

COST IMPLICATIONS OF DESIGN/BUILD CONCURRENCY

 *Donald Birchler, Gary Christle, and Eric Groo*

Developing a weapon while in production does increase program risk and is sometimes cited as a reason for cost growth. This article explores the relationship between concurrency and cost growth in large weapon programs. The authors defined concurrency as the proportion of research, development, and test and evaluation appropriations authorized during the same years in which procurement appropriations are authorized. Their results strongly indicate that concurrency does not necessarily predict cost growth. Using classical regression techniques, the authors found no evidence supporting this relationship. To investigate other relationships between cost growth and concurrency, they also used a smooth curving technique. These experiments showed that, although the relationship is not strong, low levels of concurrency are more problematic than higher levels.

Keywords: *Weapon System; Cost Growth; Concurrency; Ordinary Least Squares (OLS) Regression; Research, Development, Test and Evaluation (RDT&E)*

General characteristics

Length: 62 ft 1 in (18.90 m)

Wingspan: 44 ft 6 in (13.56 m)

Height: 16 ft 8 in (5.08 m)

Empty weight: 840 lb (381 kg)

Maximum takeoff weight: 10,000 lb (4,536 kg)

Maximum speed: 1,000 mph (1,609 km/h)

Service ceiling: 65,000 ft (19,812 m)

Range: 2,000 mi (3,218 km)

Maximum payload: 18,000 lb (8,165 kg)

Maximum fuel capacity: 18,000 lb (8,165 kg)

Maximum cruise speed: 1,000 mph (1,609 km/h)

Maximum climb rate: 100 ft/s (30.5 m/s)

Maximum turn rate: 10 g

Maximum roll rate: 10 deg/s

Maximum altitude: 65,000 ft (19,812 m)

Maximum range: 2,000 mi (3,218 km)

Maximum service ceiling: 65,000 ft (19,812 m)

Maximum payload: 18,000 lb (8,165 kg)

Maximum fuel capacity: 18,000 lb (8,165 kg)

Maximum speed: 1,000 mph (1,609 km/h)

Maximum climb rate: 100 ft/s (30.5 m/s)

Maximum turn rate: 10 g

Maximum roll rate: 10 deg/s

Maximum altitude: 65,000 ft (19,812 m)

Maximum range: 2,000 mi (3,218 km)

Maximum service ceiling: 65,000 ft (19,812 m)

Maximum payload: 18,000 lb (8,165 kg)

Maximum fuel capacity: 18,000 lb (8,165 kg)

Maximum speed: 1,000 mph (1,609 km/h)

Maximum climb rate: 100 ft/s (30.5 m/s)

Maximum turn rate: 10 g

Maximum roll rate: 10 deg/s

Maximum altitude: 65,000 ft (19,812 m)

Maximum range: 2,000 mi (3,218 km)

Maximum service ceiling: 65,000 ft (19,812 m)

Maximum payload: 18,000 lb (8,165 kg)



Typically, defense programs experience some level of concurrency; that is, production of the weapon system happens while some portions of the design are still being completed. Many people within the defense acquisition community argue that high levels of design/build concurrency ultimately lead to cost growth, as it implicitly creates a greater level of risk. For example, a memorandum from the Assistant Secretary of the Navy for Research, Development and Acquisition (ASN-RDA) identified the high degree of concurrency in the Littoral Combat Ship as being a large contributor to the program's overall cost growth (DoD, 2006).

In a zero-risk world, the requirements, concept of operations, and substantial prior development would be completed before the release of the Request for Proposal (RFP) for the design phase. In addition, 100 percent of the design would be complete before the release of the production RFP; and all the initial material/components would always be procured and available before production started. Moreover, requirements would not change once design started, design would not change once production started, and production would flow smoothly without delays caused by late software or hardware. Thus, in a zero-risk world we would say programs have zero overlap, or concurrency, and virtually no production risk.

Unfortunately, this zero-risk approach to production planning is impossible to achieve, and even if it were, many would argue that it is not desirable. The Japanese, for example, pioneered the "just-in-time" inventory strategy, where materials essential for production are not only unavailable before production start, they are deliberately fabricated and delivered at the last possible moment to reduce in-process inventory, thus reducing storage and finance costs associated with inventory beyond what is immediately needed. No financial or technical reasons preclude production in one portion of the program while a design is completed on an unrelated portion.

Other reasons to inject plans with some design/build concurrency, despite potential increases in risk of cost growth, include (a) urgent need for the product, (b) maintaining the industrial base, (c) avoiding obsolescence, and (d) reducing exposure to requirements changes.

Consequently, major programs always retain some level of concurrency, much of which is actually an integral part of the plan (see [a] through [c], previous paragraph).

