Article from:

Risks & Rewards

August 2013 – Issue 62

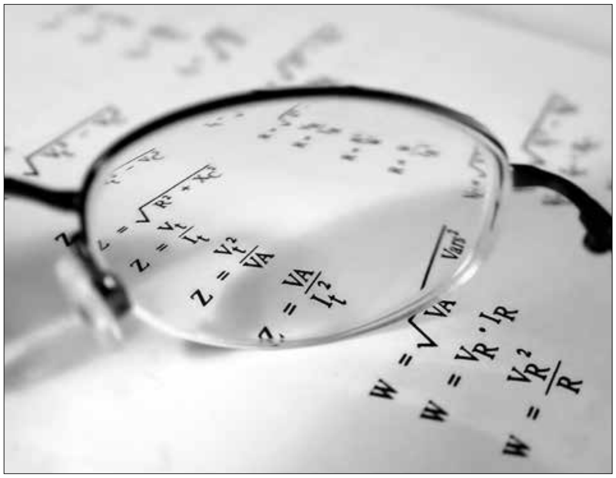

AUMANN-SHAPLEY VALUES: A TECHNIQUE FOR BETTER ATTRIBUTIONS

By Joshua Boehme

"HAPPY IS THE ONE WHO KNOWS THE CAUSES OF THINGS." –VIRGIL[1]

Actuaries encounter attribution problems on a regular basis. Indeed, any situation where results change, whether due to changes in assumptions, market conditions, or even just the passage of time, often leads to the natural follow-up question: why did the results change? To answer this question, actuaries use various techniques, each with its own strengths and weaknesses. However, a technique not widely known among actuaries—the Aumann-Shapley value from game theory—in many cases can produce attributions that better satisfy our intuitive expectations of a good attribution.

Some authors have already applied the Aumann-Shapley approach to various financial problems in non-actuarial settings. Denault (1999), for example, used it to allocate margin requirements among portfolios of options. However, the Aumann-Shapley approach has yet to receive widespread exposure within the actuarial community.

FORMALIZING THE PROBLEM

Suppose we have a multivariate function f and two vectors of parameters u and v representing the previous and latest parameters respectively. In the most general form of the problem we place almost no restrictions on the function f or its parameters. In some applications f may contain discontinuities or may lack a closed-form solution. The function f could take non-continuous parameters as well. For example, a binary variable could indicate whether to use one method or another, such as curtate versus continuous mortality.

In an attribution problem we seek to explain the difference f(v)-f(u) by assigning to the i^{th} variable an amount a_i representing its contribution to the difference, where i ranges from 1 to the number of inputs to the function f.

Ideally the total of the a_i values would equal f(v)-f(u), but in practice that does not always happen. Any remaining difference, which sometimes goes by the term untraced or unexplained, represents some portion of the change that the attribution method in question could not allocate to one of the input variables.

Although likely few of us have ever tried to formally list the properties we want a "good" attribution to satisfy, intuitively we have an idea of how a reasonable method should behave. For example, if the i^{th} variable did not change (so $u_i = v_i$), we would expect its contribution to the difference to equal zero. Similarly, if the i^{th} variable has no impact on the value of f (meaning that f(u) = f(v) whenever $u_i \neq v_i$ and $u_j = v_j$ for all $j \neq i$), then we again expect its contribution to equal zero.

ATTRIBUTION TECHNIQUES

Aumann-Shapley

The technique that this article focuses on, the Aumann-Shapley value, requires f and its parameters to satisfy a few conditions. Specifically, f must have partial derivatives[2] in all of its parameters along the vector between u and v. We do not need a closed-form version of f, but we must have a way to compute its partial derivatives at any given point on the path.

For attribution problems that satisfy these requirements, the Aumann-Shapley approach produces some valuable results. It always produces an attribution with no unexplained amount.[3] As we will see in later examples, its results also show a certain desirable stability with respect to how we set up the problem.

For each variable we calculate the attributed amount a_i as:

$$a_i = (v_i - u_i)\int_0^1 \frac{\partial f}{\partial x_i}\big((1-z)u + zv\big)dz$$

The resulting integral does not always have a closed-form solution, but we can evaluate it numerically.

Alternatively, we can view the Aumann-Shapley approach as a three-step process:

1. Find the partial derivative of f with respect to its i^{th} parameter, denoted $\frac{\partial f}{\partial x_i}$ here.

