

December 2007

Beyond Chaos: Examining IT Project Performance - Research in Progress

Paul Johnson
Simon Fraser University

Blaize Reich
Simon Fraser University

Christopher Sauer
University of Oxford

Follow this and additional works at: <http://aisel.aisnet.org/irwitpm2007>

Recommended Citation

Johnson, Paul; Reich, Blaize; and Sauer, Christopher, "Beyond Chaos: Examining IT Project Performance - Research in Progress" (2007). *International Research Workshop on IT Project Management 2007*. 3.
<http://aisel.aisnet.org/irwitpm2007/3>

This material is brought to you by the International Research Workshop on IT Project Management (IRWITPM) at AIS Electronic Library (AISeL). It has been accepted for inclusion in International Research Workshop on IT Project Management 2007 by an authorized administrator of AIS Electronic Library (AISeL). For more information, please contact elibrary@aisnet.org.

Beyond Chaos: Examining IT Project Performance

Andrew C. Gemino*
Simon Fraser University
gemino@sfu.ca

Blaize Horner Reich
Simon Fraser University
BReich@sfu.ca

Chris Sauer
Saïd Business School
University of Oxford
chris.sauer@sbs.ox.ac.uk

ABSTRACT

Management information systems researchers have not yet developed a generally accepted method to measure project performance. The performance measures developed by a consulting firm, the Standish Group, remain the most frequently cited indicators of IT project performance. This paper examines three classes of reason for observed differences in IT project performance including: the subjects and projects in the sample, the data collection method and the analysis method used. Summaries for differences across published methods are provided along with discussion of the pros and cons related to each measurement method. Our comparison suggests that both fast and simple methods and more rigorous methods can provide value. However, to achieve the ability to compare across studies, we should be clear as to which methods are used for a given research study.

Keywords

IT Project Management, Project Management Performance

INTRODUCTION

Although information technology (IT) projects are an important mechanism for delivering value from the IT function, researchers have not yet developed a generally accepted method to measure project performance. Therefore, there is little basis of comparison from study to study. The performance measures developed by a consulting firm, the Standish Group (www.standishgroup.com) remain the most frequently cited indicators of IT project performance. The CHAOS chronicles (Standish Group, 1994) has therefore become a touchstone for researchers concerned with information technology (IT) project performance. The oft-quoted success rate of 16% and average budget and schedule overruns of 189% and 222% respectively are numbers that are difficult to ignore and suggest a relatively poor level of performance for many IT projects. The Standish Group, a consulting firm that collected the data, has repeated the sampling over a decade and by 2007 reported rates for successful projects have risen to 35%, failure rates have lowered to 19% and projects are estimated to have an average cost overrun of approximately 54% (Rubinstein, 2007). Our observations, in interviewing project managers and sponsors, largely agree with the improvement that IT project performance has been making, but it remains unclear what the “real” levels of performance are.

Despite the positive trend, a growing number of researchers have voiced skepticism regarding the original CHAOS numbers and the methods and analysis techniques used to create the initial document and subsequent reports (Glass, 2006, Jørgensen and Moløkken-Østfold, 2006). The reasons for the skepticism have included: 1) a lack of clarity in the sampling and analysis method, 2) a lack of clarity on the nature of respondents and whether they are considering one or many projects,

Our research group has analyzed results from two separate surveys, one in the UK and one in the US focused on IT project performance. This paper focuses on results from the UK study only. The results from the UK survey (Sauer et. al, 2007) suggest that IT projects perform significantly better than is popularly believed. In a reversal of reports from the Standish Group 2006 findings that approximately 35% of information technology projects are successful and 65% are challenged or abandoned (Rubenstein, 2007), we have found 67% of projects delivered close to budget, schedule and scope expectations and only 33% were significantly challenged or abandoned. We have found similar numbers in the US study. However, these results are not presented in this paper.

There are two classes of reason why such differences might be found: methodological differences and substantive differences. While methodological differences remain across studies, we can never be sure whether the apparent performance differences are an artifact of the methods. Our purpose in this study is to examine the methodological differences between our own work and that of Standish as a means to help articulate the critical issues that need to be resolved if cross-study comparability is to be achieved.