In sum, there seems to be a good case to be made that concurrency is actually desirable and possibly reduces cost, and another equally good case that it adds risk, which ultimately leads to cost growth. Unfortunately, despite decades of interest in concurrency within the acquisition community, the literature on concurrency is surprisingly thin. We found only one study conducted by the

Congressional Budget Office (1988) that specifically looked at the correlation between concurrency and cost growth. Another more recent study by RAND touched on concurrency, but was primarily about other factors that lead to cost growth (Arena, Leonard, Murray, & Younossi, 2006). However, in both cases, the studies found only a very weak correlation between concurrency (defined as the overlap between operational test and evaluation and production) and cost growth. Our study, done on behalf of the ASN-RDA, examines this relationship in more detail using a slightly different definition of concurrency and a larger data set.

Definitions

In general, a lack of consensus prevails regarding the meaning of concurrency in acquisition programs. We chose a definition that reflects the most general use of concurrency and was tractable for analysis given the data available for large acquisition programs. Other definitions for concurrency exist and likely have different implications in those contexts. Our definitions for concurrency and cost growth follow:

- *Concurrency* is the proportion of research, development, test and evaluation (RDT&E) appropriations that are authorized during the same years that procurement appropriations are authorized. This proportion is further restricted to the first 95 percent of total RDT&E spending.
- *Cost growth* is, after adjusting for quantity changes and inflation, the proportional increase of the final cost to the initial cost estimate.

We chose 95 percent of the total RDT&E appropriation because RDT&E monies continue throughout the life of the program, albeit at a much reduced rate toward the design-complete/testing-complete phase of the program. This is usually due to the ongoing need for updates and modifications, but has little bearing on concurrency issues. We were satisfied, after a little experimentation, that the 95 percent cutoff addressed this for most programs.

This measure is not a perfect proxy for concurrency. If anything, the 95 percent cutoff likely overstates concurrency. Moreover, it also misses concurrency in related programs that can have a significant effect on cost and schedule of an item such as concurrency of weapon production with development of items designated for the

weapon, but being developed under other programs (such as radar, sonar, etc., which are being developed for more than one platform).

The definition for cost growth may initially appear to be overly broad, allowing for the inclusion of costs that are completely unrelated to concurrency. However, adjustments for these costs would have been much more complex, requiring systemic changes in both the initial estimates and final profiles tailored to each program. Out of concern that this process would become ad hoc, we left the definition broadly defined with adjustments made for quantity and inflation only.

Data

For measures of cost growth and concurrency, we gathered data from Selected Acquisition Reports (SARs) on the procurement and RDT&E profiles. The SARs are available in the Defense Acquisition Management Information Retrieval System (DAMIR). We reviewed this list and selected programs based on their maturity and availability of data.

On some occasions, we needed to drop new lines of production from final SARs that were absent from the first SAR. This was necessary to make apples-to-apples comparisons, particularly when controlling for quantity changes, due to the tendency of some programs to add on additional lines of production to existing procurement programs.

An illustrative case is the V-22. The initial estimate for the program was essentially for a single airframe for use by the Marines. During the course of the program, the Air Force Special Forces ordered a modified version of the airframe. This new line of production, however, cost dramatically more than the Marine version, presumably because of modifications and enhancements necessitated by the requirements for Special Forces operations. This growth in unit costs was obviously due to scope changes and not incidental to changes in program quantity or concurrency—our primary controls in this study. Fortunately, the additional RDT&E and procurement costs associated with these units were entirely funded out of Air Force appropriations, making it relatively easy to exclude these costs from the cost growth and concurrency calculations. For other programs with similar issues, where the distinction was less apparent, we reviewed budget exhibits and other publicly available budget justification materials for information to tease new subprograms away from historical program plans. This was not always possible, leading us to drop several programs from the analysis.

Method

Our approach was driven by two primary questions:

- Relative to cost growth, is there an ideal amount of concurrency that should be programmed for large acquisitions?
- If there is no “ideal,” what is the relationship, if any, between cost growth and concurrency?

These questions suggested a hybrid approach, employing traditional statistics and hypothesis testing methods as well as more modern methods of data exploration. First, using Ordinary Least Squares (OLS) regression, we fit a global quadratic function to the data. The quadratic model did not fit the data well, which led us to consider a second approach—locally weighted scatterplot smoothing (LOESS). LOESS is a nonparametric regression method, which allows the data to express and inform without restricting the data to fit some function. We assessed the results of this second approach with a bootstrapping technique (Efron & Tibshirani, 1998).

