Marquette University
e-Publications@Marquette

Management Faculty Research and Publications   Business Administration, College of

4-1-2006

An Empirical Study of Operational Performance Parity Following Enterprise System Deployment

Mark Cotteleer
Marquette University, mark.cotteleer@marquette.edu

An Empirical Study of Operational Performance Parity Following Enterprise System Deployment

Mark J. Cotteleer, Marquette University, P.O. Box 1881, Milwaukee, Wisconsin 53201, USA

mark.cotteleer@marquette.edu

Abstract: *This paper presents an empirical investigation into whether the implementation of packaged Enterprise Systems (ES) leads to parity in operational performance. Performance change and parity in operational performance are investigated in three geographically defined operating regions of a single firm. Order lead time, the elapsed time between receipt of an order and shipment to a customer, is used as a measure of operational performance. A single ES installation was deployed across all regions of the subject firm's operations.*

Findings illustrate parity as an immediate consequence of ES deployment. However, differences in rates of performance improvement following deployment eventually result in significant (albeit smaller than pre-deployment) performance differences. An additional consequence of deployment seems to be an increased synchronization of performance across the formerly independent regions.

1. Introduction

The productivity paradox of information technology is largely resolved. Properly implemented, information technology can effectively be used to pursue improved performance. The efficacy of information technology has been demonstrated at the level of the economy (Dewan and Kraemer 1998), industry (Weill 1992), and firm (McAfee 2002) using measures that include GNP, productivity, consumer surplus, and operational performance. But as issues of general effectiveness recede, new questions emerge. One of these is whether parity in operational performance is observed across sites that implement enterprise-level packaged software systems.

The importance of this question derives from the increasing use of such systems. In 2001, firms invested more than $47 billion in Enterprise Systems (ES) applications (AMR 2001) and billions of additional dollars in the services and computer hardware required to implement and maintain them (Mabert, Soni, and Venkataramanan 2000; Willcocks and Sykes 2000). SAP, the largest ES provider, reported 2004 software and consulting revenues in excess of €4.3 billion with maintenance, training, and other services contributing another €3.1 billion. Overall the firm

claims more than 12 million users across 88,700 installations at 26,150 customers in 120 countries. Oracle, SAP's main ES competitor, reported overall revenue in excess of $10 billion, with business applications and consulting accounting for more than $2 billion of that total.[1] Major information technology categories falling within the ES domain include Enterprise Resource Planning (ERP), Customer Relationship Management (CRM), and Advanced Planning and Scheduling (APS) systems.

Both proponents and critics of ES investment maintain that these technologies enable standardization of business process execution. Proponents view this as a positive outcome likely to improve processes and promote consistently high operational performance throughout firms. Critics hold that business process standardization will occasion competitive parity across firms and is thus non-strategic. Both sides of the debate, however, lack a clear understanding of whether parity in operational performance is actually a consequence of ES implementation. This paper contributes to the development of that understanding.

This paper investigates operational performance patterns resulting from ES software package implementation in the context of three geographically defined operating regions of a single firm. "Parity", defined literally as "the quality or state of being equal or equivalent" (Webster's 1977), is assessed from both statistical and managerial (financial) perspectives. Operational performance is measured in terms of order lead time, the elapsed time between receipt of an order and shipment to the customer. The analysis was performed in the context of a computer peripherals manufacturer for which order lead time was a key determinant of competitive success. The firm deployed a single ES installation, simultaneously in all regions of its global operations.

Findings demonstrate parity in operational performance immediately following deployment of the new ES. That is, performance across the regions becomes very similar from a managerial perspective and indistinct from a statistical perspective. Parity was, however, not a long-run consequence of the ES deployment. Differences began to emerge in the months following implementation such that performance across the three operating regions by the end of the study period was managerially and statistically distinct, albeit more similar than during the pre-deployment period.

Findings suggest that parity in operational performance may arise initially as a function of "system effects" related to the specific functionality, process assumptions, and algorithms embedded in the ES itself. However, over the longer term, "implementation effects" related to the unique context and user choices associated with the individual regions potentially contributed to emergent performance differences.

An additional finding of the study relates to the apparent synchronization of regional lead time performance in the period following ES deployment. Whereas operating performance across regions had been uncorrelated in the pre-deployment period, post-deployment performance was observed to be highly correlated. This phenomenon is potentially explained using the lenses of extant ES literature and the effects of inventory pooling.

The rest of the paper proceeds as follows. Section 2 presents an overview of the literature relevant to the discussion of information technology-enabled improvement, ES deployment, and the potential for subsequent parity in operational performance. Section 3 relates the context of the study and describes the data that were used. In Section 4, the data are analyzed. Section 5 provides a discussion of the results and some suggestions for future research.

2. Enterprise Systems, Performance, and Parity

Traditional research on the value of information technology has established that firms generally benefit from IT investments (Brynjolfsson and Hitt 1998; Dewan and Kraemer 1998). More recent studies have extended this understanding, arguing that ES adoption has the potential to improve productivity, stock market premiums, and operational performance. Findings from Hitt and Wu's (2002) longitudinal study comparing ERP adopters and non-adopters indicate that ERP implementation is associated with improvements in productivity, and in Tobin's q, a measure of surplus stock market returns. These findings are consistent with longitudinal firm-level research that finds associated improvements in operational performance (McAfee 2002; Cotteleer and Bendoly 2006).

Despite these findings, it remains unclear whether ES confers long-term advantage. ES implementation, though widespread, is not yet ubiquitous. Moreover, many firms' ES are not fully deployed. The surplus stock market returns that have been identified might be transitory effects that represent a sort of early-mover rather than sustainable competitive advantage.

2.1. Enterprise Systems as a Driver of Performance Convergence

Enterprise Systems are software packages designed to integrate information and processes among the functional areas of a business, including: Finance and Accounting, Order Management, Manufacturing, Inventory Management, and Procurement (Kumar and Van Hillegersbert 2000). This paper reports operational performance outcomes of an ES implementation and examines whether the parity sometimes assumed to result from the

business process standardization imposed by ES actually occurs. Standardization is frequently cited as a goal of managers who lead firms' ES implementations (Austin and Cotteleer 1999).

Advocates of process standardization argue that it can be a major component of a firm's operations strategy, promoting the development of process knowledge (Hayes, Wheelwright, and Clark 1988), facilitating cross-site learning (Pisano and Rossi 2001), and enhancing the ability to measure, compare, and improve (Garvin 1998). Critics view standardization as an enabler of process and performance duplication, affecting the firm's overall ability to differentiate it from competitors. Their view holds that this duplication fosters the migration of competitive knowledge across firms, particularly when process knowledge is embedded in technology, which has been identified as an enabler of the transfer of organizational learning (Epple, Argote, and Murphy 1996).

Business process standardization in the context of ES deployment is occasioned by the process templates that ES providers embed in packaged applications (Scheer and Habermann 2000), and through the use of the standardized decision-support tools (e.g., demand planning and forecasting, inventory management) used by managers to improve process performance (Razi and Tarn 2003).

