Measuring the Value from Improved Predictions of Software Process Improvement Outcomes Using Risk-Based Discount Rates

Warren Harrison
David Raffo
John Settle
Portland State University

Abstract

It is desirable to evaluate software process improvements in terms of their expected benefits. However, such benefits cannot be considered "sure things". There is always an element of risk involved regarding whether the investment will pay off. We can imagine techniques that (for a price) may provide more accurate predictions of the benefits of a process change. Such improved predictions can reduce the perceived risk of the change. In this paper, we describe a way in which we can establish a monetary value for this reduction in perceived risk. Once we know how much a certain reduction in risk is worth to us, we have a convenient way for establishing the amount we are willing to invest in the improved prediction capability.

Introduction

The motivation for process improvements is better quality, more timely delivery or reduced effort. All of these can be viewed as positive cash flow due to cost savings. For instance, defects found early in the life cycle, reduce the cost of finding and fixing errors during testing. Naturally, most people accept the fact that in order to introduce process improvements, there are some up front costs for things such as training, process documentation, etc. These costs can be considered negative cash flows. Thus we have the classical investment model:

Present cash outflow ➔ future cash inflow

Obviously, one always has the opportunity to invest in alternate projects. Comparison of projects to be funded is traditionally known as Capital Budgeting. Most contemporary Capital Budgeting techniques utilize the concept of the "Time Value of Money", as embodied in one of the most common techniques: *Net Present Value* (NPV).

Briefly, the idea of Net Present Value analysis involves "discounting" a future cash flow at some discount rate, then subtracting out the current cash outlay, resulting in the expected value of the investment in today's dollars. For instance, it may require an investment of $50,000 to introduce design reviews into a project. In return, we may expect to receive $100,000 in benefits two years later when the product enters the testing phase. Thus, applying a 10% discount rate to the benefits expected to accrue two years in the future, we obtain the following value for applying a design review to the current project:

$$\text{NPV} = PV_{2,10\%}(100,000) - 50,000$$
$$32,645 = 100,000 / (1 + 0.10)^2 - 50,000$$

In most treatments of NPV, the discount rate is taken to be a given. However, in reality, the discount rate is comprised of two parts: a base, "risk-free" rate which reflects the value of money to be received in the future with certainty, and a "risk premium" that reflects the uncertainty of a future pay-off. In the case of a software process improvement such as a design review, the risk derives from the fact that the benefits of the process change are not a "sure thing". They may be either more or less beneficial than we expect. For instance, there may be few latent design errors to find, the review process itself may be ineffective, or the follow-through may be lacking so even when a problem is found, it is not addressed.

Obviously, the more variable the future return, the more risky the investment, the higher the risk premium and the greater the discount rate since a rational investor will expect a greater potential payback from a riskier investment.

Assessing the Value of "Refined Prediction Techniques"

The heart of deciding on a specific process change lies in predicting the payback from its use. These concepts lead to a logical way to evaluate the "value added" of techniques intended to better predict the effect of various process improvement projects. Suppose a firm has an opportunity to invest in a process improvement that requires an identifiable investment today, in return for a risky payoff in one year. Using Net Present Value Analysis, we discount the future cash inflow using a discount rate *k*, which is related to the uncertainty of the project. In turn, *k* depends on risk:

$$k = r_f + \phi = r_f + p(\text{Risk})$$

that is, the risk-free rate of return r_f (say the rate of return on Treasury Bills, currently about 5%), plus a risk premium ϕ, the product of a universal price of risk *p* times a risk measure.

For instance, say we have an opportunity to implement design reviews. Assume the cost of implementing this activity is $10,000. Further, assume that each flaw found during an inspection will result in avoiding a future correction that would cost $1,500 two years from now. Based on past data, we believe that there is a 5% chance we'll find no flaws, a 20% chance we'll find 10 flaws, a 50% chance we'll find 20 flaws, a 20% chance we'll find 30 flaws and a 5% chance we'll find 40 flaws. This is summarized as follows:

Number of Flaws	Probability	Net Value (undiscounted)
0	5%	(10,000)
10	20%	5,000
20	50%	20,000
30	20%	35,000
40	5%	50,000

The "expected undiscounted net value" is $20,000 (we do not consider the case where the expected cash flow actually changes in this paper), however, the standard deviation of the outcomes is $13,484. This information can be used to derive a measure of the risk, Risk_{wo} which would yield a risk-based discount rate k_{wo}.

Assume an improved prediction of the expected outcome of implementing design reviews were available for some additional expense. This technique could involve better/additional information, or an improved estimation model. For our current purposes, the nature of the technique itself is of no importance, and we will simply refer to it as "the technique". Now using this "new technique", we obtain the following expected results:

Number of Flaws	Probability	Net Value (undiscounted)
0	5%	(10,000)
10	10%	5,000
20	70%	20,000
30	10%	35,000
40	5%	50,000

The "expected undiscounted net value" is still \$20,000, however due to the refinement, the uncertainty of the outcome has been reduced, with the new standard deviation of the outcomes being only \$11,677.