To ensure that we were using completed cost growth profiles, we sampled from mature programs, defined as programs that had begun Initial Operating Capability, contained in DAMIR. Of these, after discarding programs for which we were unable to locate initial baseline cost estimates, we were left with an initial set of 43 programs. For these complete programs, we used the procurement and RDT&E acquisition profiles to calculate cost growth and concurrency.

To facilitate making statistical inferences about concurrency, we first needed to directly control a few known, significant influences. First, to control changes in base years dollars between SARs, we rebaselined all the reported costs in constant 2009 dollars using the appropriate inflation indices in the National Defense Budget Estimates, commonly referred to as the “Green Book.”¹ The Green Book published indices only out to 2014, so to adjust programs that were funded past this year, we extended the indices at a fixed rate of 2 percent per year.

Second, we needed to adjust procurement cost growth to reflect changes in quantity. When the first SAR is published, a procurement profile shows how many units will be purchased and in what years they will be purchased. The amount of units purchased affects the procurement cost estimates via the learning curve effect. That is, the marginal costs of production drop with quantity as, with each additional unit, workers become more efficient, manufacturing

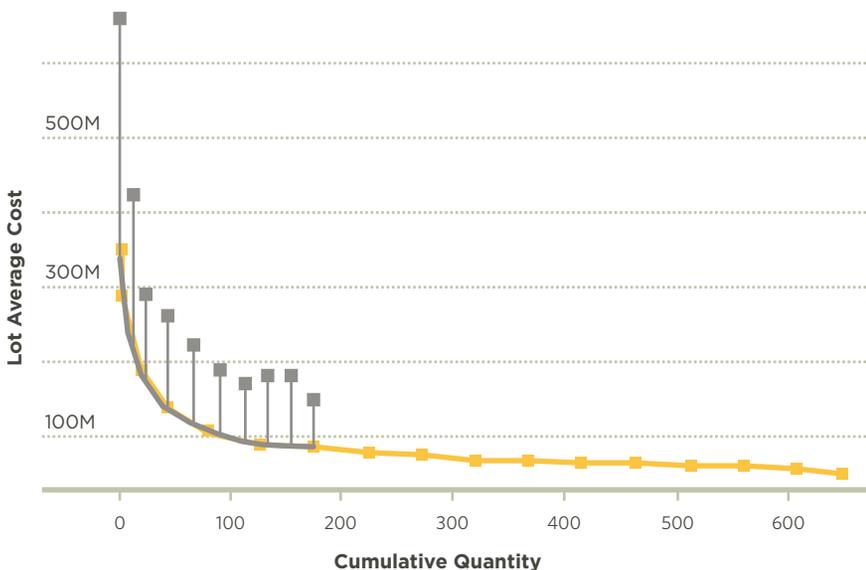
processes are refined, and quality control improves. This process is incorporated into every baseline cost estimate.

Thus, it is important to adjust the original cost estimates reflected in the first SAR to account for the changes in the quantity procured. For example, a program that was originally going to purchase 100 units at a total procurement cost of \$1,000 faces a budget cut by Congress leading to only 50 units being bought. A new baseline estimate could be calculated by simply taking the average cost per unit (i.e., \$10) and subtracting these 50 units out of the procurement funds. Using that method, we would simply multiply 50 units by \$10 to get an adjusted original cost baseline of just \$500.

But that is not satisfactory. In fact, by not buying the other 50 units, the program does not experience the same level of learning, and the average cost per unit actually rises as a result. In our example cited previously, the average cost per unit would rise to something over \$10, and the procurement savings would be less than \$500. Using a technique pioneered by Goldberg and Touw (2003), we were able to estimate the learning curve effects for the programs in our data set and adjust the original cost baseline up or down, depending on whether fewer or more units were procured.

Figure 1 illustrates the learning curve adjustment for the F-22 program. The gray squares correspond to the quantities and costs reported in the first SAR. Notice that the gold squares curve sharply downward, but then flatten out as the total quantity increases. This pattern corresponds to an anticipated initial period of intensive

FIGURE 1. LEARNING CURVE ADJUSTMENT ILLUSTRATION



learning, which progressively tapers as the gains from learning disappear. The gold line is the estimated learning curve. What is most striking about the line is how closely it appears to fit the data, without any additional modification.

Reducing the procurement quantities increases the average costs of the units purchased, as the lower cost units at the end of the production run are not added into the total production run. Thus, if the program had followed the initial learning curve, then the lot average cost would have fallen along the upper portion of the initial estimate. These adjusted lot average costs, indicated by the gray line, form the new baseline for measuring cost growth. The gray squares correspond to the actual quantities and costs reported in the final SAR profile for the F-22 program. Despite higher than expected total costs, the average unit costs decline at a rate reflective of the original estimate. Comparing the gray squares to the gray line, we can measure cost growth as the difference between the adjusted initial estimate and the final reported cost profile for a program. This is literally the area demarcated by the horizontal dotted lines.

