Article from:

ARCH 2014.1 Proceedings

July 31-August 3, 2013

Portfolio Management with the Critical Event Cost Method

John A. Major, ASA[1] and Sherry L. Thomas[2]

Abstract

The critical event cost method is a technique for assessing the impact that changes to the composition of a portfolio have upon the probability distribution of financial outcomes. The method focuses on certain "critical" events or scenarios underlying a financially interesting range of outcomes in the loss distribution. The method allows for real-time impact estimation in a distributed underwriting environment, prior to effecting such changes. It also allows for planning and budgeting the totality of such impacts. The method is defined and illustrated in the context of property catastrophe portfolio management using commercial catastrophe models, for both occurrence (OEP) and aggregate (AEP) measures of loss. Generalizations of the critical event cost are seen to correspond to estimators of the change (gradient) in various popular portfolio risk measures, including Value at Risk (VaR, also known as Probable Maximum Loss, PML), Tail Value at Risk (TVaR, also known as Conditional Tail Expectation, CTE), and, more generally, so-called spectral measures or distortion measures. In its elementary binary form, the critical event cost can be viewed as an estimate of the change in expected payout of a certain type of (possibly hypothetical) portfolio reinsurance program. As a result, the method provides a logical basis for technical risk pricing.

Keywords: Catastrophe, Portfolio Management, Value at Risk, Probable Maximum Loss, Gradient

1. Introduction

Portfolio management is the process of monitoring and changing the composition of a portfolio of assets or liabilities that are subject to random changes in value, in order to control its risk and expected return characteristics. We will focus on portfolios of exposed properties at risk of loss from natural perils such as wind and earthquake in the context of property/liability insurance. One challenge is to find a risk accumulation control approach that can:

- be implemented within an existing business process or framework;
- allocate risk-bearing capacity, and its consumption, to underwriting unit;
- monitor the evolution of the portfolio during the underwriting period, using simple query tools and reports;
- be integrated into pricing before, as opposed to after, binding of contracts; and
- directly contribute to improving portfolio risk-adjusted returns.

The critical event cost ("CEC") method is such a technique.

This paper is organized as follows. The remainder of the introduction provides definitions of key terms and introduces a simplified, artificial example of catastrophe model output. Section 2 defines and

[1] Senior Vice President, Director of Actuarial Research, Guy Carpenter & Company, LLC; 1166 Avenue of the Americas, New York, NY, 10036; john.a.major@guycarp.com

[2] Managing Director, Head of Catastrophe Management - Americas, Guy Carpenter & Company, LLC; 3600 Minnesota Dr. Suite 400, Minneapolis, MN 55435; sherry.l.thomas@guycarp.com

illustrates using critical event cost calculations to assess the portfolio. Section 3 discusses the relation between CEC and other risk metrics, and how CEC and its generalizations estimate changes in those risk metrics. Section 4 returns to a discussion of applying the CEC method to underwriting and pricing. Section 5 addresses how to apply the CEC method to annual aggregate losses. Section 6 discusses the limits imposed by uncertainty in the underlying risk models. Section 7 concludes.

1.1 Definitions

"Event" refers to an occurrence of a natural disaster (catastrophe) which has the potential to inflict losses on a portion of the portfolio of properties. Because the method is applied as part of simulation-based risk management, these events will typically be simulated events. As far as the method is concerned, each event is fully characterized by its loss impact on each element (insured property) of the portfolio. In practice, the portfolio may be disaggregated into components (groups of elements) as the fundamental unit. For example, an analysis might be desired that focuses on business operating groups, lines of business, or geographic units. In that case, event loss impacts on elements are aggregated up to the component level and the method deals with the broader aggregates as fundamental units of analysis. Another application of the method is to analyze the types of events themselves. In this case, losses are distinguished not by the elements upon which they fall but by the nature of the event causing them. Combinations of event types and portfolio components may also be analyzed.

"Cost" refers to the actuarial expected value of the losses inflicted by the events. Fundamentally, the method is about apportioning the cost to the units of analysis. In this way, the total cost may be understood and managed. The sources of historical total cost or the future expected total cost (if business continues as usual) may be identified, and their relative importance measured. To effect a change in the future cost, management may design changes to the business process aimed at limiting the cost in key components of the portfolio.

"Critical" refers to the identification of certain events as more important than others. In its simplest form, the critical event cost method identifies a subset of events and measures the cost of *only those events* on the portfolio components. More generally, the method applies weights to the events, making some more and some less important on a sliding scale. For example, events causing a loss under some threshold amount may be ignored. Events in a band around a key risk metric, such as a ratings downgrade, may be identified as highly important. Events causing extremely large losses ("doomsday scenarios") might be considered less important, but perhaps not ignored altogether.

The method is intended to augment existing risk management practices. We assume an existing infrastructure that includes the mathematical simulation modeling of portfolio losses with identifiable events having well-articulated loss consequences to relevant components of the portfolio.

1.2 Example cat model output

Exhibit A illustrates, in a simplified and artificial form, what typical cat model output looks like. Each row except for the first represents a single simulated event and its impact on a portfolio of insured properties in the United States. In real cat models, there are tens or hundreds of thousands of events.

Events are identified by a unique numeric identifier ("EVT ID") and are described in terms of type of natural peril and principal geographic area affected ("Description").

Each event has an associated probability ("Prob") in this table. The loss ("Loss") is the simulated total of insured claims paid because of the event, in millions of dollars. In some cat models all the events have the same probability, and in others they differ. Both approaches are valid.

In this table, events have been sorted on loss, from low to high. The cumulative probability ("Cum Prob") in any row is the sum of the probabilities of all events causing that amount of loss or less. Often you will see tables sorted from high to low; instead of cumulative probability, the running sum of probabilities is the exceedance probability.

The first row represents an aggregate of the many events with losses less than $2mm. Collectively, those events occur with 71% probability and their average loss is $0.5mm.

This table represents only a tiny bit of what is available from a typical catastrophe model. For each event, a wealth of detail is potentially available. Losses to defined subsets of the portfolio, perhaps even down to the policy and site level, can often be obtained. This example table will be expanded below.

From this table we can compute the overall expected loss, also known as Average Loss (AL). It is calculated as the sum-product of probability and loss across all of the events: AL = (0.71x0.5) + (0.03x2) + (0.02x3) +…+ (0.01x9.5) = 1.886. To satisfy the actuarial requirement that premiums cover losses and expenses, premiums attributed to catastrophe losses need to be at least as high as the overall AL. However, risk margins are also required to compensate investors for the possibility that premiums will not cover particular losses in some events. Pricing will be discussed in more detail in section 4.2.

