# Comparing the MAVF and RSE with the proposed Cost-benefit framework

Max Henrion for Level 4 Ventures
Draft 3 Aug, 2022 Summary

- **Level 4 Recommendation:** Level 4 recommends adoption of cost-benefit framework (CBF)to replace MAVF because it is substantially simpler, clearer, and can be based on standard numbers for value of a statistical life (VSL) from US Federal agencies and value of reliability, instead of each utility having to select their own attribute ranges and weights.
- **Simplicity**: The CBF is much simpler to specify and communicate for a framework with *n* attributes, requiring only *n-1* trade-off values compared to *n* upper bounds plus *n* weights for the MAVF.
- **Clarity**: The meaning of the trade-off values, such as Value of a Statistical Life (VSL) or Value of Loss of Load (VLL), are much clearer and easier to explain than the combined ranges and weights needed for the MAVF.
- **BCR shows net benefit**: Ranking mitigations by RSE gives you their relative cost-effectiveness, but it gives no guidance on whether the benefits of risk reduction exceed the costs. The Benefit-cost ratio (BCR) – ratio of benefits to costs -- is analogous to the RSE but has the advantage that BCR>1 implies that benefits exceed costs.
- **Standard guidance:** Having each utility developing their own MAVF, it makes their risk units and RSE difficult to compare. With CBF, the trade-off values, such as VSL, on can be based on standard guidance from US Federal agencies such as DOT and EPA or widely used tools such as LBNLs ICE model for the cost of reliability, instead of each utility, or the CPUC, having to estimate from scratch.
- **Equivalence**: The MAVF and CBF are mathematically equivalent if the CBF uses the trade-off values (such as value of a statistical life or VSL) that are only implicit in the MAVF with linear scaling functions. The difference is that CBF quantifies the value of risk reduction in dollars, where MAVF uses a scaled risk index that varies from one utility to another.
- **Nonlinear economic functions**: Both MAFV and CBF can use nonlinear value functions for attributes–e.g., for reliability, where cost may be a nonlinear function of outage length.
- **Nonlinear risk attitudes:** You can use nonlinear functions with MAFV or CBF to represent ratepayers' attitudes to risk, such as risk aversion.
- **Reference points:** With a nonlinear risk attitude, it is important to choose a clear reference point for comparison – typically, total current risk with approved mitigations, aggregated over all tranches, years, and risk types.
- **Uncertainties:** With either framework it is desirable and practical to propagate uncertainties about risks and mitigation effects through the entire framework. In either case, it is important to explicitly model the dependencies among risks and mitigations.

## Background

In the Settlement Agreement (S-MAP) of 2016, the California Public Utility Commission agreed with Investor-Owned Utilities (IOUs) and intervenors on the *Risk Decision Framework* (RDF) to quantify the benefits of proposed risk mitigations. The RDF includes a *multi-attribute value function* (MAVF) to quantify and combine the benefits of a mitigation in terms of risk reduction in attributes, including safety, reliability, and cost. It uses *Risk-Spend Efficiency* (RSE) – the ratio of the risk reduction to investment cost of a mitigation – as a guide to prioritize proposed mitigation projects.

The CPUC asked Level 4 Consulting to review the RDF and RSE as used in recent RAMP filings by the IOUs. In their response, Level 4 proposed simplifying the MAVF by monetizing each attribute – e.g., converting fatalities into equivalent monetary cost using the value of a statistical life (VSL) – and combining these equivalent costs in a *cost-benefit framework* (CBF). As part of the CBF we suggest using a *benefit-cost ratio* (BCR)–the ratio of monetized benefits of risk reductions to the investment cost of a mitigation – as an alternative to the RSE as a guide to prioritize mitigation projects.

This paper illustrates how to apply the CBF and compare its results with the existing MAVFs used by IOUs in their recent RAMP submissions. Our goal is to clarify the proposed benefit-cost approach, show that it can give equivalent results to the existing MAVF, and compare the two approaches in terms of simplicity and clarity. We end by summarizing ways to handle nonlinear value functions, risk attitude, and modeling uncertainty within these frameworks.