The contribution of this research-in-progress paper is

1. to initiate a discussion about the pros and cons of the metrics and procedures currently used for measuring IT project performance
2. to identify critical issues that need to be resolved if the research community is to accurately and consistently measure IT project performance

RESEARCH OBJECTIVES AND QUESTIONS

In this research in progress paper, we are interested in considering methodological reasons why such differences in IT project performance might be observed. We examine three classes of reasons including: 1) the subjects and projects in the sample, 2) the performance data measurement method, and 3) the analysis method. In the presentation at the workshop, we will combine the results from both the UK and US data to present a more complete view of issues arising from the measurement of performance in IT projects. The three classes noted above are discussed individually below. For each, we identify the critical issues, and discuss the pros and cons of the different resolutions adopted in the comparison studies.

Subjects and Projects in the Sample

The choice of sample frame is critical to ensure representativeness of the sample to the population unit of analysis (Pinsonneault and Kraemer 1993). An important consideration in any differences observed in performance levels are the types of projects considered and the participants who provided the data. Twelve years of data collection have provided Standish with over 50,000 completed IT projects (Hartmann, 2006). The 50,000 projects divided by seven biennial reports equals roughly 7,000 projects per report or approximately 3,500 projects per year. The 1994 CHAOS report (“Standish Group, 1994”, p. 2) suggests that the survey targets IT executive managers and focuses on IT projects across a wide variety of industries in small medium and large organizations:

“...respondents were IT executive managers. The sample included large, medium, and small companies across major industry segments, e.g., banking, securities, manufacturing, retail, wholesale, health care, insurance, services, and local, state, and federal organizations. The total sample size was 365 respondents and represented 8,380 projects.”

Since the survey contains fewer respondents than projects, it is assumed that each respondent is asked to comment on one or more projects that their organization is involved in. In the 1994 report, this averaged over 22 projects per respondent. The original report (“Standish Group 1994”) did not provide demographics of respondents. However, it can be assumed that the executive managers participating in the study would have had a high level of experience and knowledge of the IT industry.

In a more recent description of the data collection process, the Standish Group indicated the organization had altered their collection procedure somewhat noting (Hartmann, 2006):

“Now, we invite people to participate in our research using our SURF database, and we have certain entry criteria. Participants must have:

- *access to certain project data,*
- *must already be running applications,*
- *must be running particular platforms.*

The database currently has around 3000 active members.”

So, while the Standish Group data remains targeted at a broad group of industries, covering a large number of projects and a large number of participants, the principal changes relate to the data collection methods, survey participants, their level in the organization and the data collected from them. The fact of such changes across the 12 year period suggests that readers should remain cautious regarding the level of comparability through time.

The UK questionnaire that we report on in this study focused on IT project managers and asked these managers to comment on “the last completed (or abandoned) project that they had worked on”. Each of the respondents therefore provided information on only one recently completed project. Our rationale was that focusing on project managers and specifically on the single, most recently completed project provided the best opportunity for collecting information of project related variances. This approach is supported by Hufnagel and Conca (1994), who caution researchers to ensure that respondents have cognitive access to all the details of their past experience that is being investigated.

Participants in the UK study were selected from registered readers of Computer Weekly, a popular UK-based, weekly newspaper for IT professionals. An initial email along with a follow-up request to participate was sent to readers registered as

project managers on the Computer Weekly site. The survey data was undertaken between October 2002 and January 2003 and collected using web-based forms. A total of 804 participants provided responses to the survey. Of these, only 418 provided full information for budget, schedule and scope variances. A further 6 projects of the 418 projects were excluded due to either an extremely large budget (over 5 standard deviations from the mean), extremely long project (over 5 years) or an anomaly such as large budget and effort, but short duration. This left 412 responses for analysis. The project managers used in the UK sample were highly experienced IT project managers with an average of approximately 17 years in the IT industry and over 9 years as a project manager.

The original Standish Group report did not provide demographics of projects so it is difficult to compare projects between studies. The projects carried out by our 412 experienced UK project managers overshot budget by about 18%, schedule by 21% and under-delivered on scope by 7%. The Standish Group (Hartmann, 2006) reported project figures indicating an average budget overrun of 43%, 82% schedule overshoot and 48% under-delivery of scope. Others have found average budget overruns in the order of 33% (Jenkins et al 1984; Phan et al 1988). By comparison, our figures indicated lower variances across the data than reported elsewhere.