The functionality embedded in ES incorporates vendor-determined assumptions about the way business should function (Austin and Nolan 1998; Davenport 2000). Enterprise Systems have been characterized as being particularly effective drivers of standardization because the applications themselves are viewed as inflexible and because employees may have little choice but to use these systems in fulfillment of their organizational roles (Davenport 2000). Furthermore, circumvention of ES technology and process may be difficult (Bendoly and Cotteleer 2004). The standardizing effect of ES, combined with their broad availability gives rise to the view that it may be a non-strategic investment (Davenport 2000; Carr 2004).

The notion that ES implementation leads to a standardization in business process execution has been borne out in field-based research (Davenport 2000) and in an emerging theory of ES implementation. Bendoly (2001) and Bendoly and Kaefer (2004) argue for the existence and influence of "system" effects on post-implementation performance. System effects are those that derive specifically from the attributes and functionality of the ES itself. These may include both the "process maps" that drive the manner in which work and information flow through the business, as well as the fundamental algorithms upon which business decisions are made. For example, system functionality may direct the user through the steps required to allocate inventory against an order. It may also impose a decision-making approach on the user (e.g., what is an allowable configuration for the product? From which stocking location should

inventory be allocated, and with what priority?). Where such effects are induced by the implementation of standardized technology applications, they may be viewed as imitable and potentially parity causing.

Research in other operational contexts mitigates such worries by raising doubts about the extent to which an ES or any other technology can occasion performance parity. In contexts as diverse as manufacturing plants (Adler and Clark 1991), and hospital operating rooms (Pisano, Bohmer, and Edmondson 2001) studies have demonstrated that the implementation of similar technologies can not be counted on to induce similar performance.

The literature on ES implementation refers to differences in the application of technology in the context of unique firms as "implementation effects" (Bendoly 2001; Bendoly and Kaefer 2004). Implementation effects reflect the unique decisions and business process changes that underpin ES deployment in a specific context. Implementation effects may result from myriad factors including differences in management skill and in task-technology fit. Implementation effects may also emerge over time as users learn about and adapt the system. Choices may be made to deviate either from expected information or work flow, or from settings and/or the recommendations of decision-support algorithms. The presence of implementation effects is expected to lead to differences in the realized benefits of ES deployment (Bendoly and Kaefer 2004).

The presence of countervailing "system" and "implementation" effects raises questions about which will dominate in the context of an ES implementation. Studies suggest that the answers may depend, in part, on when the question is asked. Bendoly and Cotteleer (2004) argue for the possibility of a "resonant dissonance" that affects post-deployment performance over an extended period. ES users are expected to accept systems and process as implemented, during the short term, until they can identify ways to circumvent the system so as to make business processes more closely match perceived local requirements.

The decision to adopt an ES as implemented early, but to circumvent that system in order to alleviate perceived task-technology mismatches in the long-run, has direct implications for expectations relating to the parity-inducing power of ES and the influence of "system" and "implementation" effects. Adoption of the system as implemented early in the post-deployment period might be expected to induce parity across implementation sites, as technology attributes and functionality dominate the definition of user behavior. Over time, however, as resonant dissonance leads to longer-term system circumventions, "implementation effects"-local process adaptations in this case-could begin to overwhelm the standardizing influence of the ES and lead to divergence in performance. This leads to the statement of two hypotheses.

H1: *Implementation of ES will induce parity in operational performance (e.g., lead time) across independent operating regions in the immediate, post-deployment period.*

H2: *Parity in operational performance (lead time) across independent operating regions will diminish over time following* ES *systems deployment.*

3. Context and Measurement of Operational Performance

3.1. Organizational Context

This study was conducted at Tristen Corporation, a $4 billion producer of peripheral equipment for servers, desktop PCs, and other computerized equipment. The company served this market through three geographically dispersed manufacturing/distribution centers, North America, Europe, and Asia, each supporting its own network of sales offices. Tristen enjoyed strong, stable relationships with suppliers and consistency in its product mix. Forecast-based production was held in a semi-finished state within each region in anticipation of demand. Products were configured and tested during order fulfillment.

In the late 1990s, Tristen management began to endorse strategic change. The anticipated commoditization of the company's products led Tristen managers to embark on a strategy that increasingly emphasized the delivery "high value-added," which the company defined as incorporating a dual emphasis on cost leadership and customer service. This strategy was a departure from the firm's historical emphasis on technology leadership.

Tristen management identified the implementation of an ES as a key requirement for its strategic objectives. At the time of the decision, the company relied on antiquated, non-integrated, region-level information technology. The ES implementation was Tristen's key strategic initiative during the study period.

Lack of systems integration prevented Tristen from providing standard service levels to multi-national customers. For example, customers often received different order and delivery commitments from Tristen depending on the geography of the sales office that was contacted or the location of the customer facility. "Many of our customers are global," explained one manager. "So, for example, a customer has operations in North America, Europe, and Asia. They expect each operation to be treated the same. Instead, the different regions were getting different levels of service." Tristen's lack of systems integration frustrated customers' supply chain management efforts and by extension the company's attempts to maximize customer satisfaction.

Tristen embarked on the implementation of enterprise systems with the intent not simply to improve service levels, but also to ensure that performance was consistent across the different geographies the company served. Consistent performance was to be achieved through the

rollout of standardized technology and business processes, a marked contrast to the pre-ES period when regional operations were essentially independent.

The company ultimately selected a set of packaged software applications sold by a leading ES vendor. The deployment included functionality for order management, inventory, logistics, and financials and was subject to three strictures: minimize customization; minimize business process change; pursue simultaneous ("big bang") implementation of all processes at all locations.

3.2. Measuring Performance

This paper focuses on operational performance in the order management process, a primary target of ES implementations (Davenport 2000). Typically, ES includes functionality for order generation, planning, production scheduling, inventory management, order picking, and shipping (McAfee 2002; Jacobs and Bendoly 2003). Enterprise applications, moreover, are specifically engineered to link business processes and information in order to improve communication and visibility (Mabert, Soni, and Venkataramanan 2000; Hult, Ketchen, and Nichols 2002; McAfee 2002) and to improve decision-making through the use of OR models and simulations (Razi and Tam 2003). "Available-to-promise" functionality, for example, was intended to afford Tristen visibility into its internal supply network and to facilitate rapid allocation and shipment of inventory.

The measure of interest to this study is "order lead time", a key performance metric for the firm. Order lead time is defined as the elapsed time between the receipt of an order and shipment of the product to the customer. This measure has established roots as an indicator of order management and supply chain performance, having been used in multiple studies of related phenomena (Mabert, Soni, and Venkataramanan 2000; Hult, Ketchen, and Nichols 2002; McAfee 2002). The timeliness of customer service has also been identified as a primary reason for IT investment (Brynjolfsson and Hitt 1998).

Performance parity is defined as "the quality or state of being equal or equivalent" (Webster's 1977). Parity is measured statistically by assessing whether operating regions exhibit significantly different levels of performance. Attention is also given to the managerial implications of differences in performance.