## Multi-Attribute Value Function approach

First let's review the MAVF approach. Figure 1 illustrates an application of an MAVF with three attributes, Safety, Reliability, and Cost. For each attribute, the MAVF specifies natural units—in this case, fatalities, customer-minutes interrupted (CMI), and dollars--a value in natural units, a lower and upper bound to the range of values, a scaling function, and a weight. The black numbers are properties or parameters of the MAVF, in this case based on the MAVF used in the recent RAMP submitted by Southern California Edison.

The blue numbers are specific to a particular risk or mitigation. The value numbers represent a risk or mitigation for each attribute in natural units. The % of range is the value as a percent of the range from zero to upper bound. Each scaling function maps from the % of range to a scaled score from 0 to 100. In this example, the scaling functions are linear (as used by all IOUs except PG&E.) The weighted score is the product of the scaled score and weight. The total risk score is simply the sum of the weighted scores over the attributes.

| Attributes | Natural units | Value | Lower bound | Upper bound | % of range | Scaling function | Scaled score | Weights | Weighted score |
|---|---|---|---|---|---|---|---|---|---|
| **Safety** | Fatalities | 20 | 0 | 100 | 20% | 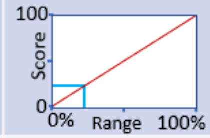 | 20 | × 50% | = 10.0 |
| **Reliability** | CMI<br>Customer-minutes interrupted | 1 billion | 0 | 2 billion | 50% | 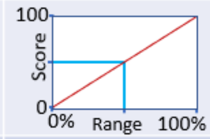 | 50 | × 25% | = 12.5 |
| **Financial** | Dollars ($) | $1 billion | $0 | $5 billion | 20% | 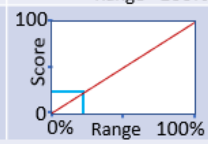 | 20 | × 25% | = 5.0 |
| | | | | | | **Total weighted risk score** | | | **27.5** |

*Figure 1: Example showing a calculation with the existing MAVF framework. The black numbers and the scaling functions for each attribute are elements of this MAVF. Blue numbers represent the specific risk case to which it is applied.*

## Cost-benefit framework

Figure 2 illustrates a corresponding cost-benefit framework. Just like the MAVF, it shows an example value for the three attributes in natural units. Instead of the upper bound and weight, it uses a single *tradeoff value* (or *monetization*) for each attribute. For Safety, it uses $100 million per fatality as the *value of a statistical life* (VSL), which is the trade-off value implied by the MAVFs of the IOUs (see Figure 3). The VSL recommended by US Dept of Transportation is around $10 million, as is the equivalent *value of mortality reduction* (VMR) used by US EPA. For reliability, it specifies a *value of loss of load* (VLL) as $2.50 dollars per customer-minute interrupted. For Financial cost, the “tradeoff value” between dollars and dollars is necessarily 1.00.

*Figure 2: Example showing application of the cost-benefit framework to quantify the total equivalent cost of a risk example. The black numbers are the trade-off values used in this framework. The blue numbers*

| Attributes | Natural units | Example value | Trade-off value | Equivalent value |
|---|---|---|---|---|
| **Safety** | Fatalities | 20 | x $100 million/fatality | = $2.0 billion |
| **Reliability** | Customer-minutes interrupted (CMI) | 1 billion | x $2.50/CMI | = $2.5 billion |
| **Financial** | Dollars ($) | $1 billion | x 1.00 | = $1.0 billion |
| | | | **Total equivalent value** | $5.5 billion |

*represent the case to which it is applied.*

The *equivalent value* for each attribute is simply the product of its value in natural units and its trade-off value. For example, the equivalent value for the Safety attribute is 20 fatalities x $100 million per statistical life = $2 billion. The *total equivalent value* (benefit or cost), $5.5 billion in this example, is the sum of the equivalent values over the attributes.

## Simplicity and clarity

A MAFV with three attributes and linear scaling functions needs specification of six parameters: an upper bound and a range for each attribute, as shown in Figure 1. (The weights must sum to 100%, so there are only five independent parameters.) The corresponding cost-benefit framework has only two parameters, as shown Figure 2--the value tradeoff for safety (VSL or VMR) in $/life and for reliability (VLL) in $/CMI. (The "value tradeoff" for Financial cost is necessarily 1.) For *n* attributes, the MAVF needs 2n-1 parameters, and the CBF needs n-1 parameters.

