Tying Project Measures to Performance Incentives

David P. Quinn
Borland Software Corporation

In the struggle to set expectations and reward good performance, organizations sometimes tie project measures to performance incentives. To avoid the pitfalls sometimes caused by this link, we must explore the nature of project measures and performance and how to possibly link project measures to incentives so the intended behavior occurs.

Any measure when viewed in isolation is open to misinterpretation and misuse. Without the full context of other measures, we may unjustly attribute good or poor performance based on viewing a single measure as an absolute. This is why linking incentives to individual measures can be detrimental to what the organization is trying to accomplish. People tend to focus on meeting a measure instead of accomplishing an outcome. People tend to focus on themselves instead of the team or the organization, causing problems for others and negatively impacting the overall organization's performance. In fact, linking incentives to measures often has some unintended consequences.

Effort Variance as an Example

Effort variance is a very useful measure when used for project management as an indicator. It can show problems with estimates, processes, project scope, and performance. These problem areas have a different meaning depending on whether effort variance is indicating less effort expended than estimated or more effort expended than estimated.

In the case of less effort expended than estimated (e.g., 100 hours estimated but only 80 expended for a –20 percent effort variance), here are some causes of why the variance may have occurred:
- Estimate
    - Engineers padded their estimates.
    - Estimation parameters were wrong.
- Scope
    - Lack of understanding of scope.
    - Scope reduced but estimate not updated.
- Process
    - Steps skipped.
    - A process improvement occurred.
- Performance
    - Hours worked were not recorded/reported.
    - Brilliant work was performed.
    - Missed one or more requirements.
    - Did not complete the task.
    - Allowed poor quality in order to meet a deadline.

Of these causes, only two (brilliant performance and process improvement) should result in a reward through an incentive. The other causes should result in some sort of remedial action.

On the other hand, when looking at more effort expended than estimated (e.g., 100 hours estimated but 120 expended for a +20 percent effort variance), here are some causes of why the variance may have occurred:
- Estimate
    - Someone lowered the original estimates to meet mandated cost/schedule.
    - Estimation parameters were wrong.

> *"The biggest potential problem with linking incentives to effort variance is that people may game the system."*

- Scope
    - Lack of understanding of scope.
    - Scope increased but the estimate was not updated.
    - Customer was indecisive on the requirements.
- Process
    - Process is inefficient.
    - The process does not match the customer's needs.
    - Unnecessary/inappropriate steps were taken.
    - The work from someone earlier in the process made this person's work more difficult.
- Performance
    - Meets initial estimate but does not meet the modified estimate to meet budget/schedule (i.e., someone changed the estimate without changing scope just to make the numbers match the preferred schedule).
    - Poor work performance.
    - Added capability the customer did not request.

As with the previous example, not all of these causes should detract from incentives. Only poor work performance and using the wrong parameters are signs that the engineer needs to adjust his/her behavior. Many of the others are organizational or managerial faults that impact the engineer's performance.

Before rewarding or punishing for effort variance, there are other factors to consider. Perhaps it is okay that effort variance is above estimate if the quality of code leads to reduced effort variance for testing (i.e., there is no real impact to schedule) or reduced rework to correct defects. Perhaps the minor changes in scope (that no one felt required changing in the estimates) really hit harder than thought and showed up in the effort variance. You would not know if the latter case was true unless requirements volatility measures were gathered.

Other Indicators

The impact of effort variance on projects and on the organization must be examined to get an adequate picture of the importance of effort variance. For instance, a team of developers may take shortcuts to reduce effort variance, but the shortcuts negatively impact quality. That may increase the effort variance of the test team who must perform more test cycles than estimated. The developers may get rewarded for their improved effort variance at the expense of the test team.

Significant effort variance should result in a change in schedule performance. If team members' effort variance is –20 percent, the schedule should see a comparable variance (i.e., the schedule variance should be roughly –20 percent). When there are major gaps between the two variances (for example, –15 percent effort variance and –5 percent schedule variance), this should signal that something may be wrong. The organization needs to investigate why the effort per-

formance is not impacting the schedule more. Perhaps other factors inhibit an improvement in schedule variance.