1.3 Probability versus frequency, occurrence versus annual

The technically informed reader will object to the previous discussion because real catastrophe models use frequency, not probability, and the total frequency typically exceeds one. For example, severe thunderstorms happen quite often, with many more than one expected in a given year. One could substitute frequencies for probabilities in the foregoing, and a meaningful approach to critical event cost would result. However, the discussion in section 3 would become somewhat tortured, because the interpretations of common risk metrics such as TVaR and VaR would become strained, possibly beyond the point of recognition.

For the purposes of the main discussion of this paper, the event table should be interpreted as representing occurrence exceedance probabilities (OEPs). Specifically, the table should be taken to mean that the *worst* event of a year (measured by portfolio loss) happens with the specified probability. For example, there is a 90% chance that the worst event of a year has a $6mm loss or less. This also gives the same probability that *all* events in the year, measured singly, each have losses of $6mm or less.

This is in contrast to aggregate exceedance probabilities (AEPs). There, the question is, what is the probability distribution of the total (aggregate) loss from all events, combined, in a year? Applying the CEC method to annual aggregate losses will be addressed in section 5.

2. The critical event cost method

In this example, the focus is on losses between $3.5 and 7.5 million. Events causing losses in this range are assumed to be the *critical events*. It can be calculated that, collectively, these events occur with 16% probability. Events with losses below $3.5mm are considered to be within the firm's ability to absorb in the normal course of business. These occur with probability 79%. Losses above $7.5mm are considered too rare and too extreme to manage; they occur with probability 5%. (Later, the discussion will turn to how these events may also enter into the analysis.) In a real application, the critical events are more likely to be in the upper 0.25% to 5% probability range.

Exhibit A: Cat Model Output

EVT ID	Prob	Cum Prob	Loss ($mm)	Description
(many)	(many)	0.71	0.5(ave)	(aggregate of many)
101	0.03	0.74	2	Tornado/Hail – TX
102	0.03	0.77	3	Earthquake – CA
103	0.02	0.79	3.2	Earthquake – Midwest
104	**0.04**	**0.83**	**4**	**Winter Storm – Northeast**
105	**0.01**	**0.84**	**5.2**	**Earthquake – CA**
106	**0.04**	**0.88**	**5.8**	**Hurricane – Gulf**
107	**0.02**	**0.9**	**6**	**Hurricane – NC**
108	**0.04**	**0.94**	**6.4**	**Earthquake – Midwest**
109	**0.01**	**0.95**	**7.2**	**Hurricane – NY**
110	0.02	0.97	8	Earthquake – CA
111	0.02	0.99	8.5	Hurricane – FL
112	0.01	1	9.5	Earthquake – WA
			1.886	Average Loss
			0.892	Critical Event Cost

Source: Guy Carpenter & Company, LLC

Let us define the key metric of this paper: the *critical event cost* (CEC). The critical event cost is the expected loss from critical events. Similar to the AL, it is calculated as the sum-product of probability and loss, but only across the critical events: CEC = (0.04 x 4) + (0.1x 5.2) + ... + (0.01 x 7.2) = 0.892.

Cat analysts are accustomed to computing the AL on subsets of the portfolio. The CEC can be decomposed in exactly the same way.

The events with losses between $3.5 and 7.5 million are "critical" because losses of this magnitude are the focus of aggregation control and underwriting capacity. **For the purposes of illustration, we will assume that management has defined "cat underwriting capacity" to mean a limit on CEC of $1.0mm.** In a real application of CEC, a more sophisticated and nuanced definition of capacity would be used, but a limit on CEC would be a key ingredient.

2.1 How do wind and earthquake events compare?

We can separately identify events as being associated with earthquake or wind. In this way we can decompose the AL and CEC into earthquake vs. wind components. This is done by the same sum-product calculation as before, however only events of the relevant type enter the calculation. For example, the earthquake component of AL uses events 102, 103, 105, 108, 110, and 112 and the earthquake component of CEC uses events 105 and 108. We will assume that among the many small events not broken out separately, the average loss to earthquake is $0.3mm and to wind is $0.2mm.

The results are summarized in table B. Overall, as measured by AL, the risk of loss is almost perfectly balanced between earthquake and wind, and 49% and 51%, respectively. However, in terms of capacity (i.e., among the critical events), wind dominates with 65% of CEC vs. earthquake's 35%. As a fraction of the total capacity budget (recall, we assume a budget constraint of CEC ≤ $1.0mm), wind takes up 58% and earthquake takes up 31%, leaving 11% remaining for growth.

Exhibit B: Earthquake vs. Wind

	Total	EQ	Wind
AL	1.886	0.930	0.956
Percent		49%	51%
CEC	0.892	0.308	0.584
Percent		35%	65%

Source: Guy Carpenter & Company, LLC

2.2 How do lines of business compare?

In order to address this question, we will need more detail in the loss table. Exhibit C disaggregates losses into three lines of business: Homeowners, Commercial Multi-Peril , and Worker's Comp.

Exhibit C: Line of Business

EVT ID	Prob	Loss ($mm)	Description	HO	CMP	WC
<many>	<many>	0.5	<aggregate of many>	0.25	0.12	0.13
101	0.03	2	Tornado/Hail – TX	1.00	0.50	0.50
102	0.03	3	Earthquake – CA	1.50	0.75	0.75
103	0.02	3.2	Earthquake – Midwest	1.60	0.80	0.80
104	**0.04**	**4**	**Winter Storm – Northeast**	**1.00**	**2.00**	**1.00**
105	**0.01**	**5.2**	**Earthquake – CA**	**2.60**	**1.30**	**1.30**
106	**0.04**	**5.8**	**Hurricane – Gulf**	**2.90**	**1.40**	**1.50**
107	**0.02**	**6**	**Hurricane – NC**	**3.00**	**1.50**	**1.50**
108	**0.04**	**6.4**	**Earthquake – Midwest**	**1.60**	**3.20**	**1.60**
109	**0.01**	**7.2**	**Hurricane – NY**	**1.80**	**3.60**	**1.80**
110	0.02	8	Earthquake – CA	4.00	2.00	2.00
111	0.02	8.5	Hurricane – FL	4.30	2.10	2.10
112	0.01	9.5	Earthquake – WA	4.80	1.40	3.30
			AL	0.823	0.578	0.486
			Percent	44%	31%	26%
			CEC	0.324	0.343	0.225
			Percent	36%	38%	25%

Source: Guy Carpenter & Company, LLC

Overall, the AL breakdown shows that the Homeowners line of business has the largest share of expected losses at 44% compared to only 31% for the runner-up CMP line. However, when it comes to capacity, as measured by the critical event cost, Home and CMP are more even, with 36% and 38% shares, respectively.