2. Integrate that partial derivative along the line segment between u and v. Here, the dummy variable z represents the linear interpolation between u (at z = 0) and v (at z = 1).

3. Multiply the result by the change in that parameter $(v_i - u_i)$

Step-Through

Many actuaries faced with an attribution problem will solve it by stepping through the parameters one at a time, a technique with several important advantages. As long as we can evaluate the function f at each combination along the step-through, f can have any number of discontinuities and it can lack a closed-form solution. We can use a step-through even when f has non-continuous inputs. Furthermore, a step-through also produces an attribution with no unexplained amount.

Step-throughs have one well-known disadvantage, though: the results depend on the arbitrary order we use to step through the parameters. In the examples in this article we will use a modified technique to overcome this issue: we will perform the step-through for every possible order, then average the attributions together.[4] This removes the dependency on an arbitrarily chosen order. As we will see with a later example, though, even this modified step-through method still has a significant weakness. Despite that, in situations where we cannot satisfy the requirements of the Aumann-Shapley approach, a step-through remains a viable alternative.

Partial Derivatives

Actuaries already frequently use derivatives or approximations to the derivative to perform attributions. Some partial derivatives come up so often that they have specific names, such as the "Greeks" (delta, gamma, vega, rho, theta, etc.) or duration. In some cases, we use formulas to directly calculate the partial derivatives; other times, we shock one of the parameters a small amount to numerically estimate the derivative.

Partial derivatives have as one major advantage their frequent ease of computation and interpretation. For example, from the duration of a bond, a relatively intuitive concept, we can quickly estimate the change in its value due to a change in interest rates.

The main difference between a partial derivative attribution and the Aumann-Shapley approach comes from *where* we evaluate the partial derivative. In the Aumann-Shapley approach, we evaluate it along the entire path between u and v. For the partial derivative, we evaluate it at a single point, usually the beginning point.[5] This difference, though, leads to the major drawback of a partial derivative approach: the attribution generally has a nonzero unexplained amount. This may suffice for a quick estimate. Other times, though, we may want a complete attribution of the difference.

EXAMPLES

Example 1: *Zero-Coupon Bond*

Consider a zero-coupon 10-year bond with a maturity value of $1 million. For a given yield to maturity y, the following formula gives its value at time t:

$$f(y,t) = e^{-y(10-t)} \cdot 10^6$$

Suppose we have the following parameter sets:

	y	t	f(y, t)
u	5.00%	1	637,628
v	8.00%	2	527,292
		Difference	-110,336

We then get the following attributions from the three methods discussed above:

	y	t	Total Attributed	Unexplained
Aumann-Shapley	-147,619	37,284	-110,336	0
Step-through	-146,952	36,616	-110,336	0
Partial derivative	-172,160	31,881	-140,278	29,942

CONTINUED ON PAGE 30

As expected, both the step-through and the Aumann-Shapley approach fully attribute the change. They also produce comparable results. The partial derivative results, though easily calculable,[6] do not accurately capture the total change in value.

Example 2: *Zero-Coupon Bond, Revisited*

Many times, we can formulate a problem in multiple ways. Suppose that instead of expressing the yield for the zero-coupon bond in terms of a single variable, as in the previous example, we express it in terms of two components: a prevailing interest rate r, and a credit spread c. As a formula:

$$f(r,c,t) = e^{-(r+c)(10-t)} \cdot 10^6$$

Given the modified parameter sets:

	r	c	t	f(r, c, t)
u	4.00%	1.00%	1	637,628
v	5.00%	3.00%	2	527,292
			Difference	-110,336

The initial and final values have not changed from the previous example—we have merely separated the yields to maturity into two components. Now that we have shifted perspective, what happens to the attributions?

	r	c	t	Total Attributed	Unexplained
Aumann-Shapley	-49,206	-98,413	37,284	-110,336	0
Step-through	-49,162	-98,027	36,853	-110,336	0
Partial derivative	-57,387	-114,773	31,881	-140,278	29,942

Note in particular the step-through values. By formulating the problem in a slightly different way, the value attributed to the time variable t has changed! In contrast, the amounts attributed to t by both the Aumann-Shapley approach and the partial derivative have not changed from before.

Now, some readers may object that this change in the attributed value for t comes from the fact that we stepped through every possible order of variables. In practice, most actuaries would use only a single order, and most likely we would step through the two interest rate components consecutively. Under those circumstances, the value attributed to t would come out equal under both formulations. For example, if we use the order y, t in the first example and c, r, t in the second, then in both cases we attribute $40,540 to t. However, any particular order comes from an arbitrary choice on our part. Nothing intrinsic in the order itself would lead us to conclude that we should choose one order over another. The equally natural order t, y (or t, c, r) would lead us to attribute $32,692 to t.