Most of the programs that we examined experienced some change to the procurement quantities. This adjustment required stable associations between procurement costs and units for programs between the first and final cost profiles. Unfortunately, this requirement reduced the data set to only 28 programs suitable for analysis (Table 1).

Results

For procurement cost growth, we wanted to see if there was any correlation to concurrency, as measured by the percent of RDT&E spending that occurs when procurement spending is happening at the same time. As mentioned in the method section of this article, we calculated concurrency in two ways. First, we used the first published SAR to determine planned concurrency. We then used the last SAR to calculate actual concurrency. Thus, for each element of cost growth, we looked for correlations with two different measures of concurrency.

Based upon the feedback that we received from various Navy and DoD acquisition officials, we decided that a good starting hypothesis was that concurrency follows a Goldilocks rule (not too much, not too little, but somewhere in the middle being optimal). Too little concurrency is bad for a program as serial design and production yields a longer duration (and thus more cost) before fielding of the weapon. Too much concurrency is also bad as it accepts too

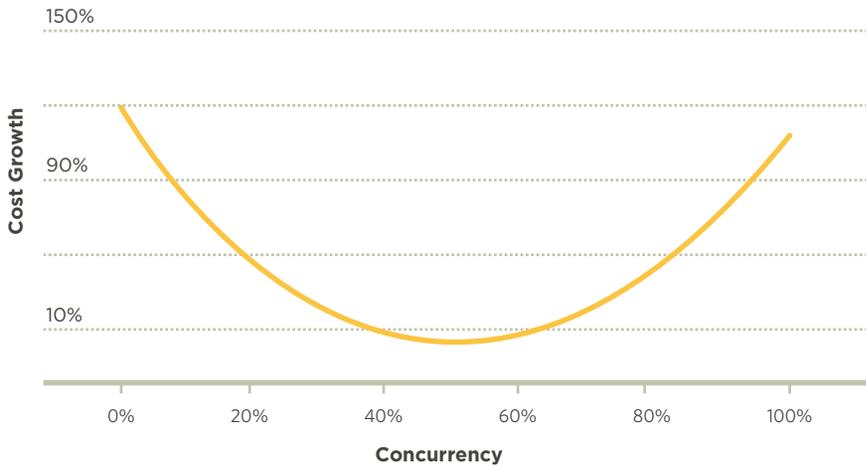
TABLE 1. LIST OF WEAPON SYSTEMS

Program	PNO
AIM 9X Sidewinder Missile	581
Air Warning and Control System Radar System Improvement Program (AWACS RSIP)	524
Bradley Upgrade	601
C-17A Globemaster III	200
CH-47F Improved Cargo Helicopter	278
EA-18G Growler	378
F-22 Raptor	265
F/A-18 E/F Super Hornet	549
Family of Medium Tactical Vehicles (FMTV)	746
High Mobility Artillery Rocket System (HIMARS)	367
Joint Air-to-Surface Standoff Missile (JASSM)	555
Joint Direct Attack Munition (JDAM)	503
Joint Primary Aircraft Training System (JPATS)	560
Longbow Apache Helicopter	831
Longbow Hellfire Missile	541
MH-60R Seahawk® Helicopter	191
MH-60S Seahawk® Helicopter	282
MHC 51 Osprey Minehunter	772
Minuteman III Guidance Replacement Program (MM III GRP)	302
Sense and Destroy Armor (SADARM)	735
Small Diameter Bomb (SDB)	354
SFW	275
SSN-21 Seawolf-class Attack Submarine	258
SSN-774 Virginia-class Attack Submarine	516
Stryker Light Armored Vehicle	299
T-45S	240
Tactical Tomahawk Missile	289
V-22 Osprey Tiltrotor Aircraft	212

much technical risk. Thus, some moderate level of concurrency would be the optimal in the sense that it minimizes cost growth. This would yield a curve similar to that shown in Figure 2.

The logic behind this approach for planned concurrency is relatively simple. Program managers plan for a certain level of funding concurrency. If they plan for too much, they may accept too much risk that could yield cost growth. On the other hand, too

FIGURE 2. HYPOTHETICAL QUADRATIC RELATION BETWEEN COST GROWTH AND CONCURRENCY



little funding concurrency forces them to create completely serial development/design and production processes that prolong program duration and also create cost growth. In sum, the planned level of concurrency forces managers to make decisions that ultimately lead to cost growth if either too much or too little concurrency is accepted.