2.3 How much underwriting capacity is consumed by a particular account?

In order to address this question, we will again need more detail in the loss table. Exhibit D shows the simulated losses to two prospective accounts, in thousands of dollars. Both are commercial accounts with multiple locations. Account AAA is located mostly in the south and west; account BBB mostly in the north and east.

Exhibit D: Two Prospective Accounts

EVT ID	Prob	Description	Acct AAA*	Acct BBB*
<many>	<many>	<aggregate of many>	5	5
101	0.03	Tornado/Hail – TX	20	-
102	0.03	Earthquake – CA	30	-
103	0.02	Earthquake – Midwest	16	34
104	**0.04**	**Winter Storm – Northeast**	**-**	**85**
105	**0.01**	**Earthquake – CA**	**52**	**-**
106	**0.04**	**Hurricane – Gulf**	**58**	**-**
107	**0.02**	**Hurricane – NC**	**60**	**32**
108	**0.04**	**Earthquake – Midwest**	**19**	**68**
109	**0.01**	**Hurricane – NY**	**-**	**76**
110	0.02	Earthquake – CA	56	-
111	0.02	Hurricane – FL	85	-
112	0.01	Earthquake – WA	-	101
		AL	**13**	**13**
		CEC	**4.8**	**7.5**
		* Losses are $000		

Source: Guy Carpenter & Company, LLC

Both accounts have exactly the same overall AL. Say the premiums and expenses are the same, so that on a purely actuarial expectations basis, the expected profitability is identical between the two prospects. Does this mean they are equally attractive?

Imagine the underwriter has been budgeted a CEC limit of 10,000, or 1% of the total CEC budget for the firm. Account AAA will consume 480, or 4.8% of her budget, whereas account BBB will consume 7.5%, over half again as much. If she was near her limit, she would certainly prefer account AAA (all else being equal).

3. Relation of CEC to other risk metrics

3.1 Risk preferences

The notion of critical events is closely related to a theory of risk aversion or risk preferences. If one were to ask management to "draw a picture" of their concerns about losses of various sizes, something like Figure E might emerge.

Figure E: Risk Aversion to Losses

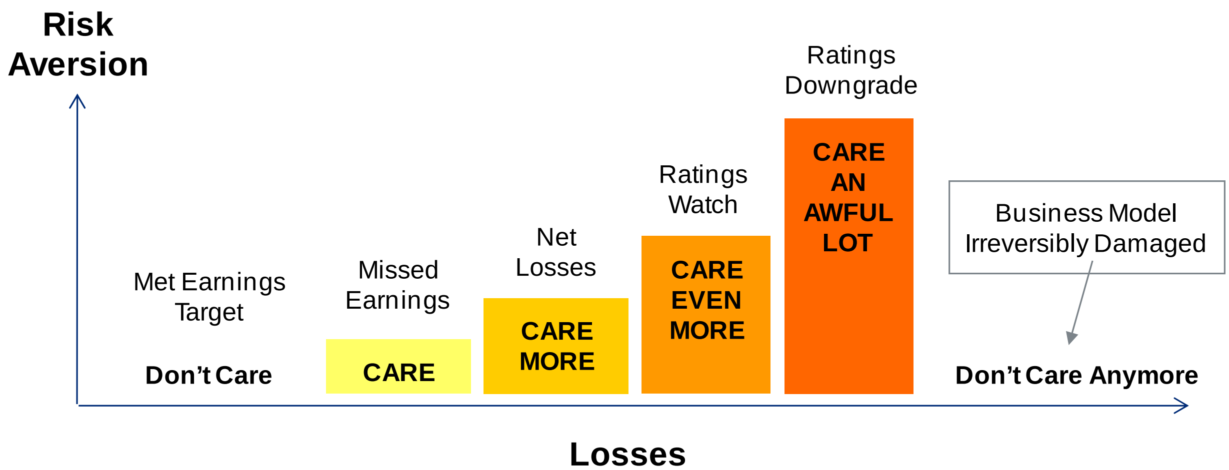

Source: Guy Carpenter & Company, LLC

Scenarios involving low levels of losses, corresponding to results near or better than expectations, are not an area of concern for management. Missing an earnings target, however, is. More extreme losses can lead to an income statement showing not profits, but net losses. This situation is even worse, and is definitely something to be avoided. But even worse is possible. A significant loss to surplus, e.g., from a sizeable catastrophe, could cause the firm's ratings status to be at risk (on watch) or even downgraded. At that point, the reputation of the firm, and its ability to continue to operate profitably, have sustained long-term damage. Beyond that point, large losses are associated with "doomsday" scenarios that are considered too big (and too improbable) to manage in any meaningful way.

One might choose to identify the events associated with ratings watch and downgrade as "critical events" and budget the company's exposure to them, as described in the previous section. However, the notion of varying degrees of risk aversion leads to a natural extension to the critical event cost method.

The first step is to recognize an alternative characterization of the AL and CEC computations. Recall that both are sum-products of loss and probability; the difference between them being the set of events over which the sum is carried. Instead, we can consider them to be sum-products over *all* events, but with a third factor multiplied with loss and probability – call it the risk coefficient. Exhibit F illustrates with the columns labeled "AL" and "CEC". For AL, the risk coefficient is one for all events. For CEC, the risk coefficient is one for the critical events and zero for all other events.

Having a third factor opens up a world of possibilities. The varying degrees of risk aversion in figure E could be represented by risk coefficients in the column labeled "Step". Risk functions of this general form arise in the theories of spectral measures, distortion functions, and utility functions.[3] This provides a more tailored expression of risk attitude than does the simple binary choice of the elementary CEC approach. For example, consider two accounts, one of whom has losses concentrated in events 104-105 and the other in events 108-109, and say they have identical CEC when calculated using the zero and one coefficients of the "CEC" column. The second account will be seen to have a higher computed cost when using the risk coefficients in the "Step" column because the "Step" coefficients of 1.0 for events 108-109 are greater than the coefficients of 0.4 for events 104-105.

The last two columns, "TVaR" and "VaR", represent commonly used portfolio risk metrics. These deserve a closer look.

VaR stands for Value at Risk, and is also known as PML or Probable Maximum Loss. It is defined as that amount of loss that is not exceeded with a specified level of probability. By examining the cumulative probability column in exhibit F, we can see that a loss of 6 is not exceeded with probability 0.9. This makes 6 the 90% VaR for this portfolio.