4. Analysis of Operating Data

Lead time data were collected for sales orders processed both before and after ES deployment. The data included information necessary to calculate order lead time as well as

related information that enabled transactions to be sorted by operating region and classified by transaction date, order quantity, and order value. Order-level data facilitates the inclusion of information on order units (i.e., number of units associated with the individual order), order value (i.e., the dollar value of the order) and a measure of business intensity or "busyness" (as represented by the number of orders within a given region per day). Although not thought to be significantly influenced by implementation plans, these controls nevertheless are viewed as potentially having impacts on fulfillment capabilities and thus lead time performance. The variables used in the analysis and their summary statistics are presented in Table 1. Table 2 provides correlations between the variables.

One desirable variable that is missing from the analysis is inventory at the firm and region level. Inventory is frequently used to buffer system shocks and to support operational performance (Hayes, Wheelwright, and Clark 1988). An analysis of quarterly inventory data suggests that Tristen attempted to buffer the shock of ES deployment by increasing inventory levels prior to implementation. However, the overall inventory pattern in the post-deployment period is marked by a rapid reduction in level. Gross and revenue-adjusted inventory fell to just 40 percent of pre-deployment levels in the year following implementation, rapidly depleting any buffer that Tristen built to support lead time performance. By the end of the study period, revenue-adjusted inventory levels remained below historic averages (and well below levels at the time the ES was deployed). It therefore seemed reasonable to conclude that any generalized improvements in lead time at Tristen would not be attributable to increases in held inventory.

Regardless of overall inventory levels, it might be argued that relative differences in regional inventory distributions could contribute to lead time performance differences. Though possible, analytical and field-based findings mitigate concerns here as well.

Tristen sought to equalize service levels across regions and customers through its ES deployment. Managers specifically wanted to eliminate instances where geographically diverse customers experienced different service-levels depending on location. The creation of policies and procedures that consistently disadvantaged specific regions would be inconsistent with that mission. Furthermore, a key objective of the ES implementation was to provide inventory visibility across the enterprise, enabling the allocation of inventory regardless of its location, and evening the playing field across regions.

Nonetheless, poor management could result in improper inventory allocation across regions, regardless of intent. An inaccurate region-level forecast could leave a region at a relative disadvantage during a particular period. However, such inaccuracies alone would not be expected to lead to consistent differences in lead time performance. If unbiased, such

inaccuracies would randomly fall across regions with both positive and negative lead time effects (Nahmias 1997), obfuscating rather than enhancing systematic differences in lead time performance. Although consistent region-level forecast bias could contribute to the appearance of differences in regional lead time performances, no evidence of this condition was detected during field analysis. Furthermore, forecast generation was subject to much greater centralization under the new ES solution. The possibility that consistent forecast bias in one region would go unnoticed seems remote.

The analysis included nearly 113,000 pre- and post-deployment order records covering a period beginning 12 months prior to ES deployment and extending 24 months after ES deployment. Order transactions for the three months prior to and five days following ES deployment are excluded from the data set in order to avoid measurement bias that would have arisen as a function of the way in which outstanding orders were converted to the new ES system. Figure 1 illustrates raw regional lead time performance during the analysis period.

4.1. Cross-Sectional Analyses

Cross-sectional data analysis evaluated differences in regional lead time performance without the imposition of a specific functional form on the data. This evaluation consisted of repeated cross-sectional ANCOVA to assess where differences existed in regional lead times. ANCOVA differs from ANOVA in so far as it facilitates the inclusion of continuous quantitative variables that might be related to the dependent variable (Neter, Wasserman, and Kutner 1990). To perform the ANCOVA, differences in average regional lead times were calculated for each day included in the analysis period. The analysis simultaneously controlled for differences in average order units, order value, and daily order volume.

ANCOVA results presented in Table 3 provide estimated mean differences in regional lead times for four distinct three-month periods: (1) prior to ES implementation; (2) immediately following ES implementation; (3) one year after ES implementation; (4) two years after ES implementation.

A graphical representation of ANCOVA results is presented in Figure 2 which uses the absolute value of the regional lead time differences in order to promote easier visual comparison.

The analyses provide preliminary support for Hypothesis 1. Visual inspection of the raw data presented Figure 1 indicates an apparent collapse of differences in lead time performance following deployment. Data from Table 3 (as depicted in Figure 2) supports this interpretation. Where significant differences in lead time appeared between two out of three region pairs (North

America vs. Europe; Europe vs. Asia) in the pre-deployment period, these differences had all but disappeared in the "At Imp" period immediately following ES deployment.

Preliminary support also exists for Hypothesis 2, which states that performance parity will diminish over time. Table 3 shows that after one year, divergence in regional lead times resulted in a significant performance difference between Europe and Asia. Two year ANCOVA results indicate significant differences in performance across all three operating regions.

Table 4 provides further support for findings related to Hypotheses 1 and 2. ANCOVA results reported in Column (1) of the table indicate a significant "difference of differences" between the pre-deployment and deployment periods for two out of three region pairs (supporting Hypothesis 1). The interpretation is that not only is the 9.47 day difference between North America and Europe performance in the "Pre-Imp" period statistically significant in and of itself (as reported in Table 3), it is also significantly different in magnitude from the 2.49 day difference between same two regions' performance in the "At Imp" period (as reported in Table 4, column 1). Significant "difference of differences" also appear between the deployment period and the two-year post-deployment period (column 3, supporting Hypothesis 2) for two out of three region pairs.

Despite the reemergence of regional differences, post-deployment disparities never rival those encountered during the pre-deployment period. A comparison of regional differences in the pre-deployment period vs. two years post-deployment (column 4 of Table 5) shows the post-deployment gaps to be significantly smaller for two of the three region pairs.

4.2. Longitudinal Analyses

ANCOVA results indicate a collapse in regional differences following ES deployment followed by a gradual increase in those differences. However, they provide little in the way of a dynamic understanding of how those changes unfolded. Longitudinal analyses were undertaken to develop this insight.

Specification of a longitudinal model for describing the operational performance changes in this case was theoretically motivated both by the extant literature on organizational learning curves and by a practical visual interpretation of the data itself.

Learning curve models generally posit that performance following implementation of a new process or technology will improve nonlinearly but monotonically with time or experience as a function of learning-by-doing (see Yelle 1979 for a review of the learning curve literature.). Specifically at Tristen, workers continued use of enterprise systems may be expected to promote improved performance, which in turn motivates future use. Through a reinforcing cycle of use

and performance, users may develop deeper understanding and familiarity with system techniques. With deeper understanding and familiarity come improved efficiency and the opportunity for even more substantial performance improvements.

In addition to this theoretical motivation, visual inspection of the dependent variable (lead time), a tactic recommended by Neter, Wasserman, and Kutner (1990), suggested that non-linear effects might be at work, but argued against strict monotonicity of the improvement function. Given this observation and anticipated learning effects, it seemed prudent not to restrict the analysis in the manner of traditional learning curve models. Instead, a more flexible analytical approach was chosen. Following the examples set in recent process studies by McAfee (2002), Banker, Field, and Sinha (2001), and Badiru (1992), a quadratic time-based term is included in the model. Such a specification provides the flexibility to capture nonlinear trends without precluding the existence of strictly monotonic functions.