Let the relative change in the risk estimate (i.e., $Risk_w / Risk_{wo}$) obtained by the "technique" be denoted λ_{Risk}. A reasonable measure of risk is the standard deviation of the estimate. Then the NPV of the proposed process improvement project estimated with and without the refinement are given, respectively, by:

$$NPV_{wo} = \text{Future Inflow} / [1 + r_f + p(Risk_{wo})]^N - \text{Current Outflow},$$

and

$$NPV_w = \text{Future Inflow} / [1 + r_f + p(\lambda_{Risk} \times Risk_{wo})]^N - \text{Current Outflow}.$$

or, for this specific example (λ_{Risk} = 11,677/13,484 = 86%):

$$NPV_{wo} = 20{,}000 / [1 + 0.05 + p(Risk_{wo})]^2 - 10{,}000,$$

and

$$NPV_w = 20{,}000 / [1 + 0.05 + p(86\% \times Risk_{wo})]^2 - 10{,}000.$$

Assuming a hurdle rate of 20%, we can conclude the price of the risk originally assumed is 15%. While this may not be explicitly articulated by management, it seems reasonable to assume that any premium over the risk-free rate is attributable to risk. Therefore the risk presented using the "improved technique" is 86% of that, or 12.9%, for a new discount rate of 17.9% instead of 20%. Under these assumptions, the Net Present Value using both valuation methods is:

$$NPV_{wo} = 20{,}000 / [1.20]^2 - 10{,}000$$
$$= 13{,}888$$

and

$$NPV_w = 20{,}000 / [1.179]^2 - 10{,}000$$
$$= 14{,}388$$

Thus, the value of the additional precision is \$500 *for this project*. In order to address the true value of the technique we would want to apply the analysis to a portfolio of projects. That said, this analysis is not overly simplified. Many project managers make process decisions based on the outcome on their current project, and have only passing interest at building the overall software engineering capability within the organization.

The risk and risk premium can be elaborated somewhat using Capital Asset Pricing Model concepts, although we do not whole-heartedly recommend the use of the CAPM. In this model, project risk β_p is defined as:

$\beta_p = \rho_{pm} \times \sigma_p/\sigma_m$.

where σ_m denotes the standard deviation of the rate of return on a market portfolio, σ_p the standard deviation of the return on the project, and ρ_{pm} the correlation coefficient between the two. In particular, this relation is defined correspondingly for the project with and without the "technique:"

$\beta_{wo} = \rho_{wo,m} \times \sigma_{wo}/\sigma_m$, and

$\beta_w = \rho_{wm} \times \sigma_w/\sigma_m$.

If we are willing to assume that $\rho_{wm} = \rho_{wo,m}$, which does not seem to be unduly restrictive, then we can relate λ_{Risk} in the above equations to the CAPM as follows:

$\lambda_{Risk} = \beta_w/\beta_{wo} = \sigma_w/\sigma_{wo}$, which is the same specification for λ as used above.

The intuitive appeal of this approach is that it is relatively easy to think in terms of standard deviations, and with a modest increase in the restrictions of assumptions, this can be made consistent with the CAPM.

Benefits of Improved Predictions

Two benefits derive from a "technique" used in this manner. One benefit is not easily quantifiable: it is the added comfort level for accepted projects. The other is eminently quantifiable: the value saved by avoiding erroneous acceptance or rejection of projects.

In Figure 1, the lines $^H k_w$ and $^H k_{wo}$ represent an NPV equal to zero, with discount rates of k_w and k_{wo}, respectively. Thus, Region I represents a project which is determined by the refinement to be less risky than the original estimate, and thus possesses an NPV greater than zero when discounted at k_w, but not at k_{wo}. Region II represents a project which both discount rates would result in an NPV greater than zero. Region III represents a project that would have an NPV less than zero regardless of which of the two discount rates were used. Region IV represents a project in which the discount rate k_{wo} provides an NPV greater than zero, but k_w does not – a project which appeared less risky than it turns out to be with the refinement. The effect of a single application of the refinement is material if the project falls either within Regions I or IV. In this case, it can be measured as the amount gained from a project that would not have otherwise been undertaken or the amount saved by not undertaking a project that would have otherwise been undertaken.

From a more global perspective, the "technique" would result in adoption of all projects that would fall in Region I of Figure 1, which would otherwise be rejected, and the rejection of all projects that fall in the Region IV which would have otherwise been accepted. Finally, management benefits from the added comfort level of knowing more about the true value of those projects in the upper right quadrant that it would have adopted anyway.

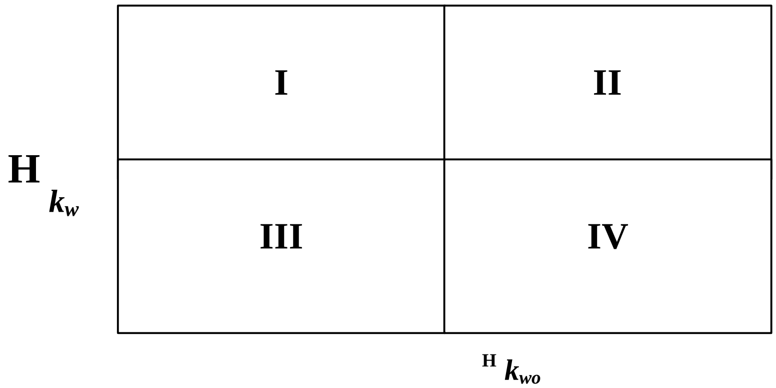

Figure 1

Ongoing Work

In this treatment, we have limited our analysis to the situation where a refined predictive capability yields the same expected value, but with a greater degree of certainty. Obviously, an improved predictive capability is likely to actually yield a different expected value. To some extent, this issue is addressed by our discussion of Figure 1. However, we are currently working on an explicit treatment of this situation.

For the purposes of this paper, we consider only process improvements. However, this approach can be used to evaluate the value added by any enhanced prediction system. So for instance, we can imagine this approach being applied to the value added by improved cost estimation techniques or metrics programs.

The issue of the cost per unit of risk is somewhat difficult when the conventional approach to making decisions about process improvements seldom deals with cash flows, net present values, discount rates and risk premiums. In this paper we have assumed that some sort of “hurdle rate” has been prescribed, and that management has, knowingly or unknowingly associated a cost of risk via an implied risk premium. Obviously practical application of this work is dependent upon some sort of systematic way of evaluating the return on investment of a process improvement.