TVaR stands for Tail Value at Risk, and is also known as CTE or Conditional Tail Expectation.[4] It is defined as the (probability-weighted) average of losses equal to or greater than the VaR (at the specified level of probability). Thus, we can compute the 90% TVaR of this portfolio as follows: [(0.02x6) + (0.04x6.4) +…+ (0.01x9.5)] / [0.02+…+0.01] = 0.873/0.12 = 7.28. Notice that the numerator is the same as the CEC computed on all events with losses greater than or equal to $6mm (i.e., the original definition). It is also equal to the three-factor sum-product using the risk coefficients in the column labeled "TVaR" (i.e. the alternative definition).

This provides a link between TVaR and CEC with the TVaR risk coefficient pattern. Except for a scaling constant (the exceedance probability), they measure risk in exactly the same way.

Exhibit F: Risk Aversion Schedules

EVT ID	Prob	Cum Prob	Loss ($mm)	Description	AL	CEC	Step	TVaR	VaR
<many>	<many>	0.71	0.5	<aggregate of many>	1	0	0.0	0.0	0.0
101	0.03	0.74	2	Tornado/Hail – TX	1	0	0.0	0.0	0.0
102	0.03	0.77	3	Earthquake – CA	1	0	0.2	0.0	0.0
103	0.02	0.79	3.2	Earthquake – Midwest	1	0	0.2	0.0	0.0
104	0.04	0.83	4	Winter Storm – Northeast	1	1	0.4	0.0	0.0
105	0.01	0.84	5.2	Earthquake – CA	1	1	0.4	0.0	0.0
106	0.04	0.88	5.8	Hurricane – Gulf	1	1	0.6	0.0	0.5
107	0.02	0.9	6	Hurricane – NC	1	1	0.6	1.0	1.0
108	0.04	0.94	6.4	Earthquake – Midwest	1	1	1.0	1.0	0.5
109	0.01	0.95	7.2	Hurricane – NY	1	1	1.0	1.0	0.0
110	0.02	0.97	8	Earthquake – CA	1	0	0.0	1.0	0.0
111	0.02	0.99	8.5	Hurricane – FL	1	0	0.0	1.0	0.0
112	0.01	1	9.5	Earthquake – WA	1	0	0.0	1.0	0.0

Source: Guy Carpenter & Company, LLC

[3] There are subtle differences between these which are outside the scope of this paper.
[4] There are also subtle differences between TVaR and CTE which are outside the scope of this paper.

In theory, we could construct the same sort of relationship between VaR and CEC with a suitably chosen pattern of risk coefficients. The risk coefficients would equal 1 for all events where the loss is exactly equal to the VaR and zero elsewhere. While this makes sense in theory, in practical cat modeling applications there will not be many events with losses *exactly* equal to the VaR. Usually it is one event or even no events at all.[5]

To remedy this difficulty, we interpret VaR as being a weighted average loss among events with losses *near* the VaR. The column labeled "VaR" provides an example of such a risk coefficient pattern. In a realistic cat model with 10,000 events, we might be considering a 99.5% VaR. There would be only 50 events with losses at least that big entering a TVaR calculation. That might not be enough for numerical stability, especially when subportfolios are examined. With 100,000 events, there would be a more comfortable 500 events in the upper tail. At least several hundred events should be used to define a band around the VaR level. If weights are used, partially-weighted events should receive a fractional count. For example, one could take the 750 events between the 99.1% and 99.85% VaRs, weight the first and last 250 by 0.5 and the middle 250 by 1.0. This would give an "effective" count of 500 events around the 99.5% VaR. The column labeled "VaR" in exhibit F illustrates this with an effective count of two events around the 90% VaR.

3.2 Changes in risk metrics

Averaging losses from events around a particular VaR, after first figuring out what that VaR is, seems like a redundant and pointless exercise. The significance of doing this comes from analysis of subportfolios, to be discussed in the next section.

Key questions in portfolio management are of the form: What happens when I change my portfolio? How does expected profitability change? How do my risk metrics change?

Consider TVaR. Many firms define their risk tolerance in terms of a limit on TVaR. Portfolio studies are often run to "what-if" adding or subtracting blocks of business, to see the resulting changes in profitability and the change in risk, i.e. TVaR. But we do not have to run entire as-if portfolios. We can instead appeal to theorems of probability theory to see that the effect of a small change in each event's loss on TVaR is just the usual CEC-style calculation of TVaR, but using *those event-by-event loss differences* instead of the original portfolio losses.

This can be justified informally as follows.[6] TVaR is the sum-product of loss, probability, and the TVaR zero-one risk coefficient. For a set of small event-by-event changes in losses, the same events will be identified as being in (coefficient=1) or out of (coefficient=0) the TVaR tail. So the new TVaR differs from the old TVaR only by the event-by-event loss changes, multiplied by the probabilities. Imagine adding account AAA (exhibit D) to the portfolio. Recall the original TVaR is [(0.02x6) + (0.04x6.4) +…+ (0.01x9.5)] / [0.02+…+0.01] = 0.873/0.12 = 7.28. Adding in account AAA, the calculation becomes [(0.02x6.0060)) + (0.04x6.4019) +…+ (0.01x9.5000)] / [0.02+…+0.01] = 0.878/0.12 = 7.32. This differs from the original TVaR by [(0.02x0.0060)) + (0.04x0.0019) +…+ (0.01x0.0000)] / [0.02+…+0.01] = 0.005/0.12 = 0.04. Calculating TVaR on the new portfolio in effect recalculates the original portfolio and adds the difference. We can instead calculate the difference directly.

---

[5] The probabilities may work out in such a way that the desired VaR must be interpolated between two simulated losses.

[6] See Major (2004) or Kreps et al. (2006) for thorough discussions and rigorous proofs.

Other firms define their risk tolerance in terms of a limit on VaR. Surprising to some, this same technique works on VaR. The CEC of the loss differences, using the VaR risk coefficient pattern, estimates the change in VaR resulting from the change in the loss distribution.

What about the original CEC? This can be shown to have the following interpretation. The firm holds capital to protect policyholders against the risk of total losses exceeding total premiums.[7] Referring back to exhibit A, say premiums are sufficient to pay losses up to $3.5mm, and there is $4mm of capital. Then capital is, in effect, funding a "4 excess of 3.5" reinsurance program. At $3.5mm, the program attaches and capital starts being used to fund losses. At $7.5mm, the capital has been exhausted and no further payments are available. The events associated with losses "inside the layer" are identified as the critical events in exhibit A.

