1

Risk-Based Pricing - An Introduction

> *"The biggest risk is not taking any risk... In a world that changing really quickly, the only strategy that is guaranteed to fail is not taking risks."*
>
> – Mark Zuckerberg

Risk-free investing simply does not exist - or if it did, it would be so overwhelmed by the inefficiencies of the investment process that the capital base of any who pursue it would quickly erode to nothingness like a sandcastle in the waves.

Taking risk is an essential part of any investment strategy for without risk, there truly is no reward. And, by investing we simply mean 'putting to work' capital that must be efficiently deployed. The burdens and responsibility of those of us who work in the financial markets encompass all efforts to ensure that investments are prudent yet profitable [1]. In the context of risk, this means that we must have a clear foundation for adjusting our pricing - the purchase price that we exchange for assets - to reflect the risks that are being taken [2].

Many investors shy away from certain investments: they simply cannot or will not entertain the slightest exposure to *verboten* assets, regardless of the potential return. For some, this might be religious; Shariah investors, for example, would not invest in pork products regardless of the potential returns [3]. Others, for moral reasons may well shy away from what they call "sin-based" products such as gambling or alcohol. Some have investment guidelines that clearly circumscribe the type of asset that is suitable for investment. But all investors would seem to have some flexibility in the assets that they choose - if not, they would quickly be replaced by a machine.

[1] e.g. Haugen, 1990 [10, chap. 5]
[2] Fong & Fabozzi, 1985 [7]
[3] see Iqbal, 2007 [14, pp. 16-22], and Ayub, 2007 [1, chap. 3]

Within the guidelines that constrain investment activity, the thoughtful investor will generally be pursuing all assets that qualify - but will be doing so with an eye to pricing in the risks associated with each specific asset [4]. We cannot turn a blind eye to risks that clearly exist and simply wish them out of existence. This was part of the meltdown of the financial markets during the downturn of 2007-2009. There were risks associated with certain assets that most investors simply refused to consider. They ignored certain risks inherent in their favorite asset categories - and despite those risks paid much more for assets than they turned out to be "worth," once push came to shove. The meltdown in asset prices that followed those initial few precarious months was not the sole cause of their loss - for much of the loss was already assured because the assets were priced incorrectly [5]. The risks were inherent in the assets from the beginning. It was simply that when the events unfolded, those who had not priced their portfolios properly suffered the consequences of their ignorance of those risks.

This series will attempt to address risks that are inherent in fixed-rate assets and it will provide a framework for properly valuing assets given these risks. It will rely upon the development of two primary lines of development. The first is a development of an understanding of risk. The other is the development of a methodology for creating cash-flow models for these assets - the primary means by which we can place a measurable and repeatable value on the assets and compare them one to another.

1.1 Risks

This particular volume (Volume 1 of the Series) begins our modeling of risk by first imagining a world in which no risk *per se* exists - at least so far as the cash-flow models of our assets are concerned. The layerings of risk will be taken up, beginning with the next volume in this series. This volume presumes a world without variability in payments, without options, without the inherent randomness of human events. We do this, at this early stage in cash-flow modeling, to clearly understand what we call the *pristine* cash flows associated with an asset. These are the cash flows that would result from a debt instrument if everything goes precisely as idealized in the basic terms of the financing contract - that all payments occur on time, that no observable events exist with regard to borrower credit, that no options - such as early repayment options - exist or are likely to be exercised. This variability-free analysis is the starting point for any cash-flow model of an asset [6]. As we must begin our risk discussions from a fixed point upon which all agree, so, too, must we must build a foundation for cash-flow modeling that presumes, initially, that there is a theoretical state of no risk, a state where all contractual obligations will be fulfilled. Only from

[4] see Smithson, 2003 [18, chap. 2]

[5] refer to Hubbard, 2009 [13, chap. 4] Also note that Hubbard refers to *catastrophic* overconfidence (p. 102), an apt term for market behavior of the time.

[6] see Fabozzi, 1997 [6] [5]

that vista can we then layer on models with probabilistic outcomes that give us the insight we need to truly understand and value an asset.

But, since we are *not* addressing the modeling of variability (i.e., risk) in this volume, let us clarify what that implies - let us identify what we are *not* addressing. We are presuming that cash flows are always received exactly as specified in the basic contract documents - with no options being exercised and no variation in human performance, no variation in mail delays and no uncertainty about credit performance; we are imagining a world that does not exist in reality. To the contrary, events in the real world cause variability in each of these things that we casually assume not to exist within this first volume. This natural variability is considered risk, risk in financial performance.

