected outside the controlled environment of randomized trials, such as electronic health records, registries, claims databases, and patient-reported outcomes. RWE can complement trial data by providing insights into effectiveness, safety, and utilisation patterns in routine practice. For a newly approved digital health app, RWE might reveal adherence rates and outcomes across diverse healthcare settings. Integrating RWE poses challenges related to data quality, standardisation, and privacy regulations.

Health economics is the broader discipline that underpins HTA, encompassing the study of resource allocation, cost measurement, and outcome valuation in health care. Health economists develop the analytical tools (e.g., cost-effectiveness models) used in HTA and advise on methodological standards. Collaboration between health economists and clinical experts is essential to ensure that models accurately reflect clinical pathways while maintaining economic rigour.

Decision-making framework provides a structured approach for translating HTA findings into policy actions. Frameworks may incorporate explicit criteria such as cost-effectiveness, disease burden, equity, and feasibility, often organised into a scoring or weighting system. NICE's appraisal framework, for instance, integrates clinical and economic evidence with patient values. Implementing a decision-making framework requires political commitment, clear governance structures, and transparent documentation of how criteria are applied.

Reimbursement refers to the payment arrangement between a health system or insurer and providers for the delivery of a technology. HTA outcomes frequently influence reimbursement decisions, determining whether a technology is covered, under what conditions, and at what price. A reimbursement policy might restrict a high-cost drug to patients who have failed first-line therapy, reflecting both cost-effectiveness considerations and clinical guidelines. Negotiating reimbursement rates can be complex, involving price-volume agreements, risk-sharing contracts, and managed entry schemes.

Managed entry scheme (MES) is a risk-sharing arrangement that allows conditional adoption of a technology while additional data are collected to confirm its value. Types of MES include coverage with evidence development (CED), performance-based contracts, and price-volume agreements. For a novel gene-editing therapy, a CED might grant provisional coverage contingent on the establishment of a patient registry to monitor long-term outcomes. Designing MES requires careful definition of outcome measures, data collection mechanisms, and criteria for continuation or termination.

Health technology horizon scanning is the systematic identification of emerging technologies that may impact health care in the near future. Horizon scanning informs priority-setting for HTA by flagging innovations that warrant early evaluation. Agencies may use literature searches, expert panels, and patent analyses to generate a pipeline of candidates. The challenge lies in balancing the need for timely assessment with the limited data typically available for nascent technologies.

Coverage decision is the formal determination by a health authority on whether a technology will be included in a publicly funded benefit package. Coverage decisions are often based on HTA reports, stakeholder input, and budgetary considerations. A positive coverage decision may be unconditional, or it may include stipulations such as restricted indications, prior-authorization requirements, or mandatory participation in post-marketing studies. The transparency of the decision-making process affects public trust and acceptance.

Clinical practice guideline (CPG) offers evidence-based recommendations for the management of specific health conditions. HTA findings frequently feed into CPG development, ensuring that recommendations reflect both clinical effectiveness and economic considerations. For example, a guideline for managing type 2 diabetes may incorporate HTA results to advise on the preferred second-line agents based on cost-effectiveness. Aligning HTA with guideline development can be difficult when evidence is rapidly evolving or when stakeholder consensus is lacking.

Regulatory approval is the authorization granted by a national or regional agency (e.g., FDA, EMA) that permits a technology to be marketed. Regulatory approval focuses on safety and efficacy, not on cost-effectiveness. Consequently, a technology may be approved but not reimbursed if HTA deems it not to represent good value. Coordinating regulatory and HTA timelines can streamline market entry, yet divergent evidentiary standards often cause misalignment.

Health system encompasses the organisations, institutions, resources, and people whose primary purpose is

to improve health. Understanding the structure, financing mechanisms, and delivery models of a health system is essential for contextualising HTA results. In a single-payer system, HTA recommendations may be implemented uniformly, whereas in a fragmented system with multiple insurers, adoption may vary widely. Tailoring HTA outputs to the specific health-system context enhances relevance but adds complexity to the analysis.

Implementation refers to the process of putting HTA recommendations into practice, which may involve policy changes, provider education, and monitoring mechanisms. Effective implementation ensures that the anticipated health gains and cost savings are realised. Barriers to implementation include lack of clinician awareness, resistance to change, and insufficient infrastructure. Implementation science offers tools to evaluate and optimise the uptake of HTA-informed interventions.

Monitoring and evaluation (M&E) assesses the performance of a technology after adoption, tracking outcomes such as utilisation, clinical effectiveness, safety, and budget impact. M&E data can feed back into future HTA cycles, supporting iterative improvement. For a newly introduced vaccine, M&E might involve surveillance of adverse events and measurement of population-level immunity. Maintaining high-quality M&E systems demands investment in data collection platforms and analytical capacity.

Health outcome denotes any measurable change in health status resulting from an intervention, including morbidity, mortality, functional status, and quality of life. Outcomes are the central focus of HTA, guiding both clinical and economic assessments. Selecting appropriate outcomes requires alignment with the disease burden and relevance to patients. In some cases, surrogate outcomes (e.g., biomarker changes) are used when long-term outcomes are not yet observable, but reliance on surrogates can introduce uncertainty regarding true clinical benefit.