The actuarial expected payout from this "program" is calculated as (0.04x(4-3.5)) + (0.01x(5.2-3.5)) +…+ (0.01x(7.2-3.5)) + (0.02x4) + (0.02x4) + (0.01x4) = 0.532. This can be rewritten as {0.05x7.5 – 0.21x3.5} + [(0.04x4) + (0.01x5.2) +…+ (0.01x7.2)] ={-0.360} + [0.892] = 0.532. The first term, in curly braces, is a function only of the attachment and exhaustion amounts and probabilities. The second term, in square brackets, we may recognize as the CEC that was originally calculated in exhibit A!

Consider a small change to the portfolio such as adding account AAA. The events between attachment and exhaustion will not change, nor will the first term. The change to the second term is precisely what is calculated in exhibit D as the CEC for account AAA. Thus, the CEC with this "interval" type of risk coefficient pattern can be interpreted as the variable contribution to the expected payout of a reinsurance layer defined by that pattern.

4. Applying the CEC method

4.1 Underwriting

CEC can be constructed to be a measure of the catastrophe risk taken on by the firm. An upper limit on that figure, defined by management, can be considered the firm's total cat underwriting capacity. The CEC method can be used to budget and manage that capacity.

At the beginning of each underwriting period, the total CEC limit is allocated to underwriting units. Along with capacity comes a total profit target (see next section about pricing).

During that period, as contracts are bound, they contribute to an ongoing growth in CEC and thereby consume that budgeted capacity in a linear fashion. For each deal being evaluated, the account expected profitability, marginal CEC, and cumulative CEC including the deal, can be calculated. The latter demonstrates whether there is "room" to write the deal and empowers underwriters to practice point-of-sale risk management while still allowing them flexibility to exercise judgment.

The entire process can be monitored via enhanced versions of standard cat modeling reports: usage to date versus allocated budget, and expected profit to date versus target. It is common for both insurance and reinsurance companies to develop a "reference portfolio" for their prospective future writings. By evaluating the held portfolio using the CEC method, and evaluating other risk appetite and risk aversion criteria, managers can inform the reference portfolio and develop a risk capacity budget. By analyzing the events contributing to the CEC, underwriting managers can also more specifically communicate perils or geographic areas of (dis)interest.

---

[7] We deliberately simplify by ignoring expenses, taxes, investment income, etc.

Unless the portfolio is growing very rapidly, it is recommended that some stability in the CEC events is utilized, such that the range of financial outcomes from the simulated events is defined on an annual basis and the CEC is managed for those same simulated events consistently throughout the year. This simplifies the reporting throughout the year and provides consistency in the firm's risk capacity messaging.

Once the CEC events have been identified, a (re)insurer who models each policy prior to binding, or receives catastrophe model input prior to binding, can evaluate the critical event cost for each prospective policy – and indeed could automate that process. For insurers who do not utilize a catastrophe model prior to each policy underwritten, the CEC events can be tracked on an annual, semi-annual, or quarterly basis as catastrophe model results become available. This allows identification of where risk capacity has exceeded plan, where there is additional capacity, and how the total risk capacity has been utilized by underwriter or underwriting unit.

To go one step further, it would be possible as well to identify scenarios under which particularly CEC-heavy risks would be acceptable and at what price, which is the subject of the next section.

4.2 Pricing and profits

Ultimately, market forces determine the price at which insurance contracts are executed. However, the starting point is a technical, actuarial evaluation of the risk. Typically, the lowest price that the seller will find acceptable (known as the reservation price) is based on a formula that includes the expected loss payment plus an extra margin to compensate the seller for risk.

Risk margins can be divided into two broad categories. They can be based on risk characteristics that are intrinsic to the contract, or they can be based on characteristics of the contract's risks that relate to a wider set of contracts. An example of the former would be the standard deviation of the losses. An example of the latter would be the CEC.

Earlier, the use of CEC as a budgeting mechanism was discussed. Accounts AAA and BBB were compared and found to have the same expected loss but significantly different contributions to the portfolio CEC. The underwriter was expected to prefer account AAA, "all else being equal". An explicit pricing rule that took CEC into account could provide a mechanism for adjudicating how much additional premium is required to make the two accounts appear equally attractive.

Looking at exhibit D, we see that both contracts incur $13,000 of expected loss, but account AAA has a CEC of $4,800 whereas account BBB's CEC is $7,500. In the previous section, we calculated that a 4-XS-3.5 reinsurance program had an expected payoff of $0.532mm. Say management has observed or calculated that the market price of such a program is $1.064mm, or two times the expected payoff. *Management feels that this is the market value of the capital guarantee it provides to policyholders and wants premiums to reflect it.* In addition, the actuaries feel it is reasonable to assume that the ratio of market price to expected payoff would remain stable over a reasonable range of expected payoffs.[8]

The following pricing rule accommodates this:[9] Premium = Expected Loss + 2xCEC. If a new contract were to bring in that much premium, it would cover its expected losses and also the market value of the additional risk that is borne by capital. Translating this into the exhibit D example, account AAA would require a premium of at least 13,000 + 2x4,800 = 22,600 whereas account BBB would require 13,000 + 2x7,500 = 28,000, or 24% more.

[8] More sophisticated models could be built by slicing the program into a sequence of narrower layers and pricing them individually, but that is unnecessarily complicated for this illustration. The interested reader is directed to Mango et al. (2013).

[9] Again, ignoring expenses, taxes, investment income,…

5. Critical event cost for annual aggregate losses

When dealing with perils such as hurricanes or earthquakes, where large portfolio losses in a year typically occur as a result of a single rare event, the OEP approach to the CEC method outlined above is both workable and adequate. However, with other perils such as severe storms, it is often the case that it is the sheer number of events during the year, not the single worst among them, that drives large annual portfolio losses. For those perils, a different implementation is needed.

Let us now assume that the probabilities in exhibit A are actually frequencies, and the fact that they sum to one is a coincidence of no significance. The "Cum Prob" column is to be ignored as meaningless now.

The process for creating simulated years of experience might proceed as follows: for each event, generate a random whole number, zero or greater, from a Poisson distribution with mean equal to the event's frequency. For frequencies like 0.01 or 0.04, these counts will mostly be zero, but occasionally they will be one, and (rarely) even greater than one. Collect all events with positive counts into one "year" of experience. The total of losses from all events (taking multiplicity into account) is the aggregate loss for the year. Repeat the process for as many simulated years as desired; typically this is many thousand. Sort the years by aggregate loss for a representation of the probability distribution of aggregate losses, also known as the AEP curve.

Exhibit G shows what a portion of the outcome might look like.