Risk, as we will come to understand it within the financial world, can be equated with variability of outcome [7]. Financial risk is considered equivalent to uncertainty about the timing and completeness of payments [8]. In this volume we are concerned with modeling idealized, risk-free assets - we then will progress in future volumes to discussing the incorporation of known risks into this risk-free, pristine model. The approach is systematic and methodical. It allows the incorporation of risks, even compound risks, into cash-flow models of assets.

1.1.1 Variability in Outcome Equates to Financial Risk

With respect to variability and risk, the common analogy involves comparing a vending machine with a slot-machine (or other gambling device). The vending machine provides a high degree of certainty about what the outcome will be - or what it should be - if the device is working properly. The slot-machine device provides a set of outcomes that is much more directed by a distribution of probabilities - it is much more variable in its behavior.

The accompanying Figure 1.1 illustrates two types of financial investments. One, represented by the dotted line, has greater variability in outcome. The other, represented by the solid line, has much less variability in outcome. In financial terms, the greater variability is directly translated into the dimension of risk. Financial risk is considered to be greater variability in outcome.

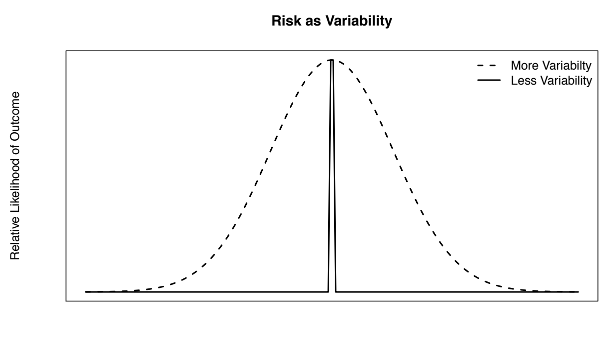

Figure 1.1: Illustration of Risk as Variability

Greater variability can be intuitively grasped with a few examples. Let's introduce a type of notation that describes a likelihood and an outcome. For

[7] e.g. Keown, 2005 [15, chap. 6]

[8] see Damodaran, 2008 [4, pp. 5-7]

example, let's imagine an investment which we will call x that has only two outcomes. One of those outcomes consists of an investment outcome returning \$1,000, and it does so with a probability of 99 chances out of a hundred (we'll write this as $p(x_1) = .99$), which allows us to represent the outcome with the notation $x_1 = \textit{\$1,000} \rightarrow p(x_1) = .99$) and the other outcome consists of returning \$999, and it does so with 1 chance out of a hundred ($x_2 = \textit{\$999} \rightarrow p(x_2) = .01$). Together, this would make the overall investment x able to be described as:

$$Investment(x) = \begin{cases} x_1 = \textit{\$1,000.00} & \rightarrow p(x_1) = .99 \\ x_2 = \textit{\$999.00} & \rightarrow p(x_2) = .01. \end{cases}$$

We can contrast the above investment x with another investment y that is:

$$Investment(y) = \begin{cases} y_1 = \textit{\$1,010.09} & \rightarrow p(y_1) = .99 \\ y_2 = \textit{\$0.00} & \rightarrow p(y_2) = .01. \end{cases}$$

Just by looking at these two investments, once can see that investment x has a very limited range of outcomes (just one dollar difference between the two alternatives), while investment y has a much greater variability - more than \$1,000 difference. We grasp intuitively that this investment y is *riskier* than investment x. The range of outcomes might be considered a simple indicator of risk.

However, there is a better way of examining this matter of risk and variability, a methodology that proves to be much more useful than just looking at the range of outcomes. To develop this methodology, though, we will need to turn to the area of statistics and probability theory.

The **Expected Return** (or average return) for any investment, z, with k discrete outcomes is defined as

$$E(z) = \sum_{i=1}^{k} z_i \times p(z_i), \tag{1.1}$$

where z_i is the value of outcome i and $p(z_i)$ is the probability of that outcome occurring. This means that both the example investments above, (x and y) have the same Expected Return $((\textit{\$1,000.00} \times .99) + (\textit{\$999.00} \times .01) = \textit{\$999.99}$ and $(\textit{\$1,01009} \times .99) + (\textit{\$0.00} \times .01) = \textit{\$999.99})$.