Models designed to illustrate learning curve effects also incorporate the same measures of operational scale and activity used in the ANCOVA analyses. Once again, these controls are considered to be exogenous and so no specific hypotheses are proposed. However, it is reasonable to expect that lead times would lengthen on busier days (as orders pile up in the system). The effect of order size and value may depend on the unique character of individual orders and regions. For example, if an order is for many of the same unit then lead time may fall as the region focuses on shipping a "big order", yet realizes economies of scale by processing many of the same unit for shipment. Alternatively, high value orders (controlling for number of units shipped) may represent "higher value added" products which require more processing and configuration prior to shipment, thus increasing lead times. Control variables for order size and value, and daily order volume, enter the model both as main effects and in interaction with regional indicator variables to account for differences in the effect each has on regional lead time performance. The resulting model is presented in Equation 1.

$$
\begin{aligned}
leadtime_i = \alpha &+ \beta_1 daycount_{ij} + \beta_2 ordqty_i \\
&+ \beta_3 ordvalue_i + \beta_4 europe_i + \beta_5 asia_i \\
&+ \beta_6 europe_i * daycount_{ij} + \beta_7 asia_i * daycount_{ij} \\
&+ \beta_8 europe_i * ordqty_i + \beta_9 asia_i * ordqty_i \\
&+ \beta_{10} europe_i * ordvalue_i + \beta_{11} asia_i * ordvalue_i \\
&+ \beta_{12} month_i + \beta_{13} europe_i * month_i \\
&+ \beta_{14} asia_i * month_i + \beta_{15} month2_i + \beta_{16} europe_i * month2_i \\
&+ \beta_{17} asia_i * month2_i + \varepsilon
\end{aligned}
\tag{1}
$$

Where *"i"* represents the customer order and *"j"* the operating region of origin.

Two additional steps are taken to promote the statistical efficiency of the model. The first addresses the presence of skewness in the raw data. The presence of such skewness increases the potential for a Type II error during hypothesis testing (Miller 1986). Box-Cox transformations (Box and Cox 1964) address this issue. In doing so, transformation increases the difficulty of interpretation of model results. To ease this interpretation, "untransformed" estimates will be provided in the presentation of results.

The second statistical step addresses a possible lack of independence among order lead times. The nature of the data (i.e., orders collected within regions on particular days) raised the concern that observed lead times might not be independent. Unseen intra-firm dynamics might reasonably be expected to influence operational performance measures such as lead time on a particular day. These dynamics might be operative at the firm, or, more conservatively, at the region level within day. Intraclass correlations verify a moderate *(p* = .19) yet significant (p < 0.01) degree of relatedness. In order to protect against any bias in the standard errors of the estimates, regressions were run with clustering at the region/day level.

Alternate models controlling for calendar month, quarter, and quarter ends were evaluated along with alternative specifications utilizing day as opposed to month as an index of time. Further analysis was completed for the post-deployment period using an enhanced model that included data on sales channel, and product type. Models were also evaluated that accounted for the possibility that extremely large or small orders might bias estimates of lead time performance. Finally, models were estimated that split the post-deployment period into multiple epochs. Each of the evaluated modeling alternatives yielded similar results. The model presented here was selected on the basis of fit, parsimony, and interpretability. Table 5 presents the statistical results.

Model results indicate a significant relationship between most model controls (main and interaction effects) and region-level order lead time. Main effects for control variables are consistent in both sign and significance in the pre- and post-deployment period. Linear hypothesis testing indicates significant contributions to the explanatory power of the model with the addition of time indices (month and month2). Linear hypothesis testing also reveals significant improvement in the explanation of variance with the addition of regional interaction terms for model controls and time indices. Results indicating differences in the effect of time on regional lead time performance suggest the possibility that relative performance may vary over the long-term.

Preparation for formal hypothesis testing commenced with the estimation of prototypical lead time performance for each of the operating regions in the post-deployment period. Prototypical lead times are calculated using regression results and region-level mean values for model controls. In addition to estimating lead time performance for each operating region in each time period, standard errors of these estimates are also calculated for each region/period. Using estimates of both regional lead time performance *and* the standard errors of those estimates, confidence intervals were calculated to see if differences in lead time performance are significant in any given period. Building on the ANCOVA results, similarities in performance are expected during early periods, followed by the emergence of significant differences in lead time over time. A summary of these differences is presented in Table 6. The estimates provided in the table indicate the number of days (untransformed) by which the first region leads the second (i.e., is faster at shipping) in lead time performance. For example, model estimates in the first operating month indicate that North America lead time performance was 3.75 days better (faster) than Europe performance. Similarly, Europe is estimated to be 1.09 days better than Asia during that same period (and thus North America is estimated to be 4.84 days better than Asia). A visual representation of the differences in estimated regional lead time performance is provided in Figure 3.

Results of the longitudinal analyses provide further support for Hypothesis 1. The period that immediately followed ES deployment was marked by similarity in order lead time performance across regions. Similar to the ANCOVA results, no significant differences existed in the lead time performance of any region pairs for at least the first three months of prototypical results.

The longitudinal analysis also provides support for Hypothesis 2. Over time, the differences in rates of performance change following ES deployment (reported in the regression results presented in Table 5) began to exert themselves. By the fourth month following ES deployment, statistically significant performance gaps had appeared between the Asia region and the other two operating regions. By the 17th month following ES deployment, a significant performance gap had also opened between North America and Europe. As depicted in Figure 3, the North America-Europe performance gap continued to widen through the end of the analysis period (month 24). Evidence suggests that the North America-Asia gap had stabilized at approximately 12.5 days difference. The Europe-Asia gap appears to be closing by the end of the analysis period.

It is noted that one possible explanation for closing performance gaps between regions may be the functional form imposed upon the data. Use of the exponential time index, although

providing more flexibility than strictly monotonic functions, results in a model that projects an indefinite period of lead time improvement for each of the regions. By definition, lead times must stabilize at some duration greater than zero, thus leveling off at some point on their improvement trajectories. Although the current functional form fits the data well in this case, future research might consider the adoption of more complex functional forms that accommodate this leveling off. In this way, an assessment might be made as to how performance gaps persist over periods extending past the two year interval examined in this study.

4.3. Managerial Differences

Differences in lead time performance at the end of the study period have significant implications for managers. Table 7 provides one estimate of the benefits that might accrue to Tristen if the Europe and Asia regions were to achieve the "best practice" performance levels of the North America region. These estimates are based on a calculation of potential working capital releases when Tristen's order-to-cash business cycle is shortened.

Estimates of working capital release are based on the average daily revenue of the firm associated with each operating region. For example, the Europe region accounted for approximately 29 percent of Tristen's $4 billion in annual sales. Assuming a 365 day year implies Europe revenue of approximately $3.178 million per day, capital that will be collected by the business when the customer pays, but that is "invested" in the business until that time. A reduction in the order-to-cash cycle (e.g., through a reduction in elapsed time between order and shipment) for the Europe region therefore stands to yield a "working capital release" of $3.178 million for every day the cycle is shortened. These estimates suggest that improving Europe's lead time performance to North America's levels would release nearly $11 million (3.44 days at $3.178 million in revenue per day) in working capital. Improving Asia's lead time performance to North America's level is estimated to be worth almost $26 million (12.44 days at $2.082 million in revenue per day). Total potential working capital releases of approximately $37 million are therefore possible. These estimates account for neither change to other possibly important performance measures, nor other benefits that might accrue to Tristen from improvements in responsiveness originally targeted by the ES implementation effort.