Exhibit G: Simulated Annual Experience

YEAR ID	EVT ID	Loss ($mm)	Description
1001	None	**0**	**TOTAL**
1002	102	3.0	Earthquake – CA
critical	101	2.0	Tornado/Hail – TX
	88	0.7	Winter Storm – Northeast
		5.7	**TOTAL**
1003	108	6.4	Earthquake – Midwest
	73	0.6	Earthquake – CA
	52	0.3	Windstorm – Northeast
		7.3	**TOTAL**
1004	103	3.2	Earthquake – Midwest
critical	101	2.0	Tornado/Hail – TX
		3.3	**TOTAL**

Source: Guy Carpenter & Company, LLC

In the context of annual aggregate losses, "critical event" now means critical *year*. Those years with total losses of a certain magnitude of interest are isolated and identified as critical (or, more generally, are assigned risk coefficients commensurate with their total losses).

The calculation of portfolio critical event cost proceeds as before. For example, say that critical years are those with losses between $3mm and $6mm. This means simulated years 1002 and 1004 are critical whereas 1001 and 1003 are not. No explicit probability has been assigned to each simulated year – they are all of equal probability, 1/N, where N (=4) is the total number of simulated years. The portfolio CEC based on these parameters is (5.7+3.3)/4 = 2.25.

We could also use alternative risk coefficients, as was done in exhibit F. In that case, the interpretation of familiar risk metrics such as TVaR and VaR would remain unaffected, except that they would be applied to the total year losses instead of worst loss in a year.

The calculation of critical event cost for portions of the portfolio or proposed additions (accounts) would seem to require a substantial amount of additional computation. Say we were interested in reworking the comparison between prospective accounts AAA and BBB that was done in exhibit D. It appears we would need to merge the two exhibits by (1) adding two columns to exhibit G for the two accounts' losses, (2) populating each exhibit G event row with the losses taken from the corresponding row of exhibit D, and finally (3) totaling each simulated year's losses from each account. Then the aggregate CEC calculation could proceed on the accounts' annual losses. This laborious process of rebuilding the annual simulation results would have to be done for *every* prospective account.

The result is shown in exhibit H. Notice that the loss experience in event 101 was duplicated in years 1002 and 1004. This can and will happen at random due to the simulation process of sampling events into years.

Exhibit H: Simulated Annual Experience of Prospective Accounts

YEAR ID	EVT ID	Description	AAA ($000)	BBB ($000)
1001	None	**TOTAL**	**0**	**0**
1002	102	Earthquake – CA	30	0
critical	101	Tornado/Hail – TX	20	0
	88	Winter Storm – Northeast	0	0
		TOTAL	**50**	**0**
1003	108	Earthquake – Midwest	19	68
	73	Earthquake – CA	0	0
	52	Windstorm – Northeast	0	0
		TOTAL	**19**	**68**
1004	103	Earthquake – Midwest	16	34
critical	101	Tornado/Hail – TX	20	0
		TOTAL	**36**	**34**

Source: Guy Carpenter & Company, LLC

The calculation of account-level aggregate CEC is as follows:

- Account AAA: (50+36)/4 = 21.5
- Account BBB: (0+34)/4 = 8.5

This is meant simply to illustrate the process and the results are not meant to be compared to the results in exhibit D. In a real application, there would likely be thousands of simulated years identified as critical.

In this example, we have taken the account loss data from the event table and propagated it forward into the annual simulation table, where it is combined with the critical year identification to compute the aggregate CEC. An alternative, computationally simpler method is to take the critical year identification from the annual simulation table and propagate it *backward* into the event table where it can be combined with the account loss data to compute the same aggregate CEC. This is illustrated in exhibit I.

Start with the event table (exhibit D).[10] First, all event probabilities are replaced by the annual probabilities (in this case, 1/4). Second, from the annual table (exhibit H), we count the number of times each event occurs in a critical year. (If more general risk coefficients were assigned to the years, we would sum them. Here we are summing zeroes and ones, i.e., counting.) This count is recorded in the appropriate event row in the AggCEC field.

Finally, for each account, compute the sum-product of probability, AggCEC, and the account loss:

- Account AAA: (0.25x2x20 + 0.25x1x30 + 0.25x1x16) = 21.5
- Account BBB: (0.25x2x0 + 0.25x1x0 + 0.25x1x34) = 8.5

These are, of course, the same results that were computed previously. A proof that this method produces the same results is presented in the Appendix.

The savings in computation is clear. Once a mechanism is in place to compute the CEC on an OEP basis, that same mechanism can be used to compute it on an AEP basis by first changing the probabilities and risk coefficients. The replacement columns only need to be computed once, based on the annual simulation and identification of critical years. All subsequent CEC calculations on portions of the portfolio or prospective accounts can be done with the new event table. Propagation of losses to the annual table is not necessary.

Exhibit I: Aggregate CEC of Two Prospective Accounts

EVT ID	Prob	Description	AggCEC	Acct AAA*	Acct BBB*
101	0.25	Tornado/Hail – TX	2	20	-
102	0.25	Earthquake – CA	1	30	-
103	0.25	Earthquake – Midwest	1	16	34
104	0.25	Winter Storm – Northeast	0	-	85
105	0.25	Earthquake – CA	0	52	-
106	0.25	Hurricane – Gulf	0	58	-
107	0.25	Hurricane – NC	0	60	32
108	0.25	Earthquake – Midwest	0	19	68
109	0.25	Hurricane – NY	0	-	76
110	0.25	Earthquake – CA	0	56	-
111	0.25	Hurricane – FL	0	85	-
112	0.25	Earthquake – WA	0	-	101

Source: Guy Carpenter & Company, LLC

[10] The first row, representing the aggregate of many events, has been deleted for clarity of exposition. In the example, none of those events contributes to the aggregate CEC.

6. The limits of certainty

If catastrophe model output was a perfect representation of reality, that is, if the scenarios were exhaustive of all that could happen and the associated probabilities and loss amounts were perfectly accurate, then the CEC method outlined above would generate results that could be relied upon with 100% confidence. Of course, neither is the case. Catastrophe models are, in a statistical sense, simply *estimates* of the underlying true reality of catastrophe risk. As such, they are subject to myriad sources of error. In particular, attempts to estimate expected losses for small geographic areas are subject to a greater degree of uncertainty than portfolio-wide estimates. See Major (2011) for a full discussion.

In this section, we outline methods for estimating the degree of uncertainty around (1) a particular CEC and (2) the difference between two CECs. The first is important in assessing the reliability of technical pricing formulae built on CEC. The second is important in assessing the reliability of CEC-based comparisons between two portfolio subunits. They are based on the one-sample and paired-sample t-tests of elementary statistics and apply only to the OEP version of the CEC. They cannot be applied to the AEP version without extensive modifications.