The thing that makes x and y, the investments from above, different is not their Expected Return, but another property of the outcomes. Above we noted that the range of outcomes for the two investments is different. Investment x will return either \$1,000 or \$999 - a difference of just \$1.00. Either of these two alternatives might be considered almost the same to some investors, a sure thing, so to speak. Investment y, on the other hand has a range of returns that is greater than \$1,000. The return could be \$1,010 or it could be \$0 - this is considered a *risky* investment by some.

The measure that we use to indicate this variability in return is called the **Variance** of the outcome distribution. For discrete investment outcomes, as in the investments x and y, above, the Variance is is measured as:

$$Variance(z) = \sum_i (z_i - \mu)^2 \times p(z_i). \tag{1.2}$$

This measure will indicate that the two investments are very different, indeed. In calculating the Variance, we use the Expected Return as our estimate of μ, which shows that the Variance of investment x is

$$\begin{aligned} Variance(x) &= ((\$1{,}000.00 - \$999.99)^2 \times .99) + ((\$999.00 - \$999.99)^2 \times .01) \\ &= .0099 \end{aligned}$$

and the Variance of investment y is

$$\begin{aligned} Variance(y) &= ((\$1{,}010.09 - \$999.99)^2 \times .99) + ((\$0.00 - \$999.99)^2 \times .01) \\ &= 10{,}100.81, \end{aligned}$$

two very different values. Just as we recognize intuitively by looking at the range of outcomes, investment y is shown to be much riskier than investment x using this Variance measure, as well; the Variance for this investment is several orders of magnitude greater than that for investment x - even though the Expected Return is the same. Variability in an investment return (i.e., the Variance of the investment outcome distribution) is the measure of risk that we utilize in this series.

1.2 Adjusting Pricing for Risk

As we have just seen, certain investments may have profiles of risk that make them difficult to compare with other assets. At the extremes some investors may be averse to any risk at all that exceeds government risk, while others may be interested in the marginal yield pick-up that can be obtained by judiciously taking certain chances with respect to credit or with respect to prepayment or other specific risk. This book, however, is not about the extremes in risk-taking. We do not address the merits of taking little risk or taking extreme risks. Instead, this book looks to address the issues associated with pricing similar investments, with comparative risk modeling. This book seeks to provide tools for comparing alternative investments that are typically quite similar to each other. To accomplish this use a risk-based pricing approach that provides a rationale or a framework for establishing relative values of the assets.

In our approach, we presume that the assets in question are within the scope of eligible assets for a diligent investor. We want to develop a methodology that allows the investor to compare assets and identify relative value. We do this within a framework that can establish a pricing point at which a willingness to buy or to sell assets will exist if the offer or bid price is more favorable than that pricing point [9]. If comparable but different assets are available at a single,

[9] see Hirshleifer, 1976 [11, pp. 164-166]

common price, then this methodology should identify those assets that exceed the average value and those that do not. Establishing relative value among common assets is the goal. And, to achieve this goal a thorough understanding of the asset is required, as is an understanding of the risks associated with that asset.

There are two very general approaches to risk-based pricing. One method adjusts pricing for risk purely as a matter of uncertainty. The other captures elements of risk by adjusting expectations. The uses and relative merits of these two approaches are addressed. And, as will be shown, many investors incorporate combinations of these two approaches when they clearly understand the risks in an asset and seek to price those risks appropriately.

1.2.1 Discount Rates Based Purely on Uncertainty

When some investors approach pricing, they look only at contractual payments associated with an asset and then at the likelihood that they will not receive all of those payments on time and in full. They use a different discount rate for different risks or uncertainties.

While this may be an appropriate way of incorporating risk into a valuation exercise in some circumstances - it is easy for this approach to become a very blunt instrument - simply using a greater discount rate for greater uncertainty. It is easy for this approach to lead to oversimplifications of the valuation process.

The accompanying Figure 1.2 illustrates the effect of discounting a set of cash flows at various discount rates. The topmost plot illustrates the hypothetical cash flows of a loan instrument through time (the grey bars) discounted at a relatively low rate (at 5% *per annum*). The effect of the discounting through time is represented by the black bars next to the cash flows, which represents the discounted value of that cash flow discounted back to time zero. In the topmost plot, the effect of the time-value discount is relatively small - whereas in the bottom-most plot the time-value discount is quite severe (at 25% *per annum*). The result is seen in the relatively small values of the later cash flows in the bottom plot.