Other Problems With Linkage

There are other problems with linking incentives to individual measures. Based on the organizational learning that occurs from doing projects, performance that once provided a bonus for people will no longer result in a bonus as the organization adjusts process performance measures. The following provides an example for effort variance.

One of the reasons to gather effort information is to ensure your estimating model is correct. If team members show they are expending far fewer hours than estimated on a regular basis (better than –10 percent effort variance), the organization should update the estimation model to reflect this performance. This will provide more accurate effort estimates in the future. However, once the organization modifies the estimation model, team members will no longer qualify for incentives, even though they perform at the same level the organization rewarded before.

But not updating the estimation model is wrong. If the customer continues to see that effort variance is significantly below estimates, the customer is likely to assume that engineers are padding their estimates. Trust is broken and the customer will require that the organization reduce its effort estimates on future projects no matter how reasonable the estimate may be. The organization must update the estimation model despite the impact on the incentives.

Performance Is Relative

Just as one measure by itself may not tell the whole story, sometimes that one measure hides the truth. While other measures may give a more accurate picture of where performance stands as a whole, measures may indicate where an individual engineer stands against his/her peers. A single measure by itself may indicate that an engineer is performing well, but in comparison to other engineers, the engineer's performance may be lacking.

While it may appear logical to reward someone who has a –10 percent effort variance, it does not make sense if the organization's average effort variance is –15 percent. Technically, this person's performance is slowing down the organization, and this person is not performing up to the level of his/her peers. Likewise, it may be good to reward someone with a +5 percent effort variance if the rest of

Measure	High	Medium	Low	Unsatisfactory
Effort Variance	Any variance better than -X percent.	Variance between -X percent and -Y percent.	Variance between -Y percent and +Z percent.	Variance greater than Z percent.
Schedule Variance	Any variance better than -X percent.	Variance between -X percent and -Y percent.	Variance between -Y percent and +Z percent.	Variance greater than Z percent.
Quality	Any variance better than X percent of the team average.	Variance between X percent and Y percent of the team average.	Variance between Y percent and Z percent of the team average.	Defect rate worse than Z percent of the team average.
Process Compliance	Recommended process improvement accepted for implementaion and complied with defined processes.	Participated in process improvement activities and complied with defined processes.	Complied with defined processes on regular basis.	Inconsistent use of defined processes.

Assumptions   X,Y, and Z may be different for each measure.
Quality may move from being relative to team members to being organizationally relative.
More measures could be added if desired.

Table 1: *Example Performance Incentive Structures*

the organization is performing at a +12 percent effort variance.

There is a problem with this comparison though. The engineer with a –10 percent effort variance while everyone else has a –15 percent effort variance may be handling all the tough, complex tasks. Comparing performance between peers using measures like effort variance is not as wise as it may appear.

Other Potential Problems

Someone who consistently outperforms the estimated effort does something different from everyone else. If that person does not share that something different with other people, the incentives are reinforcing the wrong behavior. As someone discovers a process improvement or other type of improvement, the person should share that information, and the organization should reward that sharing. Someone who does not share improvement information is not acting in the organization's best interest.

The biggest potential problem with linking incentives to effort variance is that people may game the system. This usually takes two forms: people work extra hours but do not record them, or people place the hours worked against a different activity. In the latter case, these hours are not necessarily charged to non-billable activities. People may charge the hours to activities on other projects with available budget.

In either case, the organization is not getting accurate information on what it takes to get a project done. The organization is no longer able to learn how accurate the estimates are or if a process improvement occurred. Organizations that find people gaming the system because of the linkage between incentives and measures usually have to take one of

two solutions. The first is to remove the linkage. The second is to have a policy that makes entering inaccurate effort data a cause for dismissal. Organizations tend to choose the former rather than the latter because the organization does not want to lose good people.

Recognizing and Rewarding Outstanding Performance

Project performance measures can be a contributor to recognizing outstanding performance, but there should not be a direct correlation between a single measure and an incentive. A measure that makes up X percent of an incentive (no matter how small X is) is more likely to be gamed, destroying any chance to accurately measure organizational performance. A better approach would be to group a number of measures together and reward them as an overall performance.