4.4. Performance Synchronization Following Deployment

Although not specifically hypothesized as part of this study, a striking finding of the analysis is the extent to which relative performance across regions synchronized following ES deployment. Synchronization implies that, independent of absolute levels, operating

performance began to improve and decline simultaneously across regions. For Tristen this synchronization is suggested in Figure 1 in that performance seems to rise and fall with greater uniformity following ES deployment. Table 8 sheds additional light on this phenomenon, providing correlation estimates between raw regional performance prior to and following deployment. The pre-deployment period was characterized by a relatively low, and sometimes negative, correlation between regional performance levels. Following deployment, correlation of regional performance increased dramatically and is uniformly positive.

Synchronization of performance might be explained by Tristen's goal to treat all regions equally with respect to serving customers, and an ES-enabled increase in inter-region inventory visibility. Tristen's implementation of ATP functionality, with visibility across the entire supply chain, can be argued to have enabled an inventory pooling opportunity for the firm. Firms typically use inventory pooling to let individual customer serving locations share inventory (and the demand variability it is intended to accommodate) as a strategy for reducing the risk of shortages while simultaneously reducing inventory levels and operational costs (Eppen 1979; Taragas 1999). However, inventory pooling may also be expected to lead to other performance changes as well.

In a setting in which all inventory is collocated (i.e., physically pooled), and in which all customers are treated equally, that customer service level (regardless of geographic area) would be similar is straightforward. That is, when the ability of the firm to serve customers in one area is diminished, it would be similarly diminished for other areas as well (assuming no single area or customer receives a preference). Thus, physically collocated inventory pooling might be expected to "synchronize" customer service levels, in this case, with regard to order lead times.

Rather than establish a single inventory location, the Tristen approach is more akin to the creation of a "Virtual Warehouse" (VW) that relies on information technology to provide operational efficiencies and global inventory visibility comparable to that achieved in single-location warehousing models (Stuart, Owen, and Landers 1995). Like other inventory-pooling approaches, VW implementations have demonstrated improved operational performance, including service levels and operational costs (Landers et al. 2000).

In short, managers' ability to allocate inventory from other stocking locations in service of regional orders is expected to yield the same opportunities and outcomes as those provided by physically pooling inventory. Thus, where performance may be uncorrelated and highly variable prior to the ES implementation at Tristen, the integration and visibility provided by those systems makes it reasonable to expect both a reduction in variability and a synchronization of performance in periods following deployment.

5. Discussion and Findings

This paper has investigated whether an ES implementation induced parity in the operational performance of three independent operating regions within a single firm. It further investigated whether that parity, if induced, persisted as a long-term consequence of ES deployment. Critics of ES have cautioned that, owing to the high mobility of IT assets, performance parity will emerge not just within, but also across firms (Davenport 2000; Carr 2004). This paper has analyzed empirical data in an effort to develop a better understanding of the extent to which performance parity is an outcome of ES deployment for at least one important measure of operational performance: order lead time.

5.1. Parity and Divergence Following ES Deployment

In the case of Tristen, managers sought to promote standardization and "common service levels" as part of an effort to deliver greater "value added" service to customers through the ES implementation. Clearly, they achieved these goals in the period immediately following ES deployment. By both statistical and managerial standards, performance levels became virtually indistinguishable across the North America, Europe, and Asia operating regions. Thus, the first hypothesis presented in this paper is supported.

Findings also provide support for the paper's second hypothesis, that parity in operational performance will tend to diminish over time. Results from both cross-sectional and longitudinal analyses indicate the emergence of performance gaps between regions over the two years following ES deployment. However, it is important to note that, despite the reemergence of regional differences, disparities in performance never reach the level exhibited in the pre-deployment period. This may lend some support to those who argue for "increased parity" but not for a "total parity" perspective.

The pattern of induced parity followed by gradual divergence lends support to emerging theories of ES deployment. In particular, empirical evidence of parity following ES deployment supports the notion of the "system" effects, those that derive specifically from the attributes and functionality of the technology itself, described in the ES literature (Bendoly 2001; Bendoly and Kaefer 2004). In the case of Tristen, it is suggested that the business process "templates" embedded in the deployed ES and the decision-making algorithms (e.g., inventory allocation and forecasting) that support them promote similarities in post-implementation business process execution. The fact that ES users may have few immediate alternatives to accepting the system

as deployed (Davenport 2000; Bendoly and Cotteleer 2004) and the expectation that similar processes yield similar results (Carr 2004), portend performance parity following ES deployment.

The implications of the long-run divergence in performance are less clear. The statistical evidence identifies a significant lead time difference and from a managerial perspective, $37 million in working capital is significant, even for a $4 billion company. The emergence of performance differences over time lends support to notions of "implementation" effects, outcomes associated with the unique decisions and business process changes that underpin ES deployment in a specific context, suggested by the ES literature (Bendoly 2001; Bendoly and Kaefer 2004).

Implementation effects might arguably be subdued in the Tristen context, as many aspects of the implementation were centrally controlled. However, even in this case, the presence of *resonant dissonance* (Bendoly and Cotteleer 2004) at the operating region level may explain the eventual ability of region-level users to influence local business process, enhancing "implementation" effects at the region level and promoting differences in regional performance. Reports from ES project managers at Tristen support the idea of *resonant dissonance* enhancing implementation effects. Specifically, managers reported a tendency of regional users to make "local adaptations" to accommodate "how they got their work done." The tendency for groups and organizations to evolve in this way, contingent on their respective resources and environments, is well accepted in the organizational literature (Lawrence and Lorsch 1967).

That performance both changes and diverges following ES deployment suggests that the implementation of these technologies does not have a deterministic impact on business process and performance. The results do, however, show the performance differential between operating regions to be smaller after than before ES deployment. This suggests that the technology may impose some boundaries on the performance of business processes. The strength of these boundaries (really, system effects under the model proposed by Bendoly (2001) and Bendoly and Kaefer (2004)) is not fully tested in this study.

For example, the context of this study imposed conditions under which performance parity, if it was going to emerge, was most likely to occur. Each operating region deployed the same technology at the same time while simultaneously adopting the same business practices as well as training and deployment approaches. Moreover, each of the regional business units sold the same products to a relatively homogenous customer group, subscribed to the same corporate performance objectives, and enjoyed the benefit of the same pre-deployment inventory buffer. To the extent that any of these factors differ across implementation sites or firms, similar levels of performance parity might be less likely to be observed. Future research

should seek to validate the findings of this study in both similar contexts and contexts in which some of the controls used here have been relaxed. Further research might also investigate whether there are differences in the system effects associated with process templates (that control work and information flow) versus those associated with the use of standardized-decision-supporting algorithms (e.g., inventory allocation and forecasting). An understanding of the differences between these effects stands to make a strong contribution to managers' ability to realize implementation benefits.

