

system is used for monitoring without modifications specifically designed for project management. For example, an existing cost tracking system oriented to shop operations would be inappropriate for a project with major activities in the area of research and development. But as we just noted, the PM's problem is to fit standard information into a reporting and tracking system that is appropriate for the project.

The real message carried by project reports is in the comparison of actual activity to plan and of actual output to desired output. Variances are reported by the monitoring system, and responsibility for action rests with the controller. Because the project plan is described in terms of scope, time, and cost, variances are reported for those same variables. Project variance reports usually follow the same format used by the accounting department, but at times they may be presented differently.

10.3 EARNED VALUE ANALYSIS

Thus far, our examples have covered monitoring for parts of projects. The monitoring of performance for the entire project is also crucial because performance is the *raison d'être* of the project. *Individual* task performance must be monitored carefully because the timing and coordination between individual tasks is important. But overall project performance is the crux of the matter and must not be overlooked. One way of measuring overall performance is by using an aggregate performance measure called *earned value*.

The Earned Value Chart and Calculations

There is a considerable body of literature devoted to earned value. To note only a few of the available items, see Anbari (2003), the Flemming references, Hatfield (1996), Project Management Institute (2004), and Singletary (1996). One must, however, exercise some care when reading any article on the subject. Various ratio index numbers have almost as many names (and hence, acronyms) as there are writers. Some authors take further license; see Brandon (1998) for instance, and also see the subsequent *Project Management Journal's* Correspondence column (September 1998, p. 53) for readers' reactions. We will adopt and stick to the PMBOK version of things, but will also note the names and acronyms used by Microsoft's Project®.* Any other names/acronyms will be identified with the author(s). A history of earned value from its origin in PERT/Cost together with its techniques, advantages, and disadvantages is reported in a series in *PM Network* starting with Flemming et al. (1994).

A serious difficulty with comparing actual expenditures against budgeted or *baseline* expenditures for any given time period is that the comparison fails to take into account the amount of work accomplished relative to the cost incurred. The earned value of work performed (*value completed*) for those tasks in progress is found by multiplying the estimated percent physical completion of work for each task by the planned cost for those tasks. The result is the amount that should have been spent on the task thus far. This can then be compared with the actual amount spent.

Making an overall estimate of the percent completion of a project without careful study of each of its tasks and work units is not sensible—though some people make such estimates nonetheless. Instead, it is apparent that at any date during the life of a project the following general condition exists: Some work units have been finished, and they are 100 percent complete; some work units have not yet been started, and they are 0 percent complete; other units have been started but are not yet finished, and for this latter group we may estimate a percent completion.

*Earlier versions of Microsoft Project® used a slightly different way to calculate earned value variances.

As we said, estimating the “percent completion” of each task (or work package) is nontrivial. If the task is to write a piece of software, percent completion can be estimated as the number of lines of code written divided by the total number of lines to be written—given that the latter has been estimated. But what if the task is to test the software? We have run a known number of tests, but how many remain to be run?

There are several conventions used to aid in estimating percent completion:

- *The 50–50 rule.* Fifty percent completion is assumed when the task is begun, and the remaining 50 percent when the work is complete. This seems to be the most popular rule, probably because it is relatively fair and doesn’t require the effort of attempting to estimate task progress. Since it gives credit for half the task as soon as it has begun, it is excessively generous at the beginning of tasks, but then doesn’t give credit for the other half until the task is finally complete, so is excessively conservative toward the end of tasks, thereby tending to balance out on an overall basis.
- *The 0–100 percent rule.* This rule allows no credit for work until the task is complete. With this highly conservative rule, the project always seems to be running late, until the very end of the project when it appears to suddenly catch up. Consequently, the earned value line will always lag the planned value line on the graph.
- *Critical input use rule.* This rule assigns task progress according to the amount of a critical input that has been used. Obviously, the rule is more accurate if the task uses this input in direct proportion to the true progress being made. For example, when building a house, the task of building the foundation could be measured by the cubic yards (or meters) of concrete poured, the task of framing the house could relate to the linear feet (meters) of lumber used, the roofing task could relate to the sheets of 4 X 8 foot plywood used, and the task of installing cabinets might be measured by the hours of skilled cabinet labor expended.
- *The proportionality rule.* This commonly used rule is also based on proportionalities, but uses time (or cost) as the critical input. It thus divides actual task time-to-date by the scheduled time for the task [or actual task cost-to-date by total budgeted task cost] to calculate percent complete. If desirable, this rule can be subdivided according to the subactivities within the task. For example, suppose progress on a task is dependent on purchasing (or building) a large, expensive machine to do a long and difficult task, but the machine itself does not make any substantial task progress. We could create a table or graph of the use of money (or time, if the machine had to be built) relative to task progress which would show a large amount of money (or time) being expended up front for the machine, but with little (or no) progress per se being made. This would then be followed by a continuing expenditure of a smaller stream of money (or time) to run the machine and finish the job, perhaps in direct proportion to the progress.