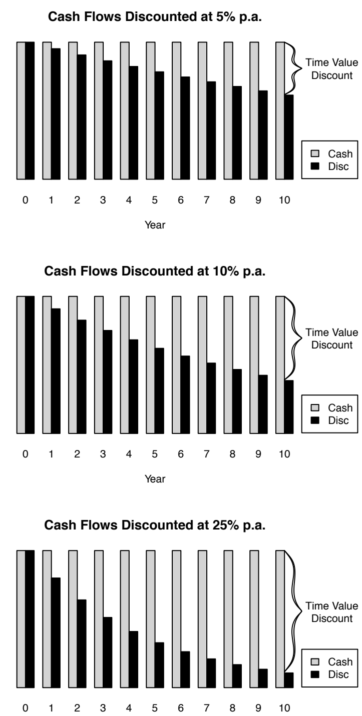

Figure 1.2: Effect of Discount Rates

In selecting a higher discount rate for valuing cash flows with greater uncertainty, the values placed on later cash flows are increasingly diminished. If the asset is seen as *risky*, then the valuer places less and less value on cash flows that are further removed from the present.

When the aggregate of these discounted cash flows is taken - that is, when the Net Present Value (or NPV) of all these cash flows are summed:

$$NPV = \sum_{p=0}^{term} \frac{Net\ Cash\ Flow_p}{\left(1.0 + \frac{Annualized\ Discount\ Rate}{Num\ Periods\ in\ Year}\right)^p}, \tag{1.3}$$

define terms we can intuitively see that the greater the discount rate, the lower the aggregate NPV for any given set of cash flows. This NPV is the price that we are willing to pay to receive these cash flows; it is the value that we place on the cash flows. (This topic is revisited in greater detail in Section 7.1 on page 90 of this volume.)

Thus, in using this blunt instrument of simply increasing the discount rate for greater uncertainty, we are essentially taking a greater and greater discount - or we are reducing the price we are willing to pay - for assets with greater uncertainty. This has the effect of devaluing riskier assets; the greater the risk in a given set of cash flows, the less we are willing to pay. This is a methodology that is often the best method that can be used for valuation, yet sometimes it is misused.

For example, some investors may look at one collection (or pool) of automobile loans and postulate a default rate, say 3% *per annum* in defaults. For this assumed default rate, they might seek to apply a discount rate on the contractual cash flows from the pool that is *greater* than if they assumed a 1% annual default rate. Or, if they assumed a 10% annual default rate, they might seek to apply an even higher discount rate. Here they equate the concept of higher expected default rate and greater risk or greater uncertainty. They believe that the greater the expected default rate, the greater the risk in the portfolio - and consequently apply a larger discount rate on the contractual cash flows from the pool with the greater expected default rate.

While this produces the desirable effect of lowering the valuation for a pool of loans the greater the expected default rate, it really confuses the point of uncertainty being related to discount rate. I could be relatively *certain* that the default rate was going to be 3% per year - perhaps because of my long experience with a certain type of borrower. If I were certain that the cash flows were going to be reduced by a particular rate of default - then by using a greater discount rate, I would not really be incorporating uncertainty (or risk), I would just be *fudging* the discount rate to generally give me a lower price. Which leads us to the next section in this chapter.

1.2.2 Adjusting Expectation to Statistical Norms

If an investor truly expects the cash flows from a pool of assets (the auto loans we envisaged in the last section) to be different from the contractual terms - then an effective risk-based pricing methodology will utilize the *expected* cash flows rather than the contractual cash flows in a pricing model. In other words, contractual cash flows are important to know, but they form only the foundation from which a more realistic, expectation-based pricing model can be developed.

The approach that this author takes in developing risk-based methodologies pricing is based upon adjusting expectations as a result of analyzing and understanding the underlying risks and factors that affect cash flows.

As shown in the nearby Figure 1.3, the *expected* cash flow may change (for example, because of an expected annual default rate). Note that the horizontal, dashed line corresponds in this figure to the contractual cash flow. Because of the expectation of defaults, the *expected* cash flows (the grey bars) are less than the contractual flows. Then, because of the discounting for time value and uncertainty, in this case at 5% *per annum*, the discounted value of the cash flows (the black bars) are lower, still. As can be seen, there are three scenarios presented; a case with 0% expected annual defaults, a case with 1% expected annual defaults and a case (at the bottom) with 3% expected annual defaults.

