# Weighted Benefit-Cost Analysis and Social Value: Understanding the Difference

Methodological Commentary

Paul Smetanin

July 2026

Canadian Centre for Economic Analysis

# Contents


**1 Context** ..... **3**
**2 What Weighted Benefit-Cost Analysis Does** ..... **3**
**3 What Social Value Measurement Requires** ..... **4**
**4 Where the Gap Lies** ..... **5**
**5 When Results Are Artifacts of Assumptions** ..... **6**
**6 Selective Scope: The Problem of Choosing What Counts** ..... **7**
**7 CANCEA's Approach** ..... **8**
7.1 Empirical well-being foundation ..... 8
7.2 Calculation sequence ..... 9
7.3 What this does not depend on ..... 9
7.4 Why individual-level estimation matters ..... 10
7.5 How weights are treated ..... 10
7.6 Counterfactual discipline ..... 10
7.7 Open scope ..... 10
**8 Applied Evidence** ..... **11**
**9 Implications for Decision-Makers** ..... **11**
**10 Conclusion** ..... **12**

# 1 Context

Decision-makers in infrastructure, transportation, energy, housing, and public policy are increasingly asked to demonstrate the social value of their investments. The expectation is reasonable: public and private capital decisions affect people, communities, and systems in ways that go beyond conventional economic returns.

However, the term "social value" is now applied to a range of methods that differ materially in what they measure, how they measure it, and what they can credibly claim. Not all approaches that use the phrase deliver the same thing. The distinction matters for any organisation choosing an evidence framework for a consequential decision.

This paper clarifies the difference between two classes of approach: weighted benefit-cost analysis and person-level empirical social value estimation.

# 2 What Weighted Benefit-Cost Analysis Does

Weighted benefit-cost analysis (weighted BCA) is an extension of conventional BCA. In a standard BCA, benefits and costs are monetised across predefined appraisal categories and compared to produce a ratio or net present value. Weighted BCA adds a distributional layer: it applies income-derived weights to the existing BCA results so that benefits received by lower-income groups count for more in the final total.

HDR, Inc.'s Social Equity Value Analysis, or SEVA, is a publicly described example of this class. As HDR presents it, SEVA applies a weighted BCA framework using project beneficiaries' incomes to compute weights, which are combined with standard BCA results to produce what HDR describes as a measure of social value. HDR positions the approach as applicable across transportation, water, energy, and broadband, and as relevant for agencies seeking federal funding where grant programs request distributional analysis alongside conventional BCA.

That is what weighted BCA does. It takes an existing benefit-cost result and adjusts it by the income profile of the affected population.

# 3 What Social Value Measurement Requires

Measuring social value is a more demanding task. It requires an evidence-based account of how a project or policy changes the lived conditions, choices, opportunities, risks, and well-being of people, relative to a credible counterfactual.

The structural requirements are:

1. **Person-level estimation.** Social value must be derived from changes experienced by actual people with heterogeneous circumstances, not from an average beneficiary or a single income cell.
2. **A credible counterfactual.** The analysis must compare the project scenario with what would otherwise have happened. Without a defined counterfactual, there is no basis for attributing change to the project.
3. **Pathway traceability.** Results must arise from represented mechanisms: employment, income, affordability, access, health, stress, stability, opportunity, or other project-specific channels. A weighting function applied to a pre-existing BCA total does not trace a pathway.
4. **Distribution shown before netting.** Positive and negative effects must be visible before being collapsed into a single number. Who gains and who loses must be identifiable.
5. **Costs alongside benefits.** Social value is not a benefits-only exercise. Construction disruption, displacement, affordability pressure, environmental effects, fiscal costs, and opportunity costs must be represented where relevant.
6. **System-level coherence.** People, households, jobs, firms, government revenues, infrastructure, geography, and time must be reconciled within one consistent representation. Economic value and social value should be derived from the same system so that double-counting is prevented structurally, not merely assumed away.
7. **Aggregation from the bottom up.** Totals should emerge from person-level and system-level results, not from a modeller-selected weighting function applied to category-level averages.

These are not theoretical preferences. They are the practical requirements for producing evidence that can be explained, challenged, and defended when real consequences depend on it.

# 4 Where the Gap Lies

Weighted BCA, as a class of methods, does not satisfy these requirements.

**It does not estimate person-level welfare change.** Weighted BCA reweights the outputs of a conventional BCA. If the BCA estimates benefits at category level or average-beneficiary level, the weighting step cannot recover person-level heterogeneity that was never estimated.