5.2. Synchronization of Operational Performance

Relative performance changes coincidental with ES implementation were also evident in increased synchronization among operating regions. The interpretation of this result as evidence of the integrating power of enterprise systems is consistent with McAfee's (2002) findings regarding the role ES plays in increasing visibility across an enterprise. It is also supported by two contextual considerations related to the study. First, Tristen's deployment approach called for implementation of a single instance of the ES applications across all sites (rather than unique implementations at each site). Second, Tristen's implementation specifically targeted improved integration and visibility across implementation sites.

Managers explicitly cited as an objective an improved ability to balance inventory across the company's internal supply network. Central to this goal was the deployment of ATP functionality. Such implementation is akin to the deployment of a "virtual warehouse", supporting Tristen's ability to realize the benefits of inventory pooling without resorting to centralized inventory management. To the extent that improved integration and visibility enabled this activity, performance synchronization might be expected. Although not specifically hypothesized as part of this study, this finding is important to managers seeking ways to improve the consistency of supply chain performance. To researchers, this finding suggests a continuing need to investigate the influence of inventory pooling benefits on the value of new alternatives for supply chain networks, echoing the call of other researchers (Kulkarni, Magazine, and Raturi 2004).

Future research might investigate the opportunity for firms to realize inventory pooling benefits through the creation of virtual warehouses. For example, studies might focus on firms that have only partially deployed ES, integrating some locations but not others. The extent to which performance is synchronized in integrated locations yet remains independent in nonintegrated locations may shed additional light on the impact of ES on the relative performance of operating units. Such studies might arguably be extended to cross-firm analyses of companies that have chosen to use ES to integrate with other participants in their supply

chains, thus contributing to the growing body of literature that centers on issues related to supply chain coordination (see, for example, Balakrishnan and Geunes 2004; Kraiselburd, Narayanan, and Raman 2004; Poundarikapuram and Veeramani 2004; Tsay and Agrawal 2004).

6.0. Conclusions

This study represents a single step in the investigation of whether enterprise systems promote parity in operational performance. Significant opportunities remain to continue this investigation. For example, where this study uses a single measure to gauge the influence of the information technology on performance convergence, other performance measures are important and might have yielded different convergence patterns. Future research can broaden definitions of operational performance to include other measures and see if similar patterns emerge.

In addition, the research reported here focuses on the performance of a single firm. To address concerns raised by critics of ES and information technology in general, future research should seek to assess the extent to which parity emerges across firms that have implemented similar enterprise technologies. Such studies might include operational measures of the kind investigated here as well as more widely available measures of financial or market performance. These studies might yet add to our growing understanding of the influences of ES on firm performance.

Notes

1 Additional details on SAP and Oracle enterprise application penetration can be found at their websites: www.sap.com and www.oracle.com.

Acknowledgments

References

Adler, P. S., K. B. Clark. 1991. Behind the learning curve: A sketch of the learning process. *Management Science* 37(3) 267-281.

AMR 2001. The report on enterprise management. AMR Research, Inc., Report, 42

Austin, R. D., M. J. Cotteleer. 1999. Current issues in IT: Enterprise resource planning and ecommerce. Harvard Business School, Boston, Massachusetts.

Austin, R. D., R. Nolan. 1998. How to manage ERP initiatives. Harvard Business School, Working Paper.

Badiru, A. B. 1992. Computational survey of univariate and multivariate learning curve models. *IEEE Transactions on Engineering Management* 39(2) 13.

Balakrishnan, A., J. Geunes. 2004. Collaboration and coordination in supply chain management and e-commerce. *Production and Operations Management* 13(1) 1-2.

Banker, R. D., J. M. Field, K. K. Sinha. 2001. Work-team implementation and trajectories of manufacturing quality: A longitudinal field study. *Manufacturing* & *Service Operations Management* 3(1) 25-42.

Bendoly, E. 2001. ERP systems and supply chain technologies: A caveat for B2B e-procurement. FedEx White Paper Series, Whitepaper, 28-36

Bendoly, E., M. J. Cotteleer. 2004. Resonant dissonance: Investigating sources of lagged operational variance following large-scale centralized system deployment. *9th Conference on Information Systems and Technology (CIST),* Denver, Colorado.

Bendoly, E., F. Kaefer. 2004. Business technology complementarities: Impact of the presence and timing of ERP on B2B e-commerce technology efficiencies. *Omega* 32(5) 395-405.

Box, G. E. P., D. R. Cox. 1964. An analysis of transformations. *Journal of the Royal Statistical Society* Series B (26) 211-243.

Brynjolfsson, E., L. Hitt. 1998. Beyond the productivity paradox: Computers are the catalyst for bigger changes. *Communications of the ACM* 41(8) 49-55.

Carr, N. G. 2004. Does IT matter?: Information technology and the corrosion of competitive advantage. Harvard Business School Publishing, Boston, Massachusetts.

Cotteleer, M. J., E. Bendoly. 2006. Order lead-time improvement following Enterprise-IT implementation: An empirical study. *MIS Quarterly* (Forthcoming).

Davenport, T. 2000. Mission critical: Realizing the promise of Enterprise systems. Harvard Business School Press, Boston, Massachusetts.

Dewan, S., K. L. Kraemer. 1998. International dimensions of the productivity paradox. *Communications of the ACM* 41(8) 56-62.

Eppen, G. D. 1979. Effects of centralization on expected costs in a multi-location newsboy problem. *Management Science* 25(5) 498-501.

Epple, D., L. Argote, K. Murphy. 1996. An empirical investigation of the microstructure of
   knowledge acquisition and transfer through learning by doing. *Operations Research* 44(1)
   77-86.

Garvin, D. A. 1998. The processes of organization and management. *Sloan Management
   Review* 39(4) 33-50.

Hayes, R. H., S. C. Wheelwright, K. B. Clark. 1988. Dynamic manufacturing: Creating the
   learning organization. The Free Press, New York, New York.

Hitt, L. M., D. J. Wu, X. Zhou. 2002. Investment in enterprise resource planning. *Journal of
   Management Information Systems* 19(1) 71-98.

Hult, G. T., D. J. Ketchen, E. L. Nichols. 2002. An examination of cultural competitiveness and
   order fulfillment cycle time within supply chains. *Academy of Management Journal* 45(3)
   577-586.

Jacobs, F. R., E. Bendoly. 2003. Enterprise resource planning: Developments and directions for
   operations management research. *European Journal of Operational Research* 146(2)
   233-240.

Kraiselburd, S., V. G. Narayanan, A. Raman. 2004. Contracting in a supply chain with stochastic
   demand and substitute products. *Production and Operations Management* 13(1) 42-62.

Kulkarni, S. S., M. J. Magazine, A. S. Raturi. 2004. Risk pooling advantages of manufacturing
   network configuration. *Production and Operations Management* 13(2) 186-199.

Kumar, K., J. Van Hillegersbert. 2000. ERP experiences and evolution. *Communications of the
   ACM* 43(4) 22-26.

Landers, T. L., M. H. Cole, B. Walker, R. Kirk. 2000. The virtual warehousing concept.
   Transportation Research Part E 36E(2) 115-125.