The logic for the actual concurrency follows along similar lines. Program managers may or may not have planned for concurrency, but events led to the situation where some level of concurrency occurred, which, if too high or too low, led to excessive cost growth. Again, the assumption is that some intermediate level of actual concurrency would be the optimum.

In all cases, this simple rule can be specified with the following function, which was estimated using OLS:

$$CostGrowth = b_0 + b_1Concurrency + b_2Concurrency^2 + e \tag{1}$$

Planned Concurrency

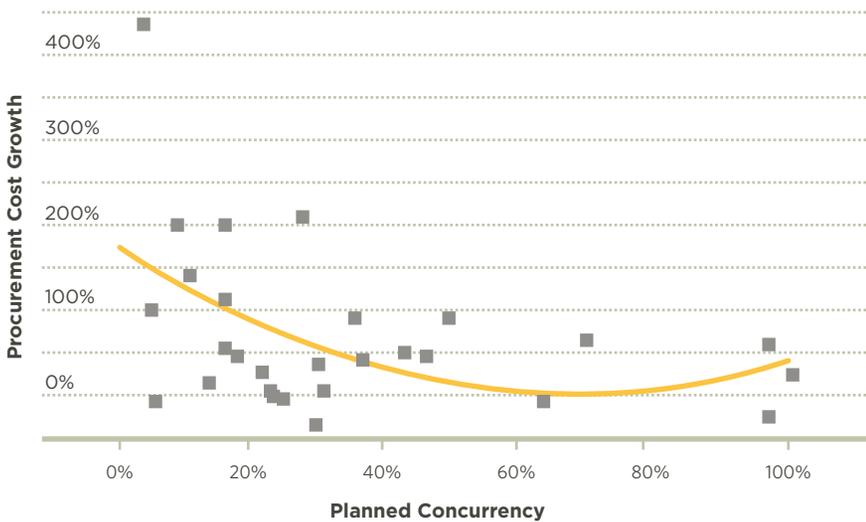
Our first model explored the relation between planned concurrency and procurement cost growth. The results are reported in Table 2.

Observe that two of the parameters are statistically significant at the .10 level (i.e., the probability that the parameters are less than 10 percent is zero), and the fitted line does give us a U-shaped curve (Figure 3). However, the adjusted R-squared is very low, which

TABLE 2. OLS RESULTS: PROCUREMENT COST GROWTH VS. PLANNED CONCURRENCY

	Estimate	Std Error	P-Value
Intercept	1.825	0.484	0.001
Concurrency	-5.014	0.465	0.052
Concurrency ²	3.667	2.273	0.119
Adjusted R ²	0.137		

FIGURE 3. FITTED CURVE: PROCUREMENT COST GROWTH VS. PLANNED CONCURRENCY



forces us to conclude that the quadratic model has little predictive power of procurement cost growth.

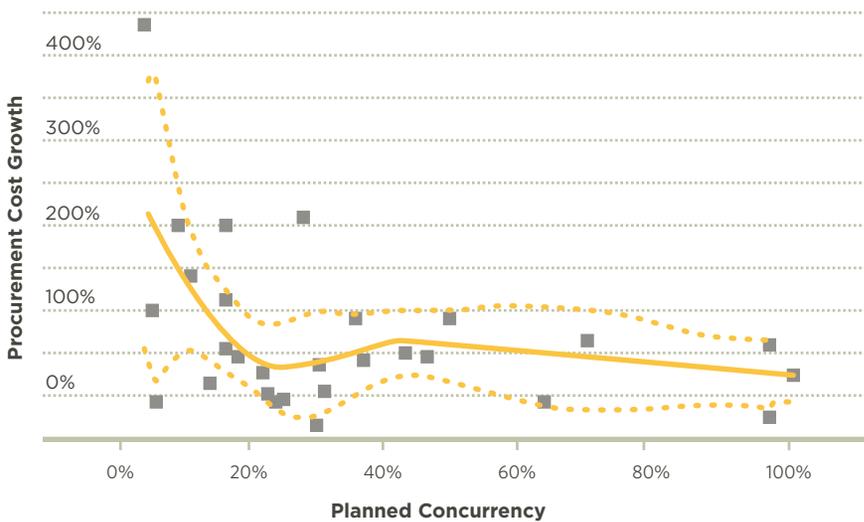
Note that much of the curvature in the model comes from one outlier. To see how well the model improves without this data-point, we ran the same model excluding the outlier (Table 3). This resulted in no improvement to the model at all and slightly less curvature.