Table 1: Comparison of Participants and Projects in Standish Group 1994 and 2004 UK Sample

Study Criterion	Chaos Report (Standish Group, (1994)	UK Study (Sauer, Gemino and Reich, 2007)
Participants Organizational Role	Executive Managers	Project Managers
Projects considered	Information technology	Information technology
Number of projects reported by participants	Multiple, completed during year	Single, last completed/abandoned project
Number of projects in sample	8,380	412
Average Budget, Median Budget	Not Reported	£9,408,000, £500,000
Average Effort	Not Reported	201.5 person months
Average Duration	Not Reported	12.4 months

In summary, significant differences exist in both the participants used in the study and the projects being considered in the sample. These differences are summarized in Table 1. Note that our UK study is contrasted with the original Standish report because we had greater access to the 1994 data sampling methods. As noted, the sampling methods Standish currently uses may have changed so no claims are made in regard to a more current comparison. Table 2 summarizes the methodological differences and their pros and con's. It is likely that the differences in study protocols contribute to at least some of the differences in observed success/failure rates.

Table 2: Comparison of Sampling Differences Between Standish Group 1994 and 2004 UK Sample

Issue	Standish method	UK Method
Type of respondent	No defined restrictions. Permits wide range of perspectives on performance and makes a larger sample easier to achieve. But no control as to the validity of their reporting of performance as they may not have access to accurate data.	Restricted to project managers and those with direct responsibility for project managers. Increases the probability that they will have access to accurate performance data. But increases the difficulty of securing large samples and the respondents may be positively biased.
Sample selection	Anyone within the sample is encouraged to enter data. Increases sample size. But, implies risk of selection bias.	Anyone within the sample is encouraged to enter data. Increases sample size. But, implies risk of selection bias.
Sampling population	Members of a network led by the Standish consulting group. Broad, cross-industry and, more recently, international population. But risk that membership of the network implies that respondents and their organizations hold a biased view of projects as a problem.	Registered project managers (and their managers) on the Computer Weekly web-site. Broad, cross industry sample. Effectively limited to a single country. Risk that registration on the web-site implies an active interest in learning and self-development which may bias responses.
Unit of analysis for data collection	One or more projects per respondent. Permits easier collection of large volumes of data. But, lack of clarity as to recency of	One project per respondent. The most recent project for which they have been responsible. Ensures no individual unduly biases the results.

	projects and respondent’s relationship to each project.	Increases the probability that the data reported are accurate and recent. But, risk of positive bias.
--	---	---

THE PERFORMANCE DATA MEASUREMENT METHOD

The Standish Group has indicated it has information for over 50,000 IT projects (Hartmann, 2006) that are collected and stored in their SURF database. A variety of data is collected in the Standish questionnaire including industry, company revenues, project type as well as budget, schedule and final deliverable variances (<http://www.standishgroup.com/gr/surf/>). We define variances here as the percentage of actual to original performance. For example a 10% budget variance would suggest that the actual project came in at 110% of the original budget estimate. Note that scope variances operate in the reverse direction so that a scope variance of 110% indicates that more scope was delivered than was initially required.

We have chosen to focus on the variance estimates and how they are collected because they are central for categorizing projects into successful/challenged/failed categories. The Standish Group collection method groups project variances into discrete “buckets”. The endpoints for these buckets are shown below in Table 3.

Table 3: Categories for Performance Variances

Bucket Minimum Ratio of Actual/Original	Bucket Maximum Ratio of Actual/Original	Theoretical Midpoint
Less than 100	100	100
101	120	110
121	150	135
151	200	175
201	500	350
501	Greater than 500	500

When data is categorized into buckets, it is not possible to directly calculate the average for a sample of items. To create, for example, an average budget overrun you need to make an assumption about how the data is distributed within the bucket. The common assumption is that the data within any one “bucket” is distributed uniformly (any number within the category is equally likely). If this is the case, then each item within a bucket can be assigned the value of the theoretical midpoint. These midpoints are shown in Table 3. To calculate the average cost overrun for the entire sample, you multiply the number of items (often referred to as the frequency) in each bucket by the theoretical midpoint of that bucket, add up the result for all of the buckets and then divide by the sample size. This is a well accepted practice for averaging categorical data.