In the end, when using a step-through, we must either accept that we have chosen an arbitrary order, or we must accept that the results could vary if we re-formulate the problem in an equivalent way. Either way, step-throughs produce non-unique results.

Example 3: *Binary Call*

Suppose we own a binary call on a particular security, with a strike K at 100. For illustrative purposes we will hold the interest rate r constant at 2 percent and the volatility σ constant at 25 percent, and we will assume the underlying security pays no dividends. For the current asset spot price S and time to option maturity t, the value of our call equals:

$$f(S,t) = e^{-rt}\Phi(d_2)$$

where

$$\Phi(x) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{x} e^{-\frac{1}{2}u^2} du$$

and

$$d_2 = \frac{\ln\frac{S}{K} + (r - \frac{\sigma^2}{2})t}{\sigma\sqrt{t}}$$

At the boundary where t = 0, its value equals

$$f(S,0) = \begin{cases} 1 \text{ if } S > K \\ 0 \text{ if } S \le K \end{cases}$$

Given the parameters:

	S	t	f(S, t)
u	90	1	0.314
v	110	0	1.000
Difference			0.686

The three methods disagree significantly about the nature and magnitude of the time component's contribution:

	S	t	Total Attributed	Unexplained
Aumann-Shapley	0.435	0.251	0.686	0
Step-through	0.653	0.033	0.686	0
Partial derivative	0.312	-0.060	0.252	0.434

The price crosses the strike during the attribution period, but the call does not reach full value immediately at that time. Its value still includes a discount for the probability of a subsequent decrease. The passage of time eventually drives that probability to zero, bringing the option to its full value. At the beginning of the attribution period, though, the opposite pattern holds: the possibility that volatility will cause the asset price to exceed the strike recedes as we approach maturity, meaning that the passage of time reduces the option's value. The partial derivative results reflect the latter effect.

Thus, f's sensitivity to time varies considerably over the attribution region. By only considering the edges of the region, the step-through does not accurately capture the full sensitivity and ends up attributing little of the change to the time component.

AN INTUITIVE ARGUMENT FOR WHY AUMANN-SHAPLEY PRODUCES A COMPLETE ATTRIBUTION

Although the examples have shown that the Aumann-Shapley approach produces a complete attribution in those cases, they do not explain why it works in general. A quick (though non-rigorous) argument will help illustrate the logic underpinning the technique.

We have defined f as a multivariate function of the x_i variables. However, over the attribution region we can also view f as a function of z alone, denoted $f^{(z)}$ for clarity. The linear interpolation between u and v connects the two: $f^{(z)}(z) = f((1-z)u+zv)$. Thus $f^{(z)}(0)$, for example, would mean to calculate the values of each x_i for z=0, then evaluate f at those values—in other words, $f^{(z)}(0)$ equals our initial value f(u). We can now write out an equation for the difference we seek to attribute:

$$f(v) - f(u) = f^{(z)}(1) - f^{(z)}(0)$$

Assuming that we have a sufficiently smooth function, we can express the difference as an integral:

$$= \int_0^1 \frac{df^{(z)}}{dz}(z)dz$$

We still need to relate this back to the original x_i variables, and we do this by applying the chain rule. Note that we have switched back to the original multivariate function f:

$$= \int_0^1 \left[\sum_i \frac{\partial f}{\partial x_i} \cdot \frac{dx_i}{dz}((1-z)u+zv) \right] dz$$

Since z represents our linear interpolation variable, the derivative of x_i equals the difference between the final and initial values. With that substitution, and separating out the individual terms inside the integral, we finally obtain:

$$f(v) - f(u) = \sum_i (v_i - u_i) \int_0^1 \frac{\partial f}{\partial x_i}((1-z)u+zv)dz$$

where the term on the right hand side corresponding to each x_i gives that variable's attribution.

From this argument, we can also see why jump discontinuities cause this approach to fail, since they cause changes in the value of f that do not get captured by the derivative. However, as long as we have a sufficiently smooth function (and a sufficiently accurate calculator for the integral) we will always get a complete attribution from the Aumann-Shapley approach.