Lawrence, P. R., J. W. Lorsch. 1967. Organization and environment. Harvard Business School
   Press, Boston, Massachusetts.

Mabert, V. A., A. Soni, M. A. Venkataramanan. 2000. Enterprise resource planning survey of US
   manufacturing firms. *Production and Inventory Management Journal* 41(2) 52-58.

McAfee, A. 2002. The impact of enterprise information technology adoption on operational
   performance: An empirical investigation. *Production and Operations Management* 11(1)
   1-21.

Miller, R. G. 1986. Beyond ANOVA: Basics of applied statistics. Wiley, New York, New York.

Nahmias, S. 1997. Production and operations analysis. Richard D. Irwin, Chicago, Illinois.

Neter, J., W. Wasserman, M. H. Kutner. 1990. Applied linear statistical models: Regression, analysis of variance, and experimental design. Richard D. Irwin, Inc., Boston, Massachusetts.

Pisano, G., S. Rossi. 2001. ITT automotive: Global manufacturing strategy-1994. Harvard Business School Case Study, Boston, Massachusetts.

Pisano, G. P., R. M. Bohmer, A. C. Edmondson. 2001. Organizational differences in rates of learning: Evidence from the adoption of minimally invasive cardiac surgery. *Management Science* 47(6) 752-768.

Poundarikapuram, S., D. Veeramani. 2004. Distributed decision-making in supply chains and private e-marketplaces. *Production and Operations Management* 13(1) 111-121.

Razi, M. A., J. M. Tarn. 2003. An applied model for improving inventory management in ERP systems. *Logistics Information Management* 16(2) 114-125.

Scheer, A. W., F. Habermann. 2000. Making ERP a success. *Communications of the ACM* 43(4) 57-61.

Stuart, D. E., J. Owen, T. L. Landers. 1995. Establishing the virtual warehouse. *1995 ASME International Mechanical Engineering Congress,* San Francisco, California.

Taragas, G. 1999. Pooling in multi-location periodic inventory distribution systems. *Omega* 27(1) 39-59.

Tsay, A. A., N. Agrawal. 2004. Channel conflict and coordination in the e-commerce age. *Production and Operations Management* 13(1) 93-110.

Webster's. 1977. Webster's New World Dictionary of the American Language. The World Publishing Company, New York, New York.

Weill, P. 1992. The relationship between investment in information technology and firm performance: A study of the valve manufacturing sector. *Information Systems Research* 3(4) 307-333.

Willcocks, L. P., R. Sykes. 2000. The role of the CIO and IT function in ERP. *Communications of the ACM* 43(4) 32-38.

Yelle, L. E. 1979. The learning curve: Historical review and comprehensive survey. *Decision Sciences* 10 302-328.

Appendix

Table 1: Model Variable Definitions and Summary Statistics*

Variable	Definition	Mean	Std. Dev.	Min	Max
leadtime	Elapsed time between order booking date and date recorded as shipped to customer.	33.44	35.65	0	434
month	A time index variable used to capture the number of months a transaction was executed prior to, or after ES deployment.	9.39	9.90	−12	24
europe	Indicator variables identifying records originating from the Europe operating region.	0.29	0.45	0	1
asia	Indicator variables identifying records originating from the Asia operating region. (records that originated in North America were identified where both Europe and Asia were set equal to zero).	0.19	0.40	0	1
daycount	The total number of order transactions processed within the operating region on the day the order was taken. A measure of business intensity	94.58	90.29	1	833
units	The total number of product units (in thousands) associated with the order. A measure of order size.	0.64	1.86	0	116
value	The total dollar value (in thousands) of the sales order. A measure of order value.	262.61	1728.00	0	126000

*$N = 112{,}818$.

Table 2: Correlation Among Model Continuous Model Variables

	Leadtime	Month	Daycount	Units	Value
Leadtime	1.000				
Month	0.012	1.000			
Daycount	0.217*	0.139*	1.000		
Units	0.045*	0.047*	0.057*	1.000	
Value	0.104*	0.080*	0.105*	0.607*	1.000

* Significant at $p < .01$.

Table 3: Mean Differences (in days) Between Operating Region Lead Times (Controlling for Differences in Order Units, Order Value, and Daily Regional Order Volume)

Lead time comparison	Pre-imp	At imp	1 year	2 years
North America vs. Europe	9.47***	2.49	2.65	−3.83*
	(2.75)	(2.23)	(2.09)	(2.13)
North America vs. Asia	−4.01	1.97	−3.93	−8.82***
	(3.96)	(3.18)	(3.23)	(3.22)
Europe vs. Asia	−15.06***	−2.06	−6.21**	−6.03**
	(3.32)	(2.55)	(2.76)	(2.62)

*** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$.

Table 4: Linear Hypothesis (F) Test Results on Equality of Order Lead Time Differences at Alternative Pre- and Post-Deployment Intervals

Lead time comparison	(1) At Imp. vs. Pre-Imp. F-stat.	(2) At Imp. vs. 1 year F-stat.	(3) At Imp. vs. 2 years F-stat.	(4) Pre-Imp. vs. 2 years F-stat.
North America vs. Europe	4.77**	0.00	4.60**	17.01***
North America vs. Asia	1.66	2.07	6.91***	1.04
Europe vs. Asia	9.68***	1.31	1.27	4.55**

*** $p < 0.01$; **$p < 0.05$; *$p < 0.1$.

Table 5: Summary of OLS Pre- and Post-Deployment Lead Time Performance (Clustered Standard Errors in Parentheses)

Deployment period	Pre-deployment coefficients (std errors)	Post-deployment coefficients (std errors)
Controls and constant		
constant	6.174 (1.283)***	−1.466 (0.464)***
daycount	0.107 (0.014)***	0.575 (0.090)***
units	−0.199 (0.060)***	−0.516 (0.027)***
value	0.737 (0.049)***	0.350 (0.025)***
Longitudinal improvements		
month	1.242 (0.302)***	0.149 (0.032)***
month2	0.079 (0.019)***	−0.008 (0.001)***
Regional differences		
europe	−3.956 (3.390)	−0.178 (0.546)
asia	−5.703 (4.256)	3.292 (0.606)***
europe*daycount	0.216 (0.055)***	0.289 (0.116)**
asia*daycount	0.545 (0.058)***	0.012 (0.157)
europe*units	−0.101 (0.304)	0.106 (0.042)**
asia*units	0.137 (0.383)	0.614 (0.042)***
europe*value	−0.328 (0.219)	−0.086 (0.035)**
asia*value	−0.430 (0.311)	−0.379 (0.029)***
europe*month	−1.366 (0.862)	−0.078 (0.045)*
asia*month	−2.506 (0.887)***	0.016 (0.065)
europe*month2	−0.086 (0.051)*	0.003 (0.001)**
asia*month2	−0.156 (0.054)***	0.002 (0.002)
Total R^2 (adj.) (w/*all* above terms)	0.222	0.114
ΔR^2 (on adding month terms, given controls)	0.009***	0.020***
ΔR^2 (on adding regional difference terms, given controls and month terms)	0.031***	0.044***

*** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$.

Table 6: Estimated Days Differences in Regional Order lead time Performance[†]