In contrast to the Standish Group collection method, our UK survey collected information about project performance variances using a continuous variable. The questions took the following general form:

“This questionnaire asks about your most recent project/programme prior to your current one. The project/programme may have been completed or terminated prior to completion. The project may have been for a previous employer.

The project/programme cost was:

exactly as initially budgeted
above the initial budget by *** %
below the initial budget by *** %”

Participants were not constrained by categories and could provide variances such as 85%, 123%, etc. We used the same concepts of variance as in Standish (ratio of actual to planned performance). The advantage of continuous data is that the average performance can be calculated without assumptions about distributions within buckets. The continuous data also let us go one step further. Since our data was continuous, it is possible for us to replicate the results using the Standish categories. We can therefore directly compare the averages created using continuous data and the averages created using categories. This comparison is shown in Table 4 below.

Table 4: Comparing Continuous and Categorical Performance Measures

Variance	(A) Continuous Measure	(B) Categorical Measure (using intervals from Standish)	(C) = (B) – (A) Difference between Measures	(D) = (C)/(A) Percent Difference
Average Budget Variance	12.87%	15.76%	+ 2.89	+ 22.5%
Average Schedule Variance	20.13%	22.63%	+2.50	+ 12.4%

When data is categorized into buckets, it is not possible to directly calculate the average for a sample of items. To create, for example, an average budget overrun you need to make an assumption about how the data is distributed within the bucket. The common assumption is that the data within any one “bucket” is distributed uniformly (any number within the category is equally likely). If this is the case, then each item within a bucket can be assigned the value of the theoretical midpoint. These midpoints are shown in Table 3. To calculate the average cost overrun for the entire sample, you multiply the number of items (often referred to as the frequency) in each bucket by the theoretical midpoint of that bucket, add up the result for all of the buckets and then divide by the sample size. This is a well accepted practice for averaging categorical data.

In contrast to the Standish Group collection method, our UK survey collected information about project performance variances using a continuous variable. The questions took the following general form:

“This questionnaire asks about your most recent project/programme prior to your current one. The project/programme may have been completed or terminated prior to completion. The project may have been for a previous employer.

The project/programme cost was:
exactly as initially budgeted
*above the initial budget by *** %*
*below the initial budget by *** %”*

Participants were not constrained by categories and could provide variances such as 85%, 123%, etc. We used the same concepts of variance as in Standish (ratio of actual to planned performance). The advantage of continuous data is that the average performance can be calculated without assumptions about distributions within buckets. The continuous data also let us go one step further. Since our data was continuous, it is possible for us to replicate the results using the Standish categories. We can therefore directly compare the averages created using continuous data and the averages created using categories. This comparison is shown in Table 5 below.

Table 5: Comparing Continuous and Categorical Performance Measures

Variance	(A) Continuous Measure	(B) Categorical Measure (using intervals from Standish)	(C) = (B) – (A) Difference between Measures	(D) = (C)/(A) Percent Difference
Average Budget Variance	12.87%	15.76%	+ 2.89	+ 22.5%
Average Schedule Variance	20.13%	22.63%	+2.50	+ 12.4%

Table 6: Comparison of Data Measurement Differences Between Standish Group 1994 and 2004 UK Sample

Issue	Standish method	UK Study Method
Continuous categorical measures	Categorical Advantage that it simplifies answering for the respondent. Disadvantages that the categories	Continuous Advantage that it is specific and permits accurate statistics. Disadvantage that it requires

	are arbitrary; can result in biased statistics if the distribution within the categories is not even; some of the categories cover a very large range.	precise knowledge of performance from the respondent.
Ratio measure	Actual/planned performance against initial targets Advantage that it provides a common metric by which to compare projects. Disadvantage that when targets change, there is no check that respondents are answering against the original targets are subsequently varied targets.	Actual/planned performance against initial targets Advantage that it provides a common metric by which to compare projects. Disadvantage that when targets change, there is no check that respondents are answering against the original targets are subsequently varied targets.