**It does not trace pathways.** Applying a distributional weight to a monetised benefit total does not explain how people's well-being changes, through which channels, over what time path, or under what conditions.

**It does not prevent omission.** If the underlying BCA omits non-market well-being effects, community pathways, labour-market dynamics, affordability pressure, system interactions, or negative channels, the weighting step cannot recover what was never measured.

**It does not show costs before netting.** A weighted BCA framed around demonstrating benefits to disadvantaged groups can obscure offsetting costs to the same or other populations.

**It does not provide system closure.** Weighted BCA does not automatically reconcile people, households, jobs, firms, supply chains, government revenues, service capacity, geography, timing, and boundary populations within one coherent representation.

**Its results can depend on the weighting function.** The choice of elasticity, income benchmark, and affected group definition are partly normative. Where the weighting function materially determines project ranking or value, the analysis becomes assumption-led rather than evidence-led.

These gaps are not minor. They determine whether an analysis can support strategic decisions, withstand technical review, and be communicated credibly to governments, communities, investors, and boards.

## 5 When Results Are Artifacts of Assumptions

There is a deeper problem with assumption-driven methods that warrants direct attention: in a weighted BCA, the results can be largely determined by the modelling choices themselves rather than by the project's actual effects on people.

If the analyst selects the weighting function, the income benchmark, the elasticity of marginal utility, the definition of who counts as a beneficiary, and the BCA categories from which benefits are drawn, those choices can predetermine the magnitude, direction, and distributional pattern of the result. The output then reflects the assumptions that went in, not an independent empirical finding about what the project does to people's lives.

This is not a theoretical risk. Academic research has demonstrated that in well-being valuation models, changes in how variables are operationalised, which activities are included, and how participation is classified can reverse key findings entirely. Oman and Taylor (2018), in a peer-reviewed replication of a prominent subjective well-being valuation model commissioned for evidence-based policy, found that the original study's headline claims about attendance and happiness were not reproducible. When alternative but equally defensible classifications were applied to the same survey data, the positive relationship disappeared or reversed direction. The authors concluded that "the way in which [variables are] operationalised have major implications for estimated effect size" and that such models are vulnerable to producing "attractive-sounding estimates" that are artifacts of coding decisions rather than evidence of causal relationships.

The implication for infrastructure evaluation is direct. When a method's results depend materially on assumed weights, selected benefit categories, or arbitrary group definitions, the analysis is not measuring social value. It is manufacturing a number whose magnitude is governed by the analyst's parameter choices. A decision-maker relying on such a result cannot know whether the number reflects the project's actual consequences or the modeller's assumptions.

This problem is compounded when the method does not require disclosure of how sensitive the result is to these choices. If a small change in the weighting function, the income benchmark, or the included benefit categories would materially alter the finding, the analysis is fragile. If the client is not shown that fragility, the analysis is misleading.

A defensible social value methodology must produce results that arise from represented system behaviour and observed data, not from normative choices embedded in the calculation structure. The results should be robust to reasonable alternative specifications, and where they are sensitive to assumptions, that sensitivity must be reported.

# 6 Selective Scope: The Problem of Choosing What Counts

A further limitation of methods that extend conventional BCA is the problem of selective scope. In a standard BCA framework, the analyst selects which benefit and cost categories to include. But when the result is presented as "social value," the scope selection becomes consequential in a different way: it determines whose well-being is represented and through which channels.

Methods that allow the analyst to choose which factors matter and which do not introduce a form of prescriptive bias. If the analysis includes travel-time savings but excludes affordability pressure; if it includes employment effects but excludes displacement; if it represents benefits to one group but does not examine costs to another, the result is partial by design.

The same problem arises when methods focus exclusively on benefits. A project that generates benefits for some people while imposing costs on others has a social value that reflects both. If the method is structured to demonstrate positive impacts without accounting for negative ones, the result is advocacy, not measurement.

Peer-reviewed research has identified this pattern in practice. Commissioned econometric models for social value have been shown to select variables, activities, and relationships in ways that reproduce existing funding hierarchies and generate positive headlines, without subjecting those selections to sensitivity testing or alternative specification (Oman and Taylor, 2018). The concern is not that analysts are dishonest, but that the method permits selective presentation to be mistaken for comprehensive evidence.