Finally, to see if some other possible relation was evident, we ran the LOESS smooth curving routine on all of the data including the outlier. We then bootstrapped the 90 percent interquartile range to see how well conditioned the data are to the original curve. If the data are from a common model, the smoothed curves generated by the repeated sampling should be similar to the original, and the confidence intervals defined should be fairly tight around the original curve.² The results of these exercises using the outlier data-point discussed previously and excluding this data-point can be seen in the figures that follow.

TABLE 3. OLS RESULTS: PROCUREMENT COST GROWTH VS. PLANNED CONCURRENCY (OUTLIER EXCLUDED)

	Estimate	Std Error	P-Value
Intercept	1.097	0.390	0.010
Concurrency	-1.907	1.934	0.334
Concurrency ²	1.157	1.764	0.518
Adjusted R ²	0.021		

FIGURE 4. PROCUREMENT COST GROWTH VS. PLANNED CONCURRENCY

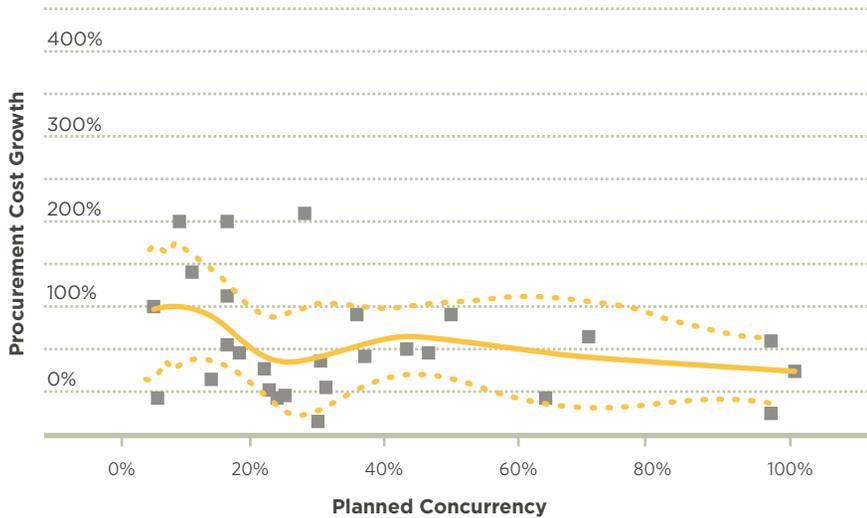


As we can see from Figure 4, the interval using the outlier data-point is extremely wide. For example, if a program had planned concurrency of .2, then, within the 90 percent interquartile range, the procurement cost growth for that program could easily range from 50 percent to over 100 percent.

To ensure that the outlier was not a significant factor in these results, we ran the same experiment excluding this data-point. This did not improve the results in any discernible way (Figure 5).

In spite of the fact that the confidence intervals around the original LOESS curves are wide, we do see a pattern in the data that suggests that low levels of planned concurrency are more problematic than higher levels of concurrency. Again, turning to the data without the outlier, we calculated the mean cost growth in procurement for those programs with planned concurrency levels under 30 percent and compared it to the means for those programs with planned concurrency over 30 percent. Those under 30 percent experienced, on average, approximately 110 percent cost growth

FIGURE 5. PROCUREMENT COST GROWTH VS. PLANNED CONCURRENCY (OUTLIER EXCLUDED)



while those over 30 percent experienced an average cost growth of approximately 50 percent. This difference was statistically different at the 95 percent confidence level.

Actual Concurrency

We next turn our attention to procurement growth as a function of actual concurrency. Table 4 shows the results of estimating the quadratic model using OLS. As in the case with planned concurrency, only the intercept parameter β_0 is significant at the .01 level. The model as a whole has an adjusted R-squared of -0.01889 indicating that the model has little explanatory power. Note also that the fitted line in Figure 6 is concave, which is the exact opposite of what our hypothesis was (i.e., a U-shaped curve).

TABLE 4. OLS RESULTS: PROCUREMENT COST GROWTH VS. ACTUAL CONCURRENCY

	Estimate	Std Error	P-Value
Intercept	1.037	0.530	0.062
Concurrency	-0.453	2.399	0.852
Concurrency ²	-0.275	2.168	0.900
Adjusted R ²	-0.019		

We also note the existence of an outlier that could exhibit a fairly large effect on the model (Figure 6). To account for this possibility,