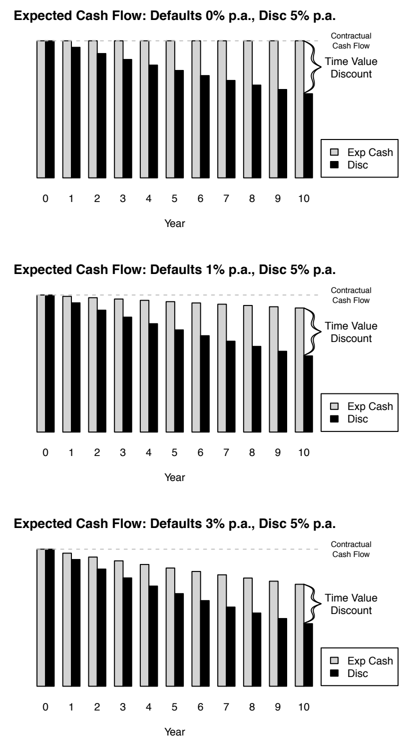

Figure 1.3: Expectation and Discount Rates

The NPV of these expected cash flows is, as indicated in Equation 1.3, above, the sum of the discounted values of the individual Net Cash Flows - but in this case, the *Net* cash flow is set to the *expected* cash flow (due to the assumed defaults, in this case). The annualized discount rate further reduces these expected cash flows to reflect the time value and uncertainty of the flows.

One might argue that the discounted values of the cash flows in this Figure 1.3 and those from Figure 1.2 (from page 14) are really not so dissimilar. One could imagine selecting a discount rate that produced nearly identical NPVs for the two approaches. That is one can imagine that selecting a greater discount rate and produce the same result as assuming a rate of default and then selecting an alternative discount rate. And, in this case, this is true. But, this simplistic approach breaks down as assumptions change.

In fact, the purpose of this writing is to illustrate that there are two mechanisms at play in risk-based pricing. The first is that there is a general relationship between risk (or variability in outcome) and the discount rate that should be used to price the cash flows from an asset. The second is that realistic models of expectation *must* be taken into account when pricing cash flows. A complete understanding of an asset and it's resulting cash flows is critical to valuing an asset. This implies that an accurate model can be built that reflects the effect of various scenarios on those cash flows. Those scenarios can result from changes in market conditions (e.g., interest rates) or changes in economic or credit conditions - or from many other factors that may exert an effect on the asset cash flows.

A critical view of these assumptions, the likelihood of various outcomes, the sensitivity of the asset to the various assumptions - these all play a role in our approach to risk-based pricing.

1.3 The Foundation for Risk-Based Pricing

The current work creates a systematic approach to risk-based pricing. That approach is based upon two main principles: risk and expectation in financial instruments. To any student of probability theory, these two points - at least as developed to this stage in our argument - will sound familiar. *Risk*, as we speak of it, is another way of describing variability in outcome. *Expectation*, as we define it, is that outcome that we really believe will happen - or that we believe to be the most likely outcome foreseen with regards to our asset cash flows.

The two concepts, perhaps ordered differently as *expected value* and *variability* are the first two moments of a statistical distribution. Statisticians usually refer to these moments as the *mean* and *variance* of a distribution of random variables [10]. Generally speaking, if we think of a random variable X as having some expected value, $E(X)$, then, for any order r we can define $E(X^r)$ as the r^{th} *moment* of a distribution. And, if we want to center these moments around the mean (which we can call μ) of the distribution, we have the r^{th} *central moment* of a distribution defined as $E((X-\mu)^r)$.

To make this more concrete, we can think of X as a discrete random variable (a random variable that only takes on specific values - such as dollars or pennies). Perhaps X is a payment that you might receive at some time in the future, valued in dollars. If we can identify the distribution of likely values, as in the accompanying Figure 1.4, we can develop this concept further. In this figure, the dashed line represents the relative likelihood of receiving a payment of amount X. The distribution is centered around the solid, vertical line at \$250 - or, in other words, the mean or *expected value* of the payment is \$250 ($E(X) = (250)$). Also, the dashed line represents some variability in the payment that is likely to be received. We