To understand the importance of an attribute to the result in CBA, you need consider only a single number for each attribute – its trade-off value in terms of dollars per natural unit (e.g., VSL or VLL). For the MAVF you need four numbers to understand the trade-off between any two attributes, as we see in the relationship:

$$VSL = \frac{\text{Weight[Safety] / UpperBound[Safety]}}{\text{Weight[Financial] / UpperBound[Financial]}}$$

Failure to understand this relationship between weights and bounds may explain why the requirement in the Settlement Agreement that the weight for safety must be at least 40% does not have the intended effect. Requiring a minimum value for VSL would be much easier to understand and implement.

## The equivalence of the two approaches

A MAVF implies a dollar trade-off value for each attribute, based on the upper bound and range of that attribute and of the Financial cost attribute. For example, the trade-off value for safety (VSL) is given by this relationship:

$$VSL = \frac{\text{Weight[Safety] / UpperBound[Safety]}}{\text{Weight[Financial] / UpperBound[Financial]}}$$

If we substitute in the MAVF values for SCE, for example, we get:

$$VSL = \frac{50\% \ / \ 100}{25\% \ / \ \$5 \ billion} = \$100 \ million$$

Figure 3 shows trade-off values for each attribute for each utility implied by the upper bounds and weights in its MAVF from their most recent RAMP submissions. PG&E uses nonlinear scaling functions, which cause the trade-off values vary over their range (as we discuss below), so we show the trade-off values averaged over the range of each attribute.

| Utility | Attribute | Units | Upper bound | Weight | Trade-off values |
|---|---|---|---|---|---|
| PG&E | Safety | Fatalities | 100 | 50% | $100,000,000 |
| | Electricity reliability | CMI | 4,000,000,000 | 20% | $1.00 |
| | Gas reliability | Customers affected | 750,000 | 5% | $1,333 |
| | Financial | USD | $5,000,000,000 | 25% | $1.00 |
| SoCal Edison | Safety | Fatalities | 100 | 50% | $100,000,000 |
| | Electricity reliability | CMI | 2,000,000,000 | 25% | $2.50 |
| | Financial | USD | $5,000,000,000 | 25% | $1 |
| SoCal Gas | Safety | Fatalities | 20 | 60% | $100,000,000 |
| | Financial | USD | $500,000,000 | 15% | $1.00 |
| | Stakeholder satisfaction | Index | $100 | 2% | $666,667 |
| | Gas meters | Gas meters w outages | 100,000 | 11.5% | $3,833.33 |
| | Gas curtailment | MMCf | 666 | 11.5% | $575,575.58 |
| SDG&E | Safety | Fatalities | 20 | 60.0% | $100,000,000 |
| | Financial | USD | $500,000,000 | 15% | $1.00 |
| | Stakeholder satisfaction | Index | 100 | 2% | $666,667 |
| | Gas meters | Gas meters w outages | 50,000 | 5.75% | $3,833 |
| | Gas curtailment | MMCf | 250 | 5.75% | $766,667 |
| | Electricity outage | SAIDI | 100 | 5.75% | $1,916,667 |
| | Electricity | SAIFI | 1 | 5.75% | $191,666,667 |

*Figure 3: The implied trade-off values ($ per unit) for each attribute of each utility, based on upper bounds and weights.*

Interestingly, all four utilities have the same implied trade-off value for safety (VSL): $100 million/fatality. But the trade-off values for other attributes vary widely. The trade-off value for electric reliability is $1.00/CMI for PG&E and $2.50/CMI for SCE. The value of gas reliability is $1,333 per customer affected for PG&E, and $3,833/gas meter outage for SoCalGas and $1,917/gas meter outage for SDG&E. The other attributes are hard to compare because they use different units.

Figure 4 compares the MAVF and CBF approaches by applying them to an example mitigation using the MAVF and trade-off values for PG&E and SoCal Edison from Figure 3. Consider a mitigation that reduces risks by 1 fatality for safety, 40 million CMI for reliability, and $50 million financial cost.[1] The columns for each utility headed MAVF show the scaled and weighted risk score for each attribute based on the upper bound and weight from in Figure 3. The row labeled Total shows the sum of these weighted risk scores. The CBF columns show the equivalent monetized value by attribute based on the trade-off values from in Figure 3. The Total is the sum of these benefits.