FIGURE 6. FITTED CURVE: PROCUREMENT COST GROWTH VS. ACTUAL CONCURRENCY

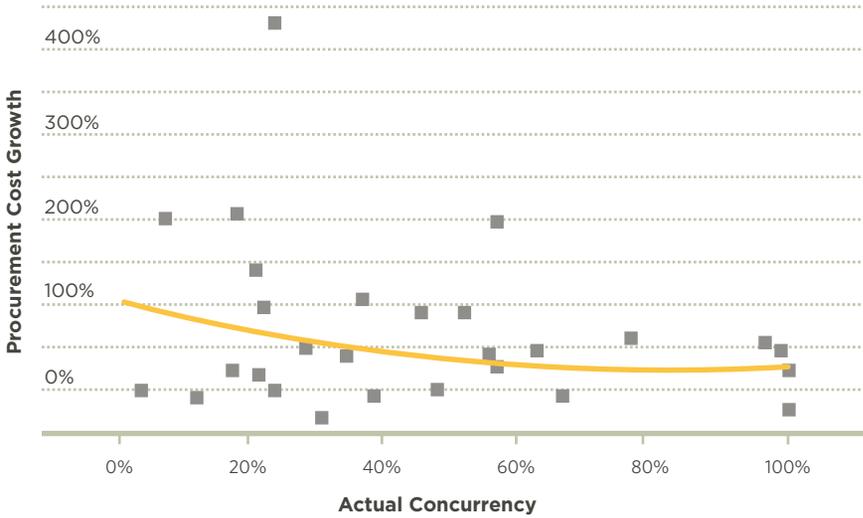


TABLE 5. OLS RESULTS: PROCUREMENT COST GROWTH VS. ACTUAL CONCURRENCY (OUTLIER EXCLUDED)

	Estimate	Std Error	P-Value
Intercept	0.762	0.376	0.054
Concurrency	-0.086	1.684	0.960
Concurrency ²	-0.319	1.521	0.836
Adjusted R ²	-0.040		

we ran the same OLS model again without this outlier. The results follow (Table 5).

Using these data, the model still performed poorly with only the intercept being significant at the .10 level (Figure 7). Further, the fitted line was still concave.

Using the LOESS smooth curving method, we examined the data to see if other relationships could possibly explain the data better than a simple quadratic function. As in the case for planned concurrency, the confidence interval is very wide, indicating that actual concurrency is also a poor predictor of procurement cost growth.

To ensure that the outlier was not a significant factor in these results, we ran the same experiment excluding this data-point. This did not improve the results in any discernible way (Figures 8 and 9).

Again, we used several statistical methods to discover any relation between actual concurrency and procurement cost growth. We specifically reject the notion that actual concurrency has a

FIGURE 7. FITTED CURVE: PROCUREMENT COST GROWTH VS. ACTUAL CONCURRENCY (OUTLIER EXCLUDED)