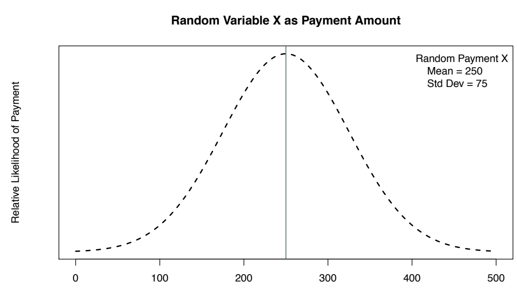

Figure 1.4: Random Variable X

[10] see Games & Klare, 1967 [8, chaps. 2 & 3]

usually represent this dashed line as $f(X)$, and have suggested that it represents the relative likelihood of observing any of the indicated values of X. This dashed line, $f(X)$, defines the moments of the distribution of X as:

$$E(X^r) = \sum_X X^r f(X) \tag{1.4}$$

and

$$E((X-\mu)^r) = \sum_X (X-\mu)^r f(X), \tag{1.5}$$

as the *central* moments of the distribution of X. In this case, Equation 1.5 is central, because the moments are centered around the population mean[11].

Furthermore, if $f(X)$ sums to 1.0,

$$\sum_X f(X) = 1.0$$

and $0 \leq f(X) \leq 1.0$ for all X, then $f(X)$ is a special type of function which we can call a probability distribution function - or more precisely a *density* function.[12]. When $f(X)$ is a probability density function, the first moment (i.e., when $r = 1$ in Equation 1.4) of the distribution of X is precisely the mean of the distribution. The second central moment of the distribution of X (i.e., when $r = 2$ in Equation 1.5) is the *variance* of the distribution. This is most easily seen if one interprets the function $f(X)$ as the relative frequency of observing the value X. In this case, if there are n total observations, then each value, $f(x)$, is just the number of observations of a particular value of X, divided by n, making this central moment $\sum \frac{(X_i-\mu)^2}{n}$, summed over all observations i - and this is the equation for variance that we all learned in school.

In other words, the foundation for risk-based pricing is statistics. While that may cause some to shiver a bit, it places everything that we do on a footing that is well understood.

We first must understand exactly what we expect (our representation of the expected value, or first moment, of our cash flow distributions, the random variables in which we have interest). We then must determine the inherent variability of the cash flows (the variance, or second moment, of those same distributions). Once these two critical measures have been estimated our analyses can follow.

1.4 A Nod to the Bayesian Approach

We must acknowledge up front, that we are partial to the Bayesian approach to probability[13]. This view provides a great deal of flexibility in the approaches

[11] Hoel, Port & Stone, (1971) introduce this concept of moments in a very intuitive way[12, p. 92]
[12] see Winkler (1972) [19, p. 57]
[13] see Winkler (1972) [19, p. 15]

that we often take. [14]

In a nutshell, the Bayesian view allows *probability distributions* to be a bit more elastic in their construct and in their use. To a strict traditionalist, a probability distribution always derives, somewhere in its origin, from something akin to an idealized experiment that relates to drawing different colored balls from an urn. Although I speak that in jest, it is not too far from the truth. The foundation for traditional probability lies in counting; how many possible outcomes exist and what is the relative frequency that can be ascribed to a particular outcome. Traditional probability theory (at least for discrete random variables) is grown from the field of counting.

In Bayesian analysis, we don't care quite so much from which field the distribution is grown. We can imagine that in some cases we really don't understand the counting process by which we could generate the probability distribution needed for the analysis of a random variable. Instead of agonizing over our inability to establish a formal counting schema - as Bayesians, we might simply turn to a nearby phenomenon that we believe to be fundamentally similar, and use a probability distribution that *seems* right, but that has no formal connection to the real variable that we are examining [15].

A Bayesian, might, for example, be faced with a need to generate likely outcomes for a financial product that has not yet been put on the market. A traditional probability approach would require that no likelihoods be estimated until an experiment of sufficient power has been performed to estimate likely outcomes. A Bayesian, on the other hand, might approach the problem by saying, "Well, we don't know for sure, but our best estimate is that this new product, X, will behave like an old product, Y. So let's use what we know about Y to estimate what we think will happen with X."

This sounds reasonable in many cases. To a Bayesian, having a formal probability distribution from which to estimate likely outcomes is nice, but not essential. The concept of probability distributions is broad enough to incorporate a "best guess" about probabilities. A Bayesian is willing to start to estimate likelihoods with a fairly broad array of 'things' that behave sort-of like probabilities, but may not really be such. Bayesians will incorporate expert guesses into an analysis - if that is all that they have to go from. They will utilize early results from experiments, long before the results should be considered statistically significant, if that is the best information that they have to use. And, even more strangely, a Bayesian has no problem updating probability distributions in the middle of an analysis once more information is known, once more results have come in. A Bayesian is willing to use whatever information about an event's likelihood may be available. There is no reliance upon a 'counting' framework from which to formally derive probabilities.