[1] These numbers are 1% of the upper bound range for PG&E where their scaling functions are linear. See below for discussion of nonlinear scaling functions.

| Attributes | Natural units | Mitigation effect | PG&E * | | SoCal Edison | |
|---|---|---|---|---|---|---|
| | | | MAVF | BCF ($millions) | MAVF | BCF ($millions) |
| **Safety** | Fatalities | 1 | 0.50 | $100 | 0.50 | $100 |
| **Electricity reliability** | CMI | 40,000,000 | 0.20 | $40 | 0.50 | $100 |
| **Financial** | USD | $50,000,000 | 0.25 | $50 | 0.25 | $50 |
| **Total** | | | **0.95** | **$190** | **1.25** | **$250** |
| **Mitigation cost** | $ millions | $100,000,000 | RSE | BCR | RSE | BCR |
| **Readability factor for RSE** | | 1,000,000 | **0.0095** | **1.9** | **0.0125** | **2.5** |
| **BCR/RSE** | $million/risk | | $200 | | $200 | |

*Figure 4: Comparison of MAVF and CBF (Cost-benefit framework) for a risk mitigation that would affect Safety, Reliability, and Financial cost based on the MAVFs of PG&E and SoCal Edison. Under MAVF it shows the risk reduction score, and under CBF the monetized dollar equivalent based on the MAVF parameters and trade-off values from Figure 3. It calculates the Risk-Spend Efficiency (RSE) and Benefit-Cost Ratio (BCR) based on the mitigation cost. For both IOUs, the ratio BCR/RSE is $200 million per risk*

There is a fixed ratio between the scaled risk units in the MAVF columns and the dollar equivalent benefits in the CBF columns. It is $200 million per risk unit for both PG&E and SoCal Edison. (This correspondence is a coincidence. The ratio varies by utility. It is $33.3 million for SoCal Gas and SDG&E.) Viewed this way, the MAVF and CBF results are mathematically equivalent, except for this fixed ratio. The key difference is that CBF quantifies the value of risk reduction in dollars, where MAVF uses a scaled risk index that varies from one utility to another.

## RSE and benefit-cost ratio

The CBF offers the benefit-cost ratio (BCR) of the monetized benefit to the cost of a mitigation, which is analogous to RSE. BCR may be used like RSE as a guide for ranking mitigations according to which is most cost effective. However, BCR has a key advantage over RSE. If the BCR is greater than one for a mitigation, it means that the dollar benefit is greater than its dollar cost. Unlike BCR, an RSE cannot tell you whether any individual mitigation is cost beneficial.

The bottom three rows of Figure 4 calculate the Risk-Spend Efficiency (RSE) and benefit-cost ratio (BCR) for a mitigation with an investment cost of $100 million. The RSE is the ratio of the total MAFV risk score to this cost, multiplied by a common readability factor of 1 million. The BCR is the ratio of the total benefits to the mitigation cost.  As with the risk and benefit components above, the ratio of BCR/RSE is $200 million per risk unit for PG&E and SoCal Edison.

Because of this fixed proportionality, the ratio of RSEs of mitigation A to mitigation B will be the same as the ratio of their BCRs. When you order mitigations by BCR or look at the ratio of their BCRs, you get it identical guidance for prioritizing mitigations as when you do this with RSEs.

## Nonlinear scaling functions and risk attitude

The Settlement Agreement (S-MAP) allows utilities to specify nonlinear scaling functions for each attribute. PG&E is the only utility to avail themselves of this option in recent RAMP submissions. Figure 5 shows an example application of the MAVF with nonlinear scaling functions. The difference from the MAVF in Figure 1 is the nonlinear scaling functions for each attribute that give a scaled score that is not directly proportional to the % of range of the value for each attribute.

| Attributes | Natural units | Value | Lower bound | Upper bound | % of range | Scaling function | Scaled score | Weights | Weighted score |
|---|---|---|---|---|---|---|---|---|---|
| **Safety** | Fatalities | 20 | 0 | 100 | 20% | 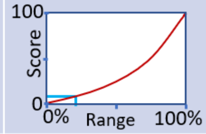 | 10 | × 50% | = 5.0 |
| **Reliability** | CMI Customer-minutes interrupted | 1 billion | 0 | 2 billion | 50% | 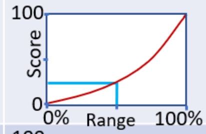 | 24 | × 25% | = 6.0 |
| **Financial** | Dollars ($) | $1 billion | $0 | $5 billion | 20% | 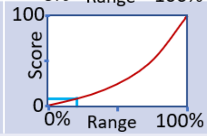 | 10 | × 25% | = 2.5 |
| | | | | | | | | **Total risk score** | 13.5 |