For instance, an organization may rate a software engineer on effort, schedule, quality, and process compliance (see Table 1). The organization may rate each factor as high, medium, low, or unsatisfactory (see Table 2). The combination of highs,

Table 2: *Compensation Formulas*

Highs	Mediums	Lows	Unsatisfactory	Percent Compensation
4	0	0	0	100%
3	1	0	0	95
3	0	1	0	90
2	2	0	0	90
2	1	1	0	80
2	0	2	0	75
1	3	0	0	75
1	2	1	0	60
0	4	0	0	50
1	1	2	0	50
0	3	1	0	25
1	0	3	0	25
0	2	2	0	15
0	1	3	0	10
0	0	4	0	0
-	-	-	1 or more	0

mediums, and lows indicate the incentive reward. Getting a high rating in all four factors would get the maximum incentive. Getting three high ratings and one medium rating would get a certain percentage. The organization would further adjust incentives for different combinations of highs, mediums, and lows. Any unsatisfactory rating would automatically eliminate the incentive.

This helps the engineer understand that each of these measures is important but no one measure is more important than the other measures. It also allows for the fact that the measures are interdependent and focuses on how the organization views technical performance as a whole.

Personal Experience With Unintended Consequences

I have two sons, David and Mark. Like most young boys, they constantly want raises in their allowances. I am a firm believer that an increase in wages results from an increase in responsibility. I also am a firm believer in trying to get the boys to do chores I do not find entirely enjoyable. Therefore, I needed to devise a way to allow them to earn more money while making my life easier.

I targeted mowing the lawn, a chore I dislike. Unfortunately, they were too young to mow the lawn. However, before I mow the lawn each week, there is another necessary chore. We have two dogs: Ace, a retired racing greyhound, and Amigo, a longhaired Chihuahua. Their daily routine generates a set of piles that someone must gather before I can mow the backyard. That became their chore.

David, the older son, gets $2 for cleaning the piles in the backyard. When David is done, Mark gets $1 to find any missed piles. However, I add an incentive to this. For every Ace pile Mark finds, Mark gains 25 cents and David loses 25 cents. For every Amigo pile Mark finds, Mark gains 10 cents and David loses 10 cents. David cannot go below $1 total but Mark has a limitless incentive. When Mark finishes his chore, I do a final inspection of the backyard. As with David, Mark loses 25 cents for every Ace pile I find and loses 10 cents for every Amigo pile I find. It seemed like a great system.

One spring I sent David out to do his task. I lost track of Mark shortly after David started but found him shortly before David was done. Since it was starting to get dark, I started Mark on his chore while David finished the last third of the yard.

As they both continued their chores, David suddenly announced that one of the Ace artifacts was covered with grass. I did not think much of it until he announced a few seconds later that other artifacts were covered by grass. Mark commented that more artifacts are likely covered by grass. I was able to put two and two together and pointedly asked Mark if he had been covering up the artifacts. His face turned white, his jaw dropped, and he meekly let out a "Yes."

I immediately sent him in the house to get his bath and go to bed. I also let him know he forfeited his money for doing the chore. As soon as the door closed, I laughed so hard I thought I would cry. I did not realize my performance incentives would create that type of behavior. You really have to be careful when linking performance measures to incentives.

Conclusion

Gathering and using measures is an essential part of business. Measures provide outstanding insight into current status, possible problems, and process improvements. However, one measure does not provide enough insight by itself. Measures must be viewed in total as the result of one measure may impact another measure. The data gathered for the measures must be unquestionably accurate. When organizations tie incentives to project measures, people often report the data inaccurately to meet the incentives. Once organizations realize data is inaccurate, they tend to break the link between incentives and measures. They recognize that data needs surpass the incentive needs.◆

About the Author

David P. Quinn is a principal consultant with Borland Software Corporation. He has more than 20 years of systems and software engineering experience. His past positions ranged from software developer to project manager to mid-level manager as well as Engineering Process Group leader. Quinn is an authorized Lead Appraiser and past member of the Software Engineering Institute's Capability Maturity Model® Advisory Board.

30 Quail Ridge RD
Hanover, PA 17331
Phone: (717) 632-1047
E-mail: david.quinn@borland.com