A credible social value methodology must be open scope. This means it must be capable of representing whatever channels through which a project affects people, whether those effects are positive or negative, expected or unexpected, immediate or delayed. The scope should be determined by the system and the affected population, not by the analyst's choice of which categories to include.

# 7 CANCEA's Approach

CANCEA's methodology is a bottom-up, empirical, person-centred method for estimating changes in well-being and economic outcomes under a project scenario compared with a counterfactual, before aggregation. It addresses each of the requirements described above.

## 7.1 Empirical well-being foundation

The methodology is grounded in satisfaction-with-life survey data covering multiple domains of an individual's circumstances including health, employment, income, housing, family and relationships, community, time use, access to services, environment, and safety. These are the domains through which people themselves report and experience changes in life quality.

Each person modelled in CANCEA's platform carries traits (fixed characteristics such as sex, birth year, and other attributes a person cannot change) and states (changeable circumstances including health conditions, employment, income, housing tenure, family composition, location, commuting pattern, and access to services). Well over a hundred traits and states combine to represent the proper heterogeneity of the individual and the contextual nature of their environment. This representation enables the method to estimate what each person's expected change in satisfaction-with-life would be under a project scenario compared with a counterfactual, given their specific circumstances.

Satisfaction-with-life is not treated as a simple number to be aggregated. It is treated as a scalar output of a wider latent evaluative system. A person's evaluation of their own well-being is shaped by their observable conditions, expectations, reference points, perceived security, opportunities, constraints, and adaptation. The methodology distinguishes between this latent evaluative state and the scalar report that surveys produce from it. The method does not assume that a one-point change in reported satisfaction carries the same welfare meaning for every person regardless of context.

The empirical calibration uses survey evidence from millions of people to estimate the conditional relationship between observable person-environment states and reported well-being. When those circumstances change under a project scenario, the expected change in well-being is estimated from that empirically grounded relationship rather than from an assumed weighting function or a transferred proxy value.

## 7.2 Calculation sequence

The approach follows a strict sequence:

1. Define the decision purpose, project scenario, and counterfactual.
2. Identify the affected population, including boundary populations such as new entrants, displaced groups, or scenario-only populations.
3. Estimate economic, social, spatial, demographic, and environmental pathways over time.
4. Estimate person-level welfare paths under the scenario and the counterfactual.
5. Calculate the difference between those paths.
6. Translate well-being changes into local monetary equivalents where appropriate.
7. Discount monetary equivalents, not raw well-being paths.
8. Aggregate from the bottom up.
9. Report positive and negative effects, affected groups, uncertainty, and sensitivity.

This sequence protects the analysis from assuming a single average beneficiary, monetising benefits without testing overlap, netting results before showing who gains and who loses, and allowing the answer to be determined by an inputted weighting function.

Economic value and social value are derived simultaneously from the same system representation. The economic and social results are different reporting views of a consistent scenario-counterfactual simulation, not independent totals layered on top of each other. This is how double-counting is prevented structurally.

## 7.3 What this does not depend on

The social value result does not depend on an analyst selecting an elasticity of marginal utility, choosing which income groups to weight, or deciding which benefit categories to include. The result emerges from empirically calibrated changes in people's circumstances and their relationship to reported well-being, estimated across the full breadth of life domains rather than a narrow subset selected by the modeller.

The method still contains assumptions. Calibration parameters, valuation functions, counterfactual design, and system boundaries all involve choices. The difference is that these assumptions are tied to observable data, scenario design, and sensitivity testing rather than to a normative weighting function that predetermines the distributional pattern of the result. Where assumptions are present, they are disclosed and sensitivity-tested rather than embedded in the calculation structure.

## 7.4 Why individual-level estimation matters

Infrastructure and public policy projects do not affect just the average person. They affect distinct individuals and households with different incomes, ages, family structures, occupations, locations, commuting patterns, health conditions, and exposure to risk. When analysis begins with averages, competing projects can appear artificially similar because their effects are compressed into average beneficiaries or average demographic cuts.

CANCEA's method preserves heterogeneity until the final reporting stage. A transit investment, a congestion-reduction strategy, a pumped hydro storage project, a housing intervention, and a health investment may all generate social value, but they do so through different people, pathways, timing, and system interactions. A method based on averages or front-end weights obscures those differences. A bottom-up simulation reveals them.

## 7.5 How weights are treated

Distributional weights can serve as a diagnostic or sensitivity test, but they should not determine the answer before project effects have been estimated.