*Figure 5: An example MAVF with nonlinear scaling functions for each attribute.*

PG&E uses the same nonlinear scaling function for each attribute. The function is designed so that the first 1% of an attribute range – e.g., the first fatality– has one tenth of the risk score each fatality between the 10$^{th}$ and 100$^{th}$. It is linear from 0 to 1% and from 10% to 100%.

This kind of scaling function represents *risk aversion* – it treats larger disasters as proportionately worse than smaller disasters between 1% and 10% of the range. Decision analysts call this kind of nonlinear function a *utility function*. In this context where key actors are utility companies, we prefer to term it a *risk-attitude function* to avoid confusion.

Nonlinear scaling functions cause the implied trade-off value to vary with the value of the attribute. Figure 2 shows the trade-off values got PG&E as an *average* over the entire range – e.g., $100 million VSL. But the actual trade-off value for an incremental fatality varies from $10 million per fatality relative to a risk with no fatalities and cost above $500 million (10% of the cost range), or almost $1 billion per incremental fatality relative to a risk with more than 10 fatalities but cost below $5 million.

## When value is not linear in natural units

For PG&E and SCE, the natural units for electric reliability are customer-minutes interrupted (CMI). This metric simply sums the lengths of short and long outages. It counts 1000 customers experiencing a 1-hour outage the same as 50 customers with a 20-hour outage. Both contribute 1000 customer hours interrupted (60,000 CMI). Surveys of electricity customers suggest that they typically consider longer outages as proportionately worse than short outages. For that reason, the ICE Calculator for the cost of electric outages uses a quadratic cost function with a term proportional to the square of the outage time.

In the future, if IOUs decide to adopt a more sophisticated metric for reliability, such as that used in ICE, they may use a non-linear function to represent such nonlinear ratepayer preferences. In such cases, they may use nonlinear scaling functions for one or more attributes, to map from natural units to dollar cost equivalent – instead of a fixed trade-off value for each attribute.

# Nonlinear functions in a Cost-benefit framework

A CBF could also use nonlinear scaling functions to represent ratepayer risk attitudes for individual attributes. This is similar to how PG&E currently uses non-linear scaling functions, except that in a cost-benefit framework, the functions would map directly from natural units to equivalent dollar value – rather than to scaled risk units. However, it would be simpler, and more consistent with standard practice in decision analysis, to use a single risk-attitude function to adjust the total equivalent value for a risk or mitigation to represent risk attitude – such as risk aversion. Figure 6 shows an example with a nonlinear trade-off function from a distribution of outages of various lengths to equivalent cost, and a non-linear risk-attitude function that adjusts the total equivalent cost.

| Attributes | Natural units | Example value | Trade-off values | Equivalent cost |
|---|---|---|---|---|
| **Safety** | Fatalities | 20 | x $100 million VMR | = $2 billion |
| **Reliability** | Customer-minutes interrupted (CMI) tabulated by outage length, *t* | 250 million (*t* = 0 to 60 min) 250 million (*t* >60 min outage) | 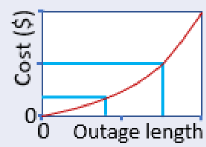  | $500 million |
| **Financial** | Dollars ($) | $500 million | x 1 | = $500 million |
| | | | **Total equivalent cost** | $3 billion |
| | | | **Risk-attitude function** | 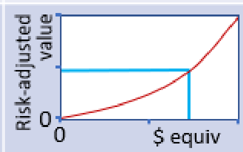  |
| | | | **Risk-adjusted value** | $3.5 billion |

*Figure 6: Cost-benefit framework with a nonlinear attribute (Reliability) and a risk-attitude function that adjusts the equivalent cost for risk.*

# Selecting a reference value with a risk-attitude function

When you use a nonlinear risk-attitude function, it is important to select the correct *reference point* against which to quantify the cost of risks or the value of a mitigation that reduce risks. The most obvious reference point might appear to be zero for each attribute. But in that case mitigations that reduce risk would have negative values. The nonlinear scaling functions are not designed to handle negative values. Preferably, the reference point should be the total risk with no mitigations (or only mitigations already approved or implemented), at some point between 0 and 100 for each attribute. In that case, the value of a single mitigation would is based on the reduction (or increase) in risk relative to that reference point. This is important because the choice of reference point will affect the risk score using a nonlinear function.