The results in Table 5 show that the assumption of uniform distribution within performance variance categories results in an upward bias in the reported variance by approximately 2.5%. This translates into a 22.5% overestimate in the reported average budget variance in the UK data and a 12.4% overestimate of the average schedule variance when compared with the same data calculated using a continuous measure. It is important to note that this difference is unlikely to be reduced with increased sample size as the sample of over 400 projects suggests categories are not uniformly distributed but rather have a significant skew towards the minimum point of the category. This suggests that if the Standish Group does not use a correction for the categorical data, the estimates Standish provides are likely overestimating the actual performance variances in the range of 10 to 20%. Table 6 summarizes the methodological issue and the pros and cons of the different approaches of capturing project variance information.

THE ANALYSIS METHOD

A third factor that may cause the observed differences in performance is the analysis method. In the original report (Standish Report, 1994, p. 2) the three categories of performance (failed, challenged and success) were created using the following project resolution types:

“For purposes of the study, projects were classified into three resolution types:

- *Resolution Type 1, or project success: The project is completed on-time and on-budget with all features and functions as initially specified.*
- *Resolution Type 2, or project challenged: The project is completed and operational but over-budget, over the time estimate, and offers fewer features and functions than originally specified.*
- *Resolution Type 3, or project impaired: The project is cancelled at some point during the development cycle.”*

This classification suggests that a successful project is on time on budget and delivers expected scope. The “ands” are important here as they suggest that all three items (and no other combination) must be achieved for success. In addition there is no room for a contingency on initial estimates. A project is challenged if it comes in as little as 1% over budget even though time and requirements expectations are met.

Continuous data from the UK sample allows us to define the three resolution types as defined above. These results are shown in column B in Table 7 using then Standish Group criteria. The results indicate that the UK sample would have provided a 17.4% success, 16.6% abandoned, and 66% challenged as defined by the Standish Group categories. These numbers can be compared with the results from the Standish Group 2004 survey (Hartmann, 2006) that provided estimates of 29% success, 18% failure and 53% challenged provided in Column A in Table 6. The results in Columns A and B show that even when using the same criteria for categorizing projects, the sample results do not correspond closely.

The results in Table 7 also show a failure rate in the UK study that is almost half the rate of the Standish Group 2004 study. A similar abandonment rate of projects in the US study (9%) suggests a significant discrepancy between the two sets of results. There are no obvious reasons for this discrepancy.

Table 7: Project Type Percentages across Different Contingency Allowances in UK Study

Project Type	Column A: Standish 2004	Column B: UK Study (2004) 100% cutoff	Column C: UK Study 5% contingency	Column D: UK Study 10% contingency	Column E UK Study 15% contingency
--------------	----------------------------	--	---	--	---

Success	29%	17.4%	25.1%	45.2 %	56.1%
Challenged	53%	74.4%	65.7%	45.6%	34.7%
Failure	18%	9.2%	9.2%	9.2%	9.2%

In addition to columns A and B, two other columns have been provided in Table 7. These columns indicate the percentage of project types that would result if the 100% on budget, time and scope constraints were relaxed. Column C reports on the percentages if a 5% contingency on top of original estimates was provided before challenged projects were determined. Likewise, Column D shows the percentages of projects types if the 100% constraints were relaxed to a contingency of 10% and column E to 15%. The large differences observed across columns B, C, D and E suggest the danger of drawing an arbitrary line to determine success and failure.

An alternative method for categorizing IT Project types is using clustering methods as demonstrated in Sauer et. al (2007). Clustering techniques are a data driven approach that uses algorithms to group projects with similar responses together in a “cluster” defined by a centroid in n dimensional space. The technique is data driven because it does not require the investigator to predetermine “reasonable” cut off points for inclusion in one or more clusters. Instead, the algorithm works to separate groups based on the data provided. Sauer et al (2007) used a used a non-hierarchical clustering method in SPSS version 14.0 to combine similar projects in the UK study based on the three performance variances: schedule, budget and scope. The results provided five distinct project types with centroids defined as follows in Table 8

Table 8: Project Types as Defined by Cluster Analysis (Sauer et. al, 2007)

Performance Variance	Type 1: Abandoned Projects n=38	Type 2: Budget Challenged n=21	Type 3: Schedule Challenged n=74	Type 4: Good Performers n=249	Type 5: Star Performers n=30
Performance Variances (Actual as % of Originally Planned) – 100%					
Schedule	N/A	+34%	+82%	+2%	+2%
Budget	N/A	+127%	+16%	+7%	-24%
Scope	N/A	-12%	-16%	-7%	+15%