In CANCEA's approach, the simulation first derives economic and well-being impacts for affected individuals. The resulting distribution can then be inspected to understand what implicit weighting or valuation pattern is implied by the simulated outcomes. A model that estimates impacts first and inspects implied weighting patterns afterward is more auditable. It lets decision-makers ask whether the outcome follows from the project, the data, and the represented system, or from modelling assumptions that are difficult to disentangle.

## 7.6 Counterfactual discipline

A credible social value estimate must answer a causal question: what changed because of the project, compared with what would have happened otherwise? The stronger requirement is that both the scenario and the counterfactual be represented in a coherent system with conserved identities, flows, transitions, and boundaries. Otherwise, the analysis may compare two sets of outputs without demonstrating that either represents a feasible system state.

## 7.7 Open scope

The method does not pre-select which factors count. If a project affects commuting stress, housing affordability, employment security, family time, health, income, and community access simultaneously, the method represents all of those channels because the person's circumstances encompass all of them. The scope is determined by the individual's actual situation and the system representation, not by a modelling convenience or an appraisal template.

# 8 Applied Evidence

CANCEA's methodology has been applied across infrastructure, transportation, energy, housing, health, regional planning, and social value evaluations. Examples include:

- **Congestion in Ontario** (*Impact of Congestion in the GTHA and Ontario: Economic and Social Risks,* CANCEA, 2024): The annual cost of congestion in Ontario was estimated at $56.4 billion, including $44.7 billion within the Greater Toronto and Hamilton Area. Of the Ontario total, $12.8 billion was estimated as economic impact and $43.6 billion as social value impact. The analysis separated economic costs from well-being costs and traced effects across commuters, regions, age groups, income groups, occupations, and industries.
- **Pumped Hydro Storage** (*Made-in-Ontario Pumped Hydro Storage: Economic and Social Value Benefits*, CANCEA, 2024): Over a 50-year analysis period, the Ontario Pumped Storage project was estimated to generate $6.8 billion in GDP contribution, 41,200 job-years, $3.9 billion in labour income, $1.4 billion in private investment, and $450 million in social value. Against a battery storage alternative over 20 years, the analysis estimated $3.5 billion more net value, including materially more jobs, rural households, households with children, and full-time employment.

These are published applications of the same methodology described in this paper, applied to real decisions with real stakeholders.

# 9 Implications for Decision-Makers

When choosing an evidence framework for a consequential infrastructure, transportation, energy, housing, or policy decision, the following questions are relevant:

- Does the method estimate how people's well-being actually changes, or does it reweight a conventional BCA total?
- Does the method trace pathways, or does it apply a weighting function to category-level averages?
- Does the method show costs and negative effects before netting?
- Does the method work from a credible counterfactual represented in a coherent system?
- Does the method derive economic and social value from the same system, or are they independent totals that may overlap?
- Can the results be explained to governments, communities, investors, and boards without depending on an opaque weighting assumption?
- Can the assumptions be challenged and tested under technical review?

These questions distinguish social value measurement from distributional adjustment.

# 10 Conclusion

Weighted benefit-cost analysis addresses the distributional question of how a conventional BCA result changes when benefits to different income groups are weighted differently. But that is the limit of its scope. It does not estimate welfare change, trace pathways, model systems, prevent double-counting, or show who gains and who loses before netting.

Social value measurement addresses the broader questions of how a project, policy, or investment changes economic outcomes and people's well-being, for whom, through which channels, over what time path, compared with what counterfactual, and with what positive and negative effects.

Decision-makers facing consequential choices in infrastructure, transportation, energy, housing, and public policy increasingly need the latter. The evidence framework must match the decision.

*CANCEA is an independent Canadian centre for economic and social analysis. CANCEA's social value methodology derives individual-level welfare impacts from an agent-based socio-economic simulation, producing economic and social value simultaneously from one coherent system representation.*

*HDR's Social Equity Value Analysis (SEVA) is referenced in this paper as a publicly described example of the weighted BCA class. The description of SEVA is based on HDR's own published materials.*

*Oman, S. and Taylor, M. (2018) "Subjective well-being in cultural advocacy: a politics of research between the market and the academy," Journal of Cultural Economy, 11(3), 225-243. Referenced for its peer-reviewed demonstration that well-being valuation models can produce results that are artifacts of operationalisation and classification choices rather than empirical findings about causal relationships.*