Surrogate endpoint is a biomarker or intermediate outcome that is intended to predict a final clinical outcome but is not itself a direct measure of patient benefit. For a cholesterol-lowering drug, LDL-cholesterol reduction serves as a surrogate for cardiovascular events. Surrogates can accelerate evidence generation, yet they may not always correlate perfectly with true clinical outcomes, potentially leading to over- or under-estimation of a technology's value.

Comparator is the alternative against which a new technology is evaluated, which may be the current standard of care, a placebo, or another active treatment. Choosing an appropriate comparator is vital for a fair assessment; using an outdated or sub-optimal comparator can bias results. In a HTA of a novel insulin pump, the comparator would typically be the most widely used insulin delivery method, such as multiple daily injections. Stakeholder opinions may influence comparator selection, especially in therapeutic areas with rapidly evolving standards.

Cost component refers to a specific element of total cost, such as drug acquisition, administration, monitoring, adverse-event management, or indirect costs like lost productivity. Disaggregating costs into components enables sensitivity analysis and highlights cost drivers. For a surgical innovation, the cost

components might include the device price, operating-room time, staff training, and postoperative complications. Accurately capturing all relevant cost components can be difficult, especially for indirect costs that rely on external data sources.

Direct medical cost includes expenses directly related to health-care services, such as hospital stays, physician visits, procedures, and medications. These costs are typically captured in payer databases and are the primary focus of many HTA analyses. For a chronic disease management program, direct medical costs could encompass routine lab tests, specialist consultations, and drug therapy. Excluding non-medical costs may underestimate the total economic impact of a technology.

Indirect cost captures the economic value of lost productivity due to morbidity, premature mortality, or caregiver time. Indirect costs are often measured using the human capital approach (valuing lost earnings) or the friction cost method (considering the time needed to replace a worker). In evaluating a workplace wellness program, indirect costs might dominate the total cost-benefit calculation. Valuing indirect costs raises methodological debates, as assumptions about labor market dynamics can substantially influence results.

Human capital approach assigns a monetary value to lost productivity based on average earnings, assuming that each hour not worked represents a loss equal to the wage rate. This approach tends to produce higher estimates of indirect costs compared with alternative methods. Critics argue that it may overstate the societal burden for populations with lower average wages, such as retirees or unemployed individuals.

Friction cost method estimates indirect costs by considering the time required to replace a sick worker, rather than the full loss of earnings. This method typically yields lower cost estimates than the human capital approach, reflecting the assumption that labour markets can absorb temporary absences. Selecting between these methods influences the magnitude of indirect costs in HTA and can affect the perceived cost-effectiveness of interventions that primarily impact productivity.

Discount rate (repeated for emphasis) is applied to both costs and health outcomes, reflecting the societal preference for present over future benefits. While standard rates are often set by health-technology agencies, some analyses explore alternative rates in sensitivity analysis. A higher discount rate diminishes the present value of long-term benefits, potentially disadvantaging preventive interventions.

Price elasticity measures the responsiveness of demand for a technology to changes in its price. Understanding price elasticity helps predict how price adjustments will affect utilisation and overall budget impact. For a high-cost oncology drug, a low price elasticity may indicate that demand will remain relatively stable despite price reductions, implying limited savings from price cuts alone. Estimating price elasticity requires robust utilisation data and may be confounded by clinical guidelines.

Health-care utilization encompasses the frequency and intensity of services used, such as hospital admissions, outpatient visits, and diagnostic testing. Utilisation data are essential inputs for budget impact analyses and for estimating real-world costs. A new telehealth platform may reduce in-person visits, altering

utilization patterns. Accurately capturing utilization trends can be impeded by fragmented data sources and coding inconsistencies across providers.

Price-volume agreement is a contractual arrangement in which the price of a technology is linked to the volume of units purchased, often providing discounts as sales increase. Such agreements can mitigate budget impact while encouraging uptake. For a vaccine, a price-volume agreement might stipulate a 10% discount after the first 100,000 doses are administered. Negotiating these agreements requires transparent forecasting and mutual trust between manufacturers and payers.

Risk-sharing contract allocates financial risk between the payer and the manufacturer based on performance outcomes. If the technology fails to meet pre-specified clinical targets, the manufacturer may provide rebates or refunds. In a risk-sharing contract for a novel biologic, the payer might receive a rebate if the drug does not achieve a certain progression-free survival threshold in the real-world population. Designing measurable and clinically meaningful performance metrics is a central challenge in risk-sharing contracts.

Health-technology evaluation (repeated for emphasis) captures the full spectrum of assessment activities, from evidence synthesis to economic modelling, stakeholder engagement, and policy recommendation. A comprehensive evaluation ensures that all relevant dimensions—clinical, economic, ethical, and organisational—are considered. Conducting a full health-technology evaluation demands multidisciplinary expertise, adequate funding, and clear timelines.