Bayesians, in other words, are just practical people who want to use whatever they can to help solve a likelihood problem. Maybe we weren't the best mathematicians - those that could derive the formal theoretical basis for even

[14] see pt. I Gelman, et al, 2004[9]

[15] see Fabozzi, 1997 [6, p. 377]

really complex problems - but we grasped the utility and the mechanisms of probability theory and we have found ways to solve practical problems using those tools. Our approach may not be *pure* to a formalist, but we are not paralyzed when actuarial data are lacking. We find a way to get things done.

The primary tools to a Bayesian are still called probability distributions. As mentioned, these can come from many different places, but they still must have the properties of probability distributions. Namely, $f(X)$ must sum to 1.0 and $0 \leq f(X) \leq 1.0$ for all X. We sometimes call them *prior* distributions (this is our way of saying to each other that we've used something before we have much data to base it on). We sometimes modify the results of early likelihood estimates by using *posterior* distributions.

As you will see, we Bayesians are quite flexible in our thinking, yet we engage in a form of rigor that allows us to speak with confidence when a problem must be solved. We have a formal method for working with the best information at hand. This series of books will present this methodology in full.

1.5 An Overview of Our Methodology

As indicated in several other places within this volume, our methodology for risk-based pricing involves a very formal and rigorous process - although we are flexible enough to work with **fuzzy sets** or *fuzzy* types of data. The process can be outlined, as follows.

1. First, we come to understand the contractual terms of an asset (think of something that generates cash flows like a loan) and identify all exogenous variables that can affect those flows.

2. Second, we create a model that generates the cash flows from that asset and create what we consider to be the *pristine* cash flows - those flows that are realized in the most 'pure' case - without call or put options exercised, without default or delinquency events and with steady-state scenarios for any set of exogenous variables - we call this our 'pristine' cash flow generator.

3. Third, we build a framework within which we can build operators to modify the contractual flows generated in our 'pristine' case - operators that represent probabilistic events like the exercise of put options or call options, or even default and delinquency events.

4. Fourth, we develop models that describe the relationships among our various input parameters - those affecting the pristine cash flows and those affecting the probabilistic events associated with our operators.

5. Fifth, we generate the pristine cash flows, and then operate upon those pristine cash flows with our probabilistic operators - connecting each input scenario to the appropriate probabilistic set of events.

6. Sixth, we value the resulting cash flows, using a discount rate that is appropriate for market conditions - which sometimes are derived from internal estimates of appropriateness and sometimes from market-based indicators.

7. We repeat items 5 and 6 repeatedly using enough different input scenarios and probabilistic operators that we come to understand both the cash flows and the valuation models of the asset for a given set of conditions, and finally come to understand the range and variations in value that can result.

This volume (Volume 1) of the Series addresses only items 1 and 2, above. In this volume we take the simple case of a fixed-rate loan and we examine how the contractual terms might be modeled. We create generators of pristine cash flows for this asset class. The second volume in this series will address items 3 and 4, and the third volume addresses the application of these in item 5. The fourth volume in this series addresses item 6 and the final volume in this series addresses item 7. The five volumes in the series will take the reader through the entire risk-based pricing process for fixed-rate loans. The series will also connect to risk-based pricing concepts, in general, as possible.

We look, in this series, at fixed-rate loans, using them as the prototype of our fixed-income asset category. But, we must remain cognizant that there are other categories of assets to which this methodology is easily applied; to floating rate loans, to adjustable rate mortgage loans, etc. We are describing a methodology that belongs, in general, to fixed-income analysis; we simply use fixed-rate loans as the vehicle by which we illustrate the application of that methodology.

Our approach involves the generation of *pristine* cash flows for a specific set of primary parameters surrounding a fixed-income asset. The parameters describe those basic terms consistent with the description of contractual loan terms, and these inputs perfectly define a set of *idealized* cash flows - cash flows that would result if everything went according to the pre-envisaged conditions or contract terms. Upon this set of pristine cash flows - or those expected based on the contractual agreements - we apply models of probabilistic events that shape and transform those pristine cash flows. It is the variability (think uncertainty) in these resulting cash flows - based upon the probabilistic events that shape the operators on the pristine cash flows - that determines the discount rate at which the flows are discounted.