The methods proposed here are very simple and are meant primarily for illustration and achieving a sense of order-of-magnitude. More sophisticated approaches are beyond the scope of this paper.

6.1 Assessing one CEC calculation

In elementary statistics, one often has a limited number of observations (known as a *sample*) of a random quantity and wishes to know the long-run average of that quantity to be expected if more samples were taken. The average of the observed values, known as the *sample mean*, is often taken as the estimate of the long-run average, known as the *population mean*. The sample mean, it is hoped, is close to the population mean, but it is unlikely to be perfectly accurate.

In order to assess the accuracy of the sample mean, one computes the *standard error of the mean* by the following formula: StandardError = SquareRoot(Variance/NumberOfObservations). The *variance* is a measure of the dispersion of the observed values and is defined as the average of the squared differences between the observed values and their overall mean.[11]

The sample mean and variance of CEC are a bit more complicated, but still straightforward. The table of scenarios gives all the necessary information. The mean is the CEC itself as calculated here. To describe the calculation of the variance, we will illustrate in exhibit J using Account AAA from exhibit D.

The first step, already done, is the calculation of CEC. For account AAA it is 4.8.

The second step is to compute, for each event, the square of the difference between the event loss contribution to CEC and the overall CEC. By "event loss contribution" we mean the loss value that gets multiplied by the probability when computing the CEC. Recall that for the CEC in Exhibit D, only loss values for events 104 through 109 contributed to the CEC. All others contributed zero.[12]

There are no contributions from events 101-103 and 110-112. That is to say, all contributions are zero. The difference between zero and 4.8 is -4.8. The square of -4.8 is 23.04. Events 104 and 109 would contribute, theoretically, but the losses are zero so the squared differences are 23.04 again. Event 105 contributes 52; 52-4.8 is 47.2; 47.2 squared is 2227.84. Exhibit J fills in the rest.

---

[11] In computing this average, it is often better to use the number of observations minus one in the denominator rather than the number of observations. When there are a large number of observations, this refinement will not make a meaningful difference to the result.

[12] If we were using risk aversion factors as in Exhibit F, we would post the loss multiplied by the associated risk aversion factor. Here, the risk aversion factors are either zero or one.

After tabulating all the squared differences, we need to calculate their weighted average, using the probabilities as weights. For the aggregated "many events" row, we use the 71% representing the sum of all their probabilities. For the other rows, use the probabilities as posted. The average squared difference is then 0.71*23.04 + 0.03*23.04… which equals 225. This is the variance.

Exhibit J: Calculation of Standard Error, One Sample

EVT ID	Prob	Description	Acct AAA* contribution	Squared Difference
<many>	<many>	<aggregate of many>	0	23.04
101	0.03	Tornado/Hail – TX	0	23.04
102	0.03	Earthquake – CA	0	23.04
103	0.02	Earthquake – Midwest	0	23.04
104	**0.04**	**Winter Storm – Northeast**	**0**	23.04
105	**0.01**	**Earthquake – CA**	**52**	2227.84
106	**0.04**	**Hurricane – Gulf**	**58**	2830.24
107	**0.02**	**Hurricane – NC**	**60**	3047.04
108	**0.04**	**Earthquake – Midwest**	**19**	201.64
109	**0.01**	**Hurricane – NY**	**0**	23.04
110	0.02	Earthquake – CA	0	23.04
111	0.02	Hurricane – FL	0	23.04
112	0.01	Earthquake – WA	0	23.04
		Average (=CEC)	**4.8**	
		Average (=Variance)		**225**
		(Variance/400)^{½} (=Std Err)		**0.75**
		* Losses are $000		

Source: Guy Carpenter & Company, LLC

The next step requires some creativity. The standard error is the square root of this after we divide by the number of observations. But what is the number of observations? Following Major (2011), we might use the following rules of thumb:

- Hurricane models are based on 100 years of observations
- Earthquake models are based on 200 years of observations
- Severe storm models are based on 100 years observations.

Since all three perils are combined in exhibit D, we will consider that there are, in effect, 400 "observations" underlying the catastrophe models that produced this table.[13]

The standard error of Account AAA's CEC is then computed in two steps. The variance, 225, is divided by 400 for a result of 0.5625. The square root of that, 0.75, is then the standard error.

[13] This step is a bit controversial. If we were doing an analysis of severity as part of an aggregate CEC, then the number of actual events (not years) would be used, and it would be appropriate to add together the number of separate types of events. Since we are doing occurrence and not aggregate CEC, it could be argued that we should use only the 100 years of experience available. Ultimately, choosing this number is a judgment call.

How do we interpret this? The usual formal statistical procedure is to apply properties of the *Student's "t"* distribution to derive a confidence interval. With at least 100 observations, and given the creative license we have taken above, we can be a little less formal and use the following guidelines:

- We can be about 70% confident that the true mean (i.e. the CEC if we computed it with perfect information) is within one standard error of the calculated CEC, that is, between 4.8-0.75 = 4.05 and 4.8+0.75 = 5.55.
- We can be about 95% confident that the true mean is within two standard errors, that is, between 4.8-2*0.75 = 3.30 and 4.8+2*0.75 = 6.30.

6.2 Assessing the difference between two CEC calculations

A similar computation on Account BBB results in its CEC being 7.52 plus or minus a standard error of 1.11. What do we make of the difference, 7.52-4.8 = 2.72? How does the standard error concept apply?

Here we make use of the paired t-test. Each scenario gives us loss figures for both Account AAA and Account BBB. It is not as if we had two separate independent samples from the two accounts. Their loss figures are coordinated, synchronized… paired.

The calculations shown in Exhibit K are almost the same as in Exhibit J except that the column representing the loss to Account AAA is replaced by the difference between the two account losses. For example, in event 107, Account AAA showed a loss of 60 and Account BBB showed a loss of 32. The difference, 60-32 = 28, is posted for event 107 in Exhibit K.