CONTINUED ON **PAGE 32**

ACCURACY OF NUMERIC INTEGRATION

Since the Aumann-Shapley results come from a numeric integration, a natural question arises: how much confidence should we have in their accuracy? In practice we can obtain very rapid convergence using Simpson's rule.[7] Example 1's results earlier used 1,000 points to ensure a highly accurate result, but even if we evaluate the integrals at just *three* points and use Simpson's rule we get almost identical results:

						Contributions to integrals	
z	y	t	f	$\frac{\partial f}{\partial y} = -(10-t)f$	$\frac{\partial f}{\partial t} = yf$	$\int \frac{\partial f}{\partial y}$	$\int \frac{\partial f}{\partial t}$
0.00	5.00%	1.00	637,628	-5,738,653	31,881	-956,442	5,314
0.50	6.50%	1.50	575,509	-4,891,829	37,408	-3,261,219	24,939
1.00	8.00%	2.00	527,292	-4,218,339	42,183	-703,057	7,031
					Total	-4,920,718	37,283
					(v_i-u_i)	3.00%	1.00
					(v_i-u_i) Total	-147,622	37,283
					Previous Results for Comparison	-147,619	37,284

Thus, even a quick calculation can produce reasonable results. Furthermore, by using a spreadsheet or programming language we can easily evaluate more points to improve the accuracy.

ADDITIONAL CONSIDERATIONS

In this article we have looked at some artificially simple examples. In real life, though, we may face more complex attribution problems. For example, we may wish to reflect the fact that interest rates vary based on the time to maturity. Even if the yield curve does not shift, the yield on an asset could change simply due to the passage of time. In that case, it may make more sense to assign the change in value solely to time, not due to a nonexistent movement in the yield curve.

To further complicate matters, when valuing options we could view volatility as depending on both the time to expiration and the moneyness of the option, leading to a two-dimensional volatility surface.

The Aumann-Shapley approach applies in those more complex cases as well. We do need a way to interpolate smoothly between the observed points. However, to perform any attribution we generally need some form of interpolation anyway since our valuation points will rarely correspond exactly to the market-observable points. Some common methods include linear interpolation and cubic splines, both of which provide differentiable interpolations.

Although this article has focused on asset valuation examples, we can use the Aumann-Shapley approach for other applications. However, we do need to verify that our function f meets the requirements. Liabilities in particular, though, often contain features that can potentially create discontinuities, including:

1. Charges, guaranteed rates of return, or other features based on market values rounded to the nearest percent, nearest 25 basis points, or some other multiple.
2. Any feature involving rebalancing something back to a target, but only if outside some tolerance band. For example, due to an investment strategy we have adopted or due to a contractual agreement we have entered into, we might rebalance a particular asset allocation back to 80 percent equity at end of each quarter if the current asset allocation deviates from that by more than 5 percent.
3. Franchise deductibles.
4. Digital payoffs.
5. Features activated or deactivated at certain thresholds, including knock-in and knock-out features.

Keep in mind, though features such as these do not *automatically* imply a problem. In example 3 our function contained a discontinuity at the point (S, t) = (100, 0), but the path between u and v did not pass through that point. In general, the specific circumstances of the problem will dictate whether we can use the Aumann-Shapley approach or not.

Despite its requirements, the Aumann-Shapley approach offers a powerful way to solve attribution problems. By adding it to their toolkit, actuaries can produce more reliable and more complete attributions, and thus move that much closer to truly understanding the causes of things.

Joshua Boehme, FSA, MAAA, works in the Asset-Liability Management department at Jackson National Life.

REFERENCES

Denault, Michel. (1999). Coherent Allocation of Risk Capital. Journal of Risk, 4, 1—34.

END NOTES

[1] Virgil. Wikiquote. Retrieved Jan. 27, 2013, from http://en.wikiquote.org/wiki/Virgil

[2] The existence of partial derivatives also implies the continuity of f. The partial derivatives can contain non-removable discontinuities, but f itself cannot contain any along the path.

[3] Except to the extent that whatever tool we use to calculate the results has finite precision, though this issue applies to any attribution technique.

[4] This corresponds to the Shapley value from game theory.

[5] Evaluating it at the end point instead or at both points does not eliminate its drawbacks. For the examples in this article we will use the partial derivative at the beginning point.

[6] The attribution for t, for example, equals 5 percent of the initial value.

[7] $\int_{x-\Delta}^{x+\Delta} f(z)dz \approx 2\Delta \frac{f(x+\Delta)+4f(x)+f(x-\Delta)}{6}$