These rough guides to “percent completion” are not meant to be applied to the project as a whole, though sometimes they are, but rather to individual activities. For projects with few activities, rough measures can be misleading. For projects with a fairly large number of activities, however, the error caused by percent completion rules is such a small part of the total project time/cost that the errors are insignificant. More serious is the tendency to speak of an entire project as being “73 percent complete.” In most cases this has no real meaning—certainly not what is implied by the overly exact number. Some authors assume that making estimates of percent completion is simple (Brandon, 1998, p. 12, col. 2, for instance). The estimation task is difficult and arbitrary at best, which is why the 50–50 and other rules have been adopted.

A graph illustrating the concept of earned value such as that shown in Figure 10-5 can be constructed using the above rules and provides a basis for evaluating cost and scope to date. If the total value of the work accomplished is in balance with the planned (baseline)




Figure 10-5 Earned value chart.

cost (i.e., minimal scheduling variance), as well as its actual cost (minimal cost variance), then top management has no particular need for a detailed analysis of individual tasks. Thus the concept of earned value combines cost reporting and aggregate scope reporting into one comprehensive chart. The baseline cost to completion is indicated on the chart and referred to as the budget at completion (BAC). The actual cost to date can also be projected to completion, as will be shown further on, and is referred to as the estimated cost at completion (EAC).

We identify several variances on the earned value chart following two primary guidelines: (1) A negative variance is "bad," and (2) the cost and schedule variances are calculated as the earned value minus some other measure. Specifically, the *cost* (or sometimes the *spending*) *variance* (CV) is the difference between the amount of money we budgeted for the work that has been performed to date, that is, the *earned value*, EV, and the actual cost of that work (AC). The *schedule variance* (SV) is the difference between the EV and the cost of the work we scheduled to be performed to date, or the planned value (PV). The *time variance* is the difference in the time scheduled for the work that has been performed (ST) and the actual time used to perform it (AT).^{*} In compact form,

$$EV - AC = \text{cost variance (CV, overrun is negative)}$$

$$EV - PV = \text{schedule variance (SV, behind is negative)}$$

$$ST - AT = \text{time variance (TV, delay is negative)}$$

Typically, variances are defined in such a way that they will be negative when the project is behind schedule and/or over cost. As we have noted, however, this practice is not universal either in the literature or in practice.

^{*}A fourth variance can be found. It is the difference between the cost that the project budget says should have been expended to date (PV) and the actual cost incurred to date by the project (AC). $PV - AC$ is what we call the *resource flow variance*. (Note that the resource flow variance is not a "cash flow" variance.)

The variances are also often formulated as ratios rather than differences so that the cost variance becomes the Cost Performance Index (CPI) = EV/AC , the schedule variance becomes the Schedule Performance Index (SPI) = EV/PV , and the time variance becomes the Time Performance Index (TPI) = ST/AT . Use of ratios is particularly helpful when an organization wishes to compare the performance of several projects (or project managers), or the same project over different time periods. As we just noted, however, the accuracy and usefulness of all these performance measures depend on the degree in which estimates of percent completion reflect reality.