Exhibit K: Calculation of Standard Error, Paired Samples

EVT ID	Prob	Description	Difference in Account contributions	Squared Difference
<many>	<many>	<aggregate of many>	0	7.40
101	0.03	Tornado/Hail – TX	0	7.40
102	0.03	Earthquake – CA	0	7.40
103	0.02	Earthquake – Midwest	0	7.40
104	**0.04**	**Winter Storm – Northeast**	**-85**	6770.00
105	**0.01**	**Earthquake – CA**	**52**	2994.28
106	**0.04**	**Hurricane – Gulf**	**58**	3686.92
107	**0.02**	**Hurricane – NC**	**28**	943.72
108	**0.04**	**Earthquake – Midwest**	**-49**	2141.84
109	**0.01**	**Hurricane – NY**	**-76**	5369.96
110	0.02	Earthquake – CA	0	7.40
111	0.02	Hurricane – FL	0	7.40
112	0.01	Earthquake – WA	0	7.40
		Average (=CEC)	**2.72**	
		Average (=Variance)		**612.68**
		(Variance/400)^{½} (=Std Err)		**1.24**
		* Losses are $000		

Source: Guy Carpenter & Company, LLC

The result is that the difference in CEC between the two accounts, 2.72, has an associated standard error of 1.24. We can be 95% confident that the true difference between them is between 2.72-2*1.24 = 0.24 and 2.72+2*1.24 = 5.20. Note the lower end, 0.24, while not there, is getting uncomfortably close to zero. If the upper and lower limits bracketed zero, we would interpret this as meaning we could not be confident (at the 95% level) about which account actually had the higher CEC. In this case, we do have confidence that Account BBB has the higher CEC.

7. Conclusion

We have seen that the critical event cost provides for a measure of risk capacity that allows for easy subdivision into peril types, geographies, lines of business, or other partitions of a portfolio of insured risks – even down to individual accounts. It leads naturally to a capacity-budgeting and capacity-consumption management mechanism, and can be incorporated into technical pricing to guide the tradeoff between risk and profitability.

CEC in its elementary binary form is closely related to the expected payoff of a reinsurance layer, and the CEC evaluated on a small change maps directly to the associated change in that expected payoff. Generalizations of CEC that use variable risk coefficients can be seen to include common risk measures such as VaR and TVaR, as well as customized "spectral" or "utility function" types of risk measures. Moreover, these generalizations also possess the budgeting/apportioning and change-analysis properties of the binary CEC.

CEC can be computed on an occurrence (OEP) or aggregate (AEP) basis, given the appropriate cat modeling inputs. Once an OEP version is at hand, the AEP version can be implemented by "reloading" the event table from a one-time additional computational step.

Assessing the reliability of a CEC figure by computing its associated standard error is an important step. This informs and guides the use of CEC in business decisions by reflecting how much confidence or credibility can be assigned to the number.

A wide range of carriers – stock and mutual, from large international firms to small regionals and specialty firms – are already using risk coefficient-based metrics to express their risk appetite and risk tolerance. For them, the CEC approach applies naturally. A smaller but growing number are already applying the CEC method to manage their underwriting. In particular, the CEC method has broad appeal to firms that are already modeling individual policies with cat model tools such as RiskLink, RiskBrowser, CatStation, etc. Those tools allow for calculation of the average loss; it is but a short step to calculating the critical event cost. The method also appeals to national carriers with targeted sales plans and a need to develop risk capacity strategies across business units.

Disclaimer

References

Kreps, Rodney E., John A. Major, and Gary G. Venter (2006) "Marginal Decomposition of Risk Measures," *ASTIN Bulletin*, Brussels, Belgium: International Actuarial Association, Vol. 36, no. 2., 375-413.

Major, John A. (2004) "Gradients of Risk Measures: Theory and Application to Catastrophe Risk Management and Reinsurance Pricing," *CAS Forum*, Arlington, VA: Casualty Actuarial Society, Winter, 45-89.

Major, John A. (2011) "Managing Catastrophe Model Uncertainty, Issues and Challenges," white paper, Guy Carpenter and Company, LLC.

Mango, Donald F., John A. Major, Avraham Adler, and Claude Bunick (2013) "Capital Tranching: A RAROC Approach to Assessing Reinsurance Cost Effectiveness," *Variance*, Arlington, VA: Casualty Actuarial Society, forthcoming.

Appendix: Aggregate Critical Event Cost Implementation[14]

Introduction

Let $e = 1,\ldots,E$ index over events. Let $y = 1,\ldots,Y$ index over simulated years. If event e occurs in year y, set $p_{e,y}$ to one, zero otherwise. The occurrence probability of event e (from the event table) is f_e.

Let $s = 1,\ldots,N_s$ index over segments of the portfolio. Event e produces loss $L_{e,s}$ in segment s. Total loss from event e is given by $L_e = \sum_s L_{e,s}$.

The risk coefficients defining critical events are represented by KE_e; for critical years by KY_y.

Critical event cost calculated based on occurrences (OEP)

The critical event cost, based on the occurrence of events, is $C^1 = \sum_e f_e \cdot KE_e \cdot L_e$. The event-based CEC allocations to specific segments are given by $C_s^1 = \sum_e f_e \cdot KE_e \cdot L_{e,s}$. Clearly, the allocations add up to the total.

Critical event cost calculated based on aggregates (AEP)

The portfolio loss experienced in simulated year y is given by $M_y = \sum_e p_{e,y} \cdot L_e$ and the loss to segment s is given by $M_{y,s} = \sum_e p_{e,y} \cdot L_{e,s}$. The aggregate critical event cost is $C^2 = \frac{1}{Y} \cdot \sum_y KY_y \cdot M_y$ and for each segment s it is $C_s^2 = \frac{1}{Y} \cdot \sum_y KY_y \cdot M_{y,s}$. This calculation can be characterized as propagating the event losses $L_{e,s}$ forward to the year totals $M_{y,s}$.

We now rewrite the formula for C^2_s; the results for the total C^2 follow in obvious parallel.

[14] The authors are indebted to Chengyou Xiao not only for the proof, but for the very idea that this could be done.

$$C_s^2 = \frac{1}{Y} \cdot \sum_y KY_y \cdot M_{y,s} = \frac{1}{Y} \cdot \sum_y \left\{ KY_y \cdot \sum_e p_{e,y} \cdot L_{e,s} \right\} = \frac{1}{Y} \cdot \sum_y \sum_e KY_y \cdot p_{e,y} \cdot L_{e,s}$$

$$= \sum_e \frac{1}{Y} \cdot \left\{ \sum_y p_{e,y} \cdot KY_y \right\} \cdot L_{e,s} = \sum_e \frac{1}{Y} \cdot T_e \cdot L_{e,s}$$

The term in braces in the next-to-last expression we have defined as T_e in the last expression. This is the AggCEC column that appeared in exhibit I. Notice how the last expression has the same form as the expression for $C^1{}_s$ but with $1/Y$ replacing f_e and T_e replacing KE_e. Here, the annual risk coefficients KY_y have been propagated backward to equivalent event coefficients T_e.

This concludes the proof that forward and backward propagation methods are equivalent.