Cost-minimisation analysis (CMA) is employed when two interventions have demonstrated equivalent effectiveness, allowing the comparison to focus solely on costs. For a generic drug that is bioequivalent to a brand-name counterpart, a CMA would assess which option is less expensive. CMA is less common than CEA or CUA because equivalence in outcomes is rarely established with high certainty.

Budget impact model (repeated) is a specific type of economic model that projects the financial effect of adopting a technology over a short-to-medium term horizon, typically 1-5 years. The model incorporates uptake rates, pricing, and changes in resource utilisation. A budget impact model for a new oral anticoagulant may reveal an initial increase in drug spend but offsetting savings from reduced stroke-related hospitalisations. Ensuring that the model aligns with the payer's budgeting cycles and reporting requirements enhances its usefulness.

Healthcare payer is the entity responsible for financing health-care services, which may be a government agency, insurance company, or employer-based plan. Payers use HTA to inform coverage decisions, negotiate prices, and manage formularies. Different payers may apply distinct thresholds and criteria, leading to variability in technology adoption across jurisdictions. Engaging payers early in the HTA process can facilitate smoother implementation.

Formulary is a list of medicines or technologies approved for use within a particular health-care system, often organised by therapeutic class and tiered according to cost-effectiveness. Inclusion on a formulary

typically requires HTA approval. For a national formulary, a new antiviral might be placed in a restricted tier, requiring specialist prescription. Formulary management must balance clinical need, budget constraints, and equity considerations.

Health-technology appraisal (HTA) (repeated) is the formal process by which an HTA agency reviews evidence and issues a recommendation. The appraisal may be accompanied by a detailed report outlining the evidence base, modelling assumptions, and the rationale for the recommendation. Transparency in the appraisal process builds credibility and facilitates stakeholder acceptance.

Decision-maker is the individual or body authorized to act on HTA recommendations, such as a health ministry, insurer, or hospital governing board. Decision-makers must weigh HTA evidence alongside political, social, and operational factors. Effective communication of HTA findings, tailored to the decision-maker's priorities, enhances the likelihood of evidence-based policy adoption.

Equity weighting adjusts the value of health gains to reflect societal preferences for reducing health disparities. Weighting may assign higher value to QALYs gained by disadvantaged groups. In a cost-utility analysis, applying equity weights could make a preventive program for low-income neighborhoods appear more cost-effective than a similar program for higher-income areas. Operationalising equity weighting requires consensus on the magnitude of weights and the groups to be prioritized.

Ethical analysis examines the moral implications of adopting or rejecting a technology, considering principles such as beneficence, non-maleficence, autonomy, and justice. Ethical considerations may influence HTA recommendations, especially for technologies that raise contentious issues, such as gene editing or end-of-life interventions. Integrating ethical analysis into HTA demands interdisciplinary collaboration and transparent deliberation.

Patient-reported outcome (PRO) captures the patient's perspective on symptoms, functional status, and health-related quality of life, often collected via questionnaires. PROs are increasingly incorporated into HTA to ensure that assessments reflect outcomes important to patients. For a chronic pain therapy, PRO measures may include pain intensity scales and interference with daily activities. Collecting high-quality PRO data can be limited by response bias and varying instrument validity across cultures.

Health-technology diffusion describes the rate and pattern by which a new technology spreads within a health system. Diffusion modelling helps predict uptake and associated budget impact. Factors influencing diffusion include clinical guidelines, reimbursement policies, physician awareness, and patient demand. Rapid diffusion of a high-cost technology may strain budgets, whereas slow diffusion can delay health benefits.

Stakeholder analysis identifies and assesses the interests, influence, and potential impact of various parties involved in the HTA process. Mapping stakeholder positions can guide engagement strategies, mitigate conflicts, and ensure that diverse perspectives are considered. In practice, stakeholder analysis may reveal that patient advocacy groups strongly support a technology, while payer organisations express concerns

about cost. Managing these dynamics is essential for credible HTA outcomes.

Implementation barrier is any factor that impedes the translation of HTA recommendations into practice, such as lack of clinician training, insufficient infrastructure, or regulatory hurdles. Identifying barriers early enables the development of targeted strategies, such as educational programmes or policy adjustments. Overcoming implementation barriers often requires coordinated effort across multiple system levels.

Policy levers are mechanisms that decision-makers can use to influence technology adoption, including pricing negotiations, formulary placement, reimbursement criteria, and incentive programmes. Selecting appropriate levers depends on the specific goals of the policy, such as controlling costs, encouraging innovation, or improving access. Effective use of policy levers can align market behaviour with HTA recommendations.

Health-technology procurement involves the processes by which health systems acquire technologies, encompassing tendering, contracting, and supply-chain management. Procurement decisions are informed by HTA evidence, but also by factors such as supplier reliability and legal considerations. Transparent procurement processes reduce the risk of corruption and ensure value for money.

Clinical pathway outlines the sequence of clinical activities for managing a specific condition, from diagnosis through treatment and follow-up. HTA models often map onto clinical pathways to ensure that cost and outcome estimates reflect real-world practice. For a cardiac care pathway, the model may include