Cost and schedule variances (or CPI and SPI) are very commonly used. A short example illustrates their application. Assume that operations on a work package were expected to cost \$1,500 to complete the package. They were originally scheduled to have been finished today. At this point, however, we have actually expended \$1,350, and we estimate that we have completed two-thirds of the work. What are the cost and schedule variances?

$$\begin{aligned}
 \text{cost variance} &= EV - AC \\
 &= \$1,500(2/3) - 1,350 \\
 &= -\$350 \\
 \text{schedule variance} &= EV - PV \\
 &= \$1,500(2/3) - 1,500 \\
 &= -\$500 \\
 \text{CPI} &= EV/AC \\
 &= \$(1,500(2/3))/1,350 \\
 &= .74 \\
 \text{SPI} &= EV/PV \\
 &= \$(1,500(2/3))/1,500 \\
 &= .67
 \end{aligned}$$

In other words, we are spending at a higher level than our budget plan indicates, and given what we have spent, we are not as far along as we should be (i.e., we have not completed as much work as we should have). We can also use SPI to calculate the time variance TV if we realize that the scheduled time, ST, should conceptually be in proportion to (EV/PV): $ST = (AT)(EV/PV)$. Since $TV = ST - AT$, then $TV = (AT)((EV/PV) - 1) = (AT)(SPI - 1)$. (This can be derived through simple trigonometry.)

It is, of course, quite possible for one of the indicators to be favorable while the other is unfavorable. We might be ahead of schedule and behind in cost, or vice versa. There are six possibilities in total, all illustrated in Figure 10-6. The scenario shown in Figure 10-5, where both SV and CV are negative, is captured in arrangement *d* of Figure 10-6. The example immediately above, which also results in negative values of SV and CV, is arrangement *c* of Figure 10-6. Barr (and others) combines the two indexes, CPI and SPI, to make a type of "critical ratio" (described further in Chapter 11) called the Cost-Schedule Index (Barr, 1996, p. 32).

$$\begin{aligned}
 \text{CSI} &= (\text{CPI})(\text{SPI}) \\
 &= (EV/AC)(EV/PV) \\
 &= EV^2/(AC)(PV)
 \end{aligned}$$

In our case,

$$\begin{aligned}
 &= \$(1,500(2/3))^2/(1,350)(1,500) \\
 &= \$1,000,000/2,025,000 \\
 &= 0.49
 \end{aligned}$$

As Barr writes, $\text{CSI} < 1$ is indicative of a problem.