Mitigations may be considered at *multiple levels or tranches*. For example, line hardening may be analyzed by individual route, by region, or aggregated over the entire territory. A nonlinear risk attitude, whether at the attribute level or for total equivalent cost, should apply to the entire territory of a utility. Thus, to properly estimate the reference point (base risks) and effect of a proposed portfolio of mitigations, it is essential to aggregate all and all mitigations over the entire territory before you apply an MAVF or a benefit-cost analysis with a nonlinear risk attitude.

The risk reduction and RSE of a mitigation—or its benefit and BCR--should be calculated using net present value of risk reduction and costs over the lifetime of the mitigation. If you are comparing multiple mitigations with different implementation times and/or lifetimes, it is important to extend the time horizon so that it covers all mitigations under consideration. The reference point should be the NPV risk with all approved or implemented mitigations within that time horizon, excluding the mitigations under consideration.

Many mitigations interact with each other. Two mitigations may be competitive – such as enhanced vegetation management and covered conductors – so that the risk reduction of both together is less than the sum of their separate contributions. Or two mitigations may be synergistic, where their total effect is greater than the sum of the separate contributions. Thus, to estimate the risk reduction (benefit) of a mitigation it is usually best to estimate the effect of adding it to a portfolio of other mitigations. In this case, the reference point for each mitigation would be the risk after implementing the portfolio.

In summary, when you use a nonlinear scaling function in MAVF or a risk-attitude function with CBF, it matters how you define the reference point. Normally, the reference point should:

- Sum risks over a combined time horizon using NPV risk
- Aggregate over all tranches up to the entire territory.
- Sum over all risks and approved or implemented mitigations and all mitigations under consideration except subtracting or adding individual mitigations for analysis

These are needed whenever we use a nonlinear risk attitude, with MAVF or CBF. On the other hand, if we restrict ourselves to linear risk attitudes, these challenges of finding the right reference value do not arise (even if we use nonlinear cost functions as for reliability).

## Treating uncertainty in MAVF or BCA

So far, we have ignored uncertainty. The estimate of the risk before the mitigations in the proposed RAMP – i.e., the reference point -- will undoubtedly be uncertain. The risk reduction due to each mitigation will also be uncertain. In recent RAMP reports, IOUs have explicitly represented uncertainty in the risk attributes for mitigations using probability distributions. Their models use Monte Carlo simulation to estimate the uncertainty in the resulting attributes. However, they then use the expected value (mean) of each sample distribution for the attribute values *before* applying MAVFs to value the risk reduction. This gives incorrect results when the MAVF contains nonlinear scaling functions – or with a cost-benefit framework with nonlinear risk attitude. To handle this properly, it is necessary to propagate the uncertainties using Monte Carlo or other technique through the MAVF or benefit-cost model. Even with linear risk attitudes, it is important to estimate the uncertainties to provide guidance on how probable it is that the RSE or BCR of one mitigation is greater than another. Explicit treatment of uncertainty also supports probabilistic analysis of uncertainties to identify which sources of uncertainty have the largest effect on these differences.

Usually, the reference risk and the risk reduction due to a mitigation will be strongly dependent. For example, uncertainty about the incidence and magnitude of wildfires will affect both the reference risk and the effect of line hardening mitigations. To properly estimate the uncertainty about the benefit of the mitigation, it is essential to model this probabilistic dependence.

Similarly, when aggregating risks or mitigations over tranches and over time periods to estimate the uncertainty about their combined effect for the entire utility, it is important to model their probabilistic dependence. For example, uncertainty about the effects of climate change creates probabilistic dependence among future temperatures, frequency of windstorms, and dryness of vegetation over time and across regions within a territory. This results in dependencies among risks and mitigation effectiveness over time and space. These dependencies must be considered during aggregation. Fortunately, if these dependencies are modeled appropriately, Monte Carlo simulation takes account of them properly during aggregation and when the distributions are propagated whether we use an MAVF or a CBA framework.