Operating period	North America vs. Europe difference in means	North America vs. Asia difference in means	Europe vs. Asia difference in means
1	3.75	4.84	1.09
2	3.29	5.25	1.97
3	2.84	5.69	2.85
4	2.42	6.15*	3.74*
5	2.03	6.63**	4.61**
6	1.68	7.13***	5.45**
7	1.38	7.64***	6.27***
8	1.13	8.16***	7.03***
9	0.94	8.69***	7.75***
10	0.82	9.21***	8.40***
11	0.75	9.73***	8.98***
12	0.75	10.23***	9.49***
13	0.80	10.71***	9.91***
14	0.92	11.17***	10.25***
15	1.08	11.58***	10.50***
16	1.28	11.95***	10.66***
17	1.52*	12.26***	10.74***
18	1.79**	12.52***	10.72***
19	2.08**	12.70***	10.62***
20	2.38***	12.82***	10.44***
21	2.67***	12.85***	10.18***
22	2.95***	12.80***	9.85***
23	3.21***	12.66***	9.45***
24	3.44***	12.44***	8.99***

*** $p < 0.01$; ** $p < 0.05$; *$p < 0.10$.

[†] Estimates calculated after "untransforming" lead time estimates for individual regions and taking differences.

Table 7: Estimated Additional Working Capital Release (in thousands) if All Regions Achieved North America "Best Practice"

Tristen revenue profile		Attributable revenue by region	%	$$$
Tristen Annual Revenue:	$4,000,000	North America	52%	5,699
Annual Days of Operation	365	Europe	29%	3,178
Tristen Revenue per Day (est.): $	11,000	Asia	19%	2,082
		Total		10,959

	End of study lead time difference from North America	Potential working capital released
Europe	3.44 days	$10,932
Asia	12.44 days	$25,900
	Total	$36,832

Table 8: Correlation of Raw Regional Performance, Pre- and Post-ES Deployment

	Pre-deployment				Post-deployment		
	North America	Europe	Asia		North America	Europe	Asia
North America	1.000			North America	1.000		
Europe	0.450	1.000		Europe	0.622***	1.000	
Asia	0.039	−0.460	1.000	Asia	0.606***	0.724***	1.000

*** $p < 0.01$; ** $p < 0.05$; * $p < 0.10$.

Figure 1: Raw Data Plot, by Month, of Regional Lead Time Performance During the Pre-deployment (Month < 0) and Post-deployment (Month > 0) Periods.
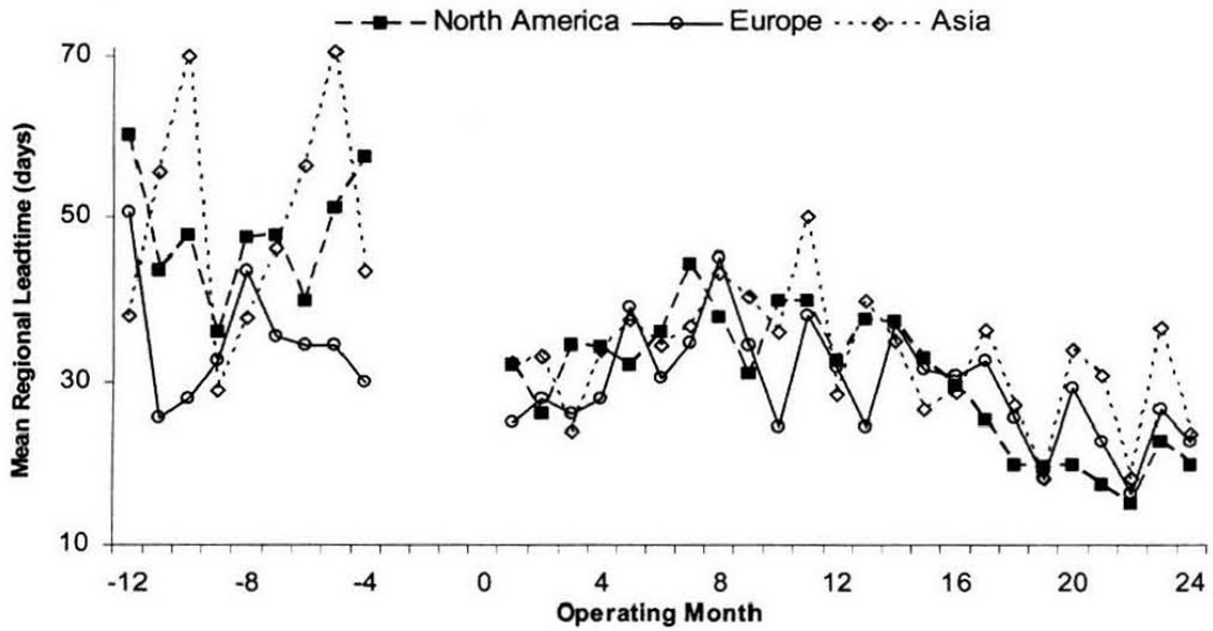
Figure 2: Absolute Value ANCOVA Plot of Lead Time Days Differences Among Regions
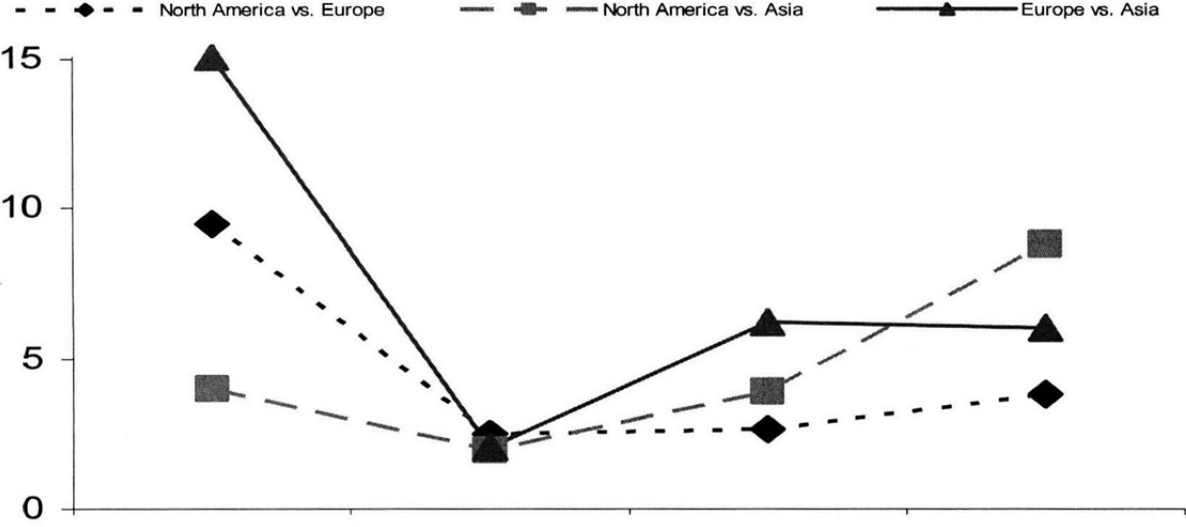
Figure 3: Estimated Days Difference in Regional Order Lead Time Performance (by Operating Period)