It is important to note that the clusters defined across different samples will likely differ. This is because the clustering analysis method is data driven and hence categorizations will change as the data changes. In regards to classification, Sauer et al (2007) went on to suggest that types 2 and 3 of IT projects were clearly underperforming, while types 4 and 5 were performing and, in the case of type 5, over performing. It was reasonable therefore to combine types 2 and 3 and refer to them as “challenged” and types 4 and 5 and refer to them as successes’. Using this categorization provided an abandoned rate of 9%, a challenged rate of 23.5 and a success rate of 67.7%.

Table 9: Comparison of Performance Categorization Differences Between Standish Group 1994 and 2004 UK Sample

Issue	Standish method	UK Study Method
Basis for performance classification	Simple, a priori classification into failed, challenged, successful. Advantage that it’s intuitively easy to understand and remains stable from study to study. Disadvantages are that the challenged category encompasses wide variation; without more detailed analysis it is not possible to identify improvement from larger performance variances to smaller performance variances.	Data driven classification evokes classes of similar performance. Advantage that it reflects what the data tells us. Disadvantage that different data sets across studies may provide incomparable categories.
Tolerance of variance	Challenged category can be defined by even a	Good performance defined by a cluster

<p>in performance classifications</p>	<p>small variance against a single target. Advantage that it reflects a widespread view that projects should meet their targets and any shortfall is underperformance. Disadvantage that it disregards the complexity, uncertainty and volatility of most IT projects and is at odds with how small variances would be judged in most organizational situations.</p>	<p>of projects that get relatively close to their targets. Advantage that moderate tolerance of variance reflects organizational reality in many cases. Disadvantage that it appears to condone small levels of underperformance.</p>
---------------------------------------	--	---

It is clear that the clustering method used in Sauer et al represents a different approach than that taken by the Standish Group in at least the 1994 study. Table 9 summarizes these differences. Clearly, other methods can be proposed for categorizing IT project performance. There is no “best” way to develop these categories. Each method will have its pros and cons. What is important in developing further research in IT project performance is that the methods used to categorize performance be clearly specified and appropriate for the research questions being proposed. Clear specification will enable the comparison of numbers across studies, and provide an improved foundation for research of factors affecting IT project performance.

CONCLUSIONS AND FUTURE WORK

We began our discussion by noting that the management information systems research has no cumulative tradition in the measurement of information technology (IT) project performance. A focus was placed on two studies, one from Standish and one from our own work in the UK, which considered IT performance. Both studies used measures based on variances on initial project targets, yet the studies reported significantly different pictures of overall project performance. The purpose of this paper was to examine the methodological differences between our own work and that of Standish in an effort to better articulate the issues that need to be resolved in developing measures of IT project performance that have potential for cross-study comparability.

Our analysis focused on three areas: the subjects and projects in the sample, the data collection method and the analysis method used. The summary provided in Table 2 of subjects and projects used in the samples suggests that choices in type of respondent, sample selection, sampling population and unit of analysis for data collection include opportunities for greater clarity but reveal risks of respondent bias. The summary provided in Table 5 of data collection methods suggests limitations to using categorical data and a potential for bias in categories that are assumed to be uniformly distributed. The summary provided in Table 8 of analysis methods suggests there are drawbacks in drawing arbitrary cutoffs for defining performance categories. Data driven methods for categorizing performance, such as clustering, also pose problems in regard to comparing across studies.

In summary, our analysis to date shows that a case can be made for quick and dirty methods as well as a case for more rigorous methods. The choice of methods depends on the question of interest and the resources available. What remains important when considering the potential for cumulative research is the need for clarity in regards to the choices being made on the following factors:

- 1) type of respondent
- 2) sample selection
- 3) sampling population
- 4) unit of analysis
- 5) data collection method, and
- 6) performance categorization technique

Our analysis in this paper has demonstrated that reported performance measures can vary widely and that they depend on the choices being made on these categories.