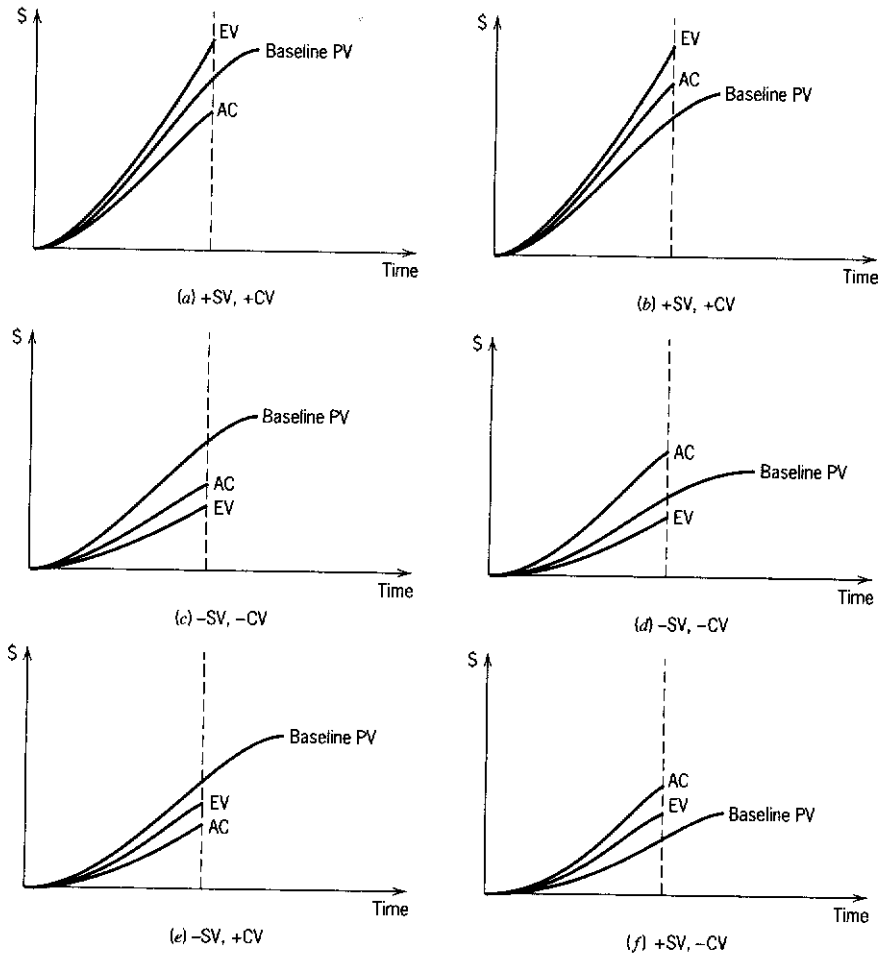


Figure 10-6 Six possible arrangements of AC, EV, and baseline PV resulting in four combinations of positive and negative schedule variance (SV) and cost variance (CV). (Figure 10-5 is arrangement *d*.)

One can continue the analysis to forecast the future of this work unit under the condition when no measures are taken to correct matters. The cost to complete the work unit can be estimated as the budgeted cost of the entire unit, less the earned value to date, adjusted by the CPI to reflect the actual level of performance. The budget at completion (BAC) in our example is \$1,500. The earned value to date (EV) is $\$1,500 \times 2/3 = \$1,000$. The estimated cost to complete (ETC), assuming the same cost efficiency level, can be projected as:

$$\begin{aligned} \text{ETC} &= (\text{BAC} + \text{EV})/\text{CPI} \\ &= (\$1,500 + 1,000)/0.74 \\ &= \$676 \end{aligned}$$

The estimated cost at completion (EAC)—and we use Barr’s term (1996) rather than Microsoft’s FAC or any of the many other names in the literature—is the amount expended to date (AC) plus the estimated cost to complete (ETC):

$$\begin{aligned} \text{EAC} &= \text{ETC} + \text{AC} \\ &= \$676 + 1,350 \\ &= \$2,026 \end{aligned}$$

Table 10-1 Earned Value Example (today is day 7)

Activity	Predecessor	Days Duration	Budget, \$	Actual Cost, \$
a	—	3	600	680
b	a	2	300	270
c	a	5	800	
d	b	4	400	
e	c	2	400	




Figure 10-7 Example AON diagram.

shown in Figure 10-7, where path **a-c-e** is the critical path, with project completion expected at day 10. What has actually happened in the project is that the first activity, **a**, took 4 days instead of the planned 3 days to complete, delaying the start of both activities **b** and **c**. Activities **b** and **d** are proceeding as expected, except of course for their one-day delay in initiation, but anyway, path **a-b-d** was not the critical path for the project.

Activities **a** and **b** are both completed and their actual costs are shown in Table 10.1. (The costs to date for activities **c** and **d** are not known.) However, due to its delay, activity **a** cost \$80 more than budgeted. Hence, the project manager is trying to cut the costs of the remaining activities, and we see that activity **b** came in \$30 under budget, which helps but does not fully offset the previous overrun.

The baseline budget (PV) using the 50–50 rule is calculated in Figure 10-8 and graphed in Figure 10-10 (solid line) where the BAC is listed as \$2,500. The project’s status and earned value (EV) as of day 7 are given in Figure 10-9 and shown in Figure 10-10 as a dotted line. Included in the figure is the actual cost (AC, the dashed line in Figure 10-10) for the two completed activities. As shown in Figure 10-10, the schedule variance is currently 0 and the cost variance is $\$1,500 - 950 = +550$.

But notice how these figures do not give a very accurate picture of project progress. The earned value up to now has been trailing the baseline and has only caught up because the 50–50 rule doesn’t have any activity beginning or ending at day 6; however, with expediting activity **c**, we may in fact be back on schedule by day 8. The cost variance, however, is highly affected by the fact that actual costs are not recorded until the activity is 100 percent complete, combined with the impact of the 50–50 rule. The result is that the baseline and earned value cost figures will tend to converge when activities begin but the actual costs will lag them considerably. Even though the proportionality rule would more accurately delay the aggregation of earned value costs, there would still be a positive bias if the actual costs were not calculated until the activities were completed. It would be more accurate, but considerably more complex, to apportion actual costs according to percentage activity completion. These effects are illustrated further in some of the problems at the end of the chapter.

rather than the original estimate of \$1,500. For a complete description of this approach to estimating the total cost of a work unit, or a set of work units, see Barr (1996) and the Flemming and Koppelman works. We also could consider the ETC as a probabilistic number, and, given upper and lower bounds and an estimated distribution for ETC, we can easily apply simulation to find a distribution for EAC.