One of the limitations of this work is the sample of informants represented in our UK data. They were drawn from the population of registered readers of Computer Weekly, and may not represent the entire population of project managers adequately, thereby introducing sampling errors (Hufnagel and Conca, 1994) and decreasing our ability to make inferences (Sivo et al. 2006).

In future work, this paper will expand to consider results from a US study completed in 2005. We are also searching the literature for other studies on IT project performance which will allow for greater comparison between techniques. These additional studies will provide a wider sample base from which to draw conclusions regarding performance levels in IT projects and may eventually allow for a meta-analysis of results to determine appropriate size and number of performance categories.

Another area of work is to further consider the construct of project variance. Our discussions have suggested that establishing “initial” and “final” project targets can be difficult for project managers. This puts into question the notion of measuring performance through projects variances alone. We have collected some initial data on what project managers interpret as “initial” and final targets which may be presented at the workshop.

Finally we end with a short discussion of the nature of project performance. Our interviews and discussions have suggested that project performance should be considered a wider construct than simply hitting targets. MIS researchers have already recognized the need to focus on both process (target) performance and product performance (Barki. et. al., 2001; Nidumolu, 1995). Interviews have revealed that practitioners are concentrating more on notions of business value and client satisfaction than simply being on time and on budget. The notion of product performance, defined here as how well the information technology addresses the needs of the company, is therefore an issue that requires significant additional research. In addition, the relationship between process/target performance and product performance also needs to be considered. These constructs have been treated largely as independent constructs in the past. It is not difficult to argue that some relationship may exist between the two performance constructs. The increasing focus on business value and recognizing the dependence between process and product performance suggests that our notions of how to measure IT project performance are likely to change as they receive increased research attention.

ACKNOWLEDGEMENTS

This research has been supported by a grant from the Social Sciences and Humanities Research Council (SSHRC) and the Natural Sciences and Engineering Research Council (NSERC) of Canada.

REFERENCES

1. Barki, H., Rivard, S., and Talbot, J. “An integrative contingency model of software project risk management”. *Journal of Management Information Systems*, 17, 4 (2001), 37-69.
2. Gemino, A., Reich, B. and Sauer, C. “A Temporal Model of Information Technology Project Management”, accepted for publication and forthcoming in *Journal of MIS*.
3. Glass, Robert, “The Standish report: does it really describe a software crisis?”, *Communications of the ACM*, Volume 49, Issue 8, August 2006, p. 15-16.
4. Hufnagel, E.M. & C. Conca, "User Response Data: The Potential for Error and Biases," *Information Systems Research*, 5, 1 (1994), pp. 48-73.
5. Jørgensen, M. and Moløkken-Østfold, K. J. “How Large Are Software Cost Overruns? Critical Comments on the Standish Group’s CHAOS Reports”, *Information and Software Technology* 48(4): 2006, p.297-301,
6. Hartmann, D., “Interview: Jim Johnson of the Standish Group”, *Infoqueue*, Aug 25, 2006, <http://www.infoq.com/articles/Interview-Johnson-Standish-CHAOS> accessed August 22, 2007.
7. Nidumolu, S. The effect of coordination and uncertainty on software project performance: Residual performance risk as an intervening variable. *Information Systems Research*, 6, 3 (1995), 191-219.
8. Pinsonneault, A. & K. Kraemer, "Survey Research in Management Information Systems: An Assessment," *Journal of Management Information Systems*, 10, 2 (1993), pp. 75-105.
9. Rubinstein, David. “Standish Group Report: There’s Less Development Chaos Today”, *Software Development Times*, March 1, 2007, <http://www.sdtimes.com/article/story-20070301-01.html>, accessed August 22, 2007.
10. Sauer, C., Gemino, A. and Reich, B., “IT Project Performance: The Impact of Size and Volatility”, *Communications of the ACM*, Vol. 50, Iss. 10, November, 2007, p. 79-84.
11. Sivo, S.A., Saunders, C., Chang, Q., Jiang, J.J., “How Low Should You Go? Low Response Rates and the Validity of Inference in IS Questionnaire Research,” *Journal of the Association for Information Systems*, (7:6), Article 17, July 2006.

12. "Standish Group 1994". "The CHAOS report, 1994, The Standish Group ,
http://www.standishgroup.com/sample_research/chaos_1994_1.php , accessed August 22, 2007.