The PMBOK also considers two other ways to calculate ETC. If you assume that the team is now going to do all the future work at the originally budgeted rate, then $ETC = BAC - EV$. Another possible assumption is that the team will do all the future work at a rate that considers both the cost and schedule deviations to date, in which case $ETC = (BAC - EV)/(CPI)(SPI)$ or just $(BAC - EV)/CSI$. Recognize that if the team is behind schedule, this implies a crash schedule to catch up, which will further inflate the costs.

Note: The planned values (PV) for each task would normally be known from the WBS and budget for the project tasks. However, when distributing PV over the scheduled time for a task (e.g., 3 weeks) for comparison to EV for monitoring purposes during the actual project, consideration should be given to how each task's EV is going to be determined. For example, if the PV is assumed to be generated in proportion to the time spent on the task, then the use of a 0–100 percent rule for EV will result in the project always appearing behind schedule. This is fine if the person monitoring the project understands that this difference of measurement methods is the cause of the “behind schedule” appearance. However, an alternative approach would be to distribute the PV for each task in the same manner that the EV is going to be measured for each task, and then the comparison of the EV to the PV will be more realistic.

Thus far, the focus has been on measuring performance on a work unit rather than on the project as a whole. Where dealing with a specific work unit, the estimates of costs and time can be fairly precise. Even the estimate of percent completion can be made without introducing too much error when using, as we did above, the proportionality rule. Given the relatively short time frame and relatively small cost compared to the whole project, errors are not apt to be significant. Random errors in estimating will tend to cancel out and we can aggregate the work unit data into larger elements, e.g., tasks or even the whole project. (Bias in estimating is, of course, a different matter.) Although the measurement error may be minimal, for most projects there is still no sound basis for estimating percent completion of the project as a whole.

Even if this aggregation is feasible, the use of earned value analysis for forecasting project schedules and costs does not mean that the forecasts will make it possible to correct malperformance. The case for remediation is not hopeful. In a study of more than 700 projects carried out under Department of Defense contracts, the chances of correcting a poorly performing project more than 15 percent complete were effectively nil (Flemming et al., 1996). The study concludes that if the beginning of the project was underestimated and took longer and cost more than the plan indicated, there was little or no chance that the rest of the project would be estimated more accurately (p. 13ff). For relatively small deviations from plan, the PM may be able to do a lot of catching up.

If the earned value chart shows a cost overrun or scope underrun, the PM must figure out what to do to get the system back on target. Options include such things as borrowing resources from activities performing better than expected, or holding a meeting of project team members to see if anyone can suggest solutions to the problems, or perhaps notifying the client that the project may be late or over budget. Of course, careful risk analysis at the beginning of the project can do a great deal to avoid the embarrassment of notifying the client and senior management of the bad news.

Example: Updating a Project's Earned Value

We use a simple example to illustrate the process of determining the baseline budget and interim earned value and actual costs for a project. Table 10-1 presents the basic project information, and updated information as of day 7 in the project. The planned AON diagram is

	Day											
Activity	0	1	2	3	4	5	6	7	8	9	10	
a		300			300							
b					150	150						
c					400				400			
d							200			200		
e										200	200	
Total		300			300	550	150	200		400	400	200
Cum. Total		300	300	600	1150	1300	1500	1500	1900	2300	2500	

Figure 10-8 Example baseline (PV) budget using the 50-50 rule.

	Day										
Activity	0	1	2	3	4	5	6	7	8	9	10
a		300			300						
b						150	150				
c						400					
d								200			
e											
EV		300			300	550	150	200			
Cum. EV		300	300	300	600	1150	1300	1500			
Actual Cost					680		270				
Cum. AC		0	0	0	680	680	950	950			

Figure 10-9 Example status at day 7.




Figure 10-10 Example earned value chart at day 7.