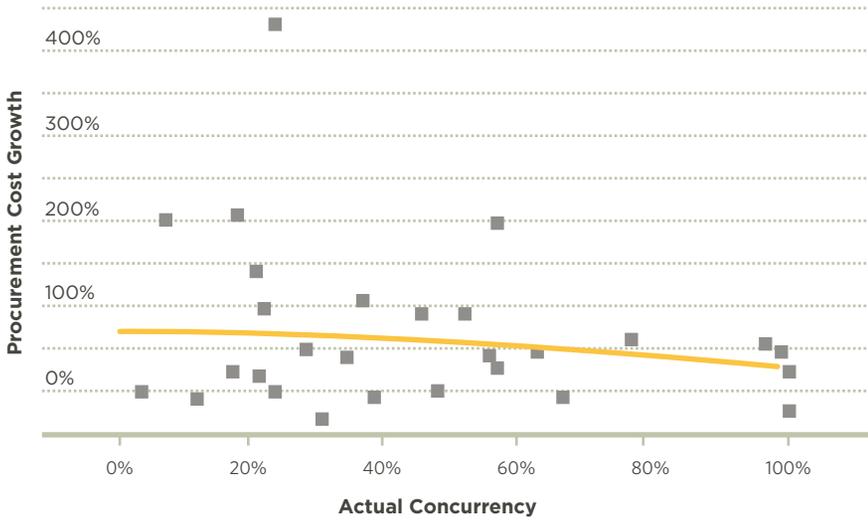
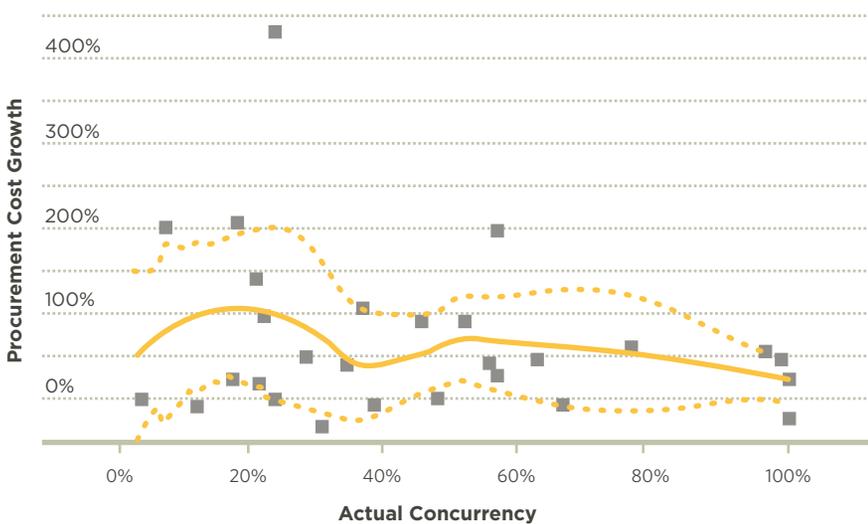
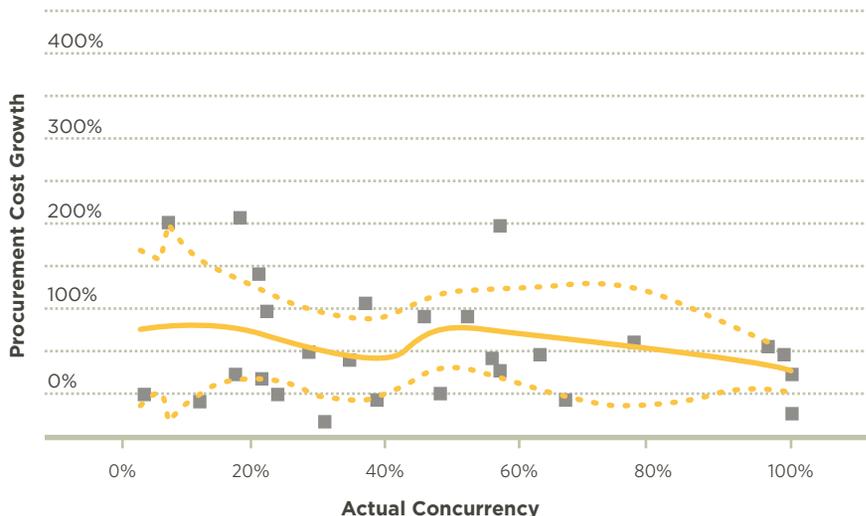


FIGURE 8. CONFIDENCE INTERVAL: PROCUREMENT COST GROWTH VS. ACTUAL CONCURRENCY



quadratic relation to procurement cost growth and find no other polynomial relationship that was consistent with the data. As is the case of planned concurrency, we do see a slight dip in cost growth for those programs with actual concurrency of approximately 30 percent although this is not as pronounced. Thus, our conclusion is that actual concurrency of RDT&E and production funding is not a strong predictor of procurement cost growth either.

FIGURE 9. CONFIDENCE INTERVAL: PROCUREMENT COST GROWTH VS. ACTUAL CONCURRENCY (OUTLIER EXCLUDED)



Conclusions for Procurement Cost Growth

In all cases, we reject the hypothesis that procurement cost growth is related to any measure of concurrency in a way described by a quadratic function. We also found no other polynomial relation that strongly supports the data. While using the LOESS curve smoothing routine on all forms of concurrency did suggest some other possible relation, bootstrapped confidence intervals indicate that any relation between the two is very weak. Thus, even if we accepted the implied curvature, the predictive power of the model for any of the concurrency measures was extremely low. In sum, we found that little if any explanatory power of concurrency by itself affects procurement cost growth. The one result that did stand out was that in the case of both planned and actual concurrency, too little concurrency was actually more problematic than too much concurrency; that is, concurrency levels under approximately 30 percent were associated with higher average levels of cost growth and higher variance as well.

Notably, our results do not indicate that concurrency is never a problem for programs and never leads to cost growth. Rather, it shows that concurrency by itself is insufficient to predict cost growth. Most likely, concurrency leads to cost growth under particular circumstances or in the presence of other factors. What these circumstances or factors are is not clear and should be examined in further research.

Acknowledgments

The authors express gratitude to Connie Custer, vice president, Communications and Public Affairs, Center for Naval Analyses (CNA), and Elana Mintz, senior research publications advisor, CNA, for their invaluable help in the preparation of this manuscript.

Author Biographies



Dr. Donald Birchler is a project director at the CNA. He has worked on acquisition and manpower-related issues for the Navy and Army for the last 7 years. His experience at CNA also includes a 2-year field billet working in support of Command Task Force 72 (CTF-72) in Kamiseya, Japan, where his work on mission readiness earned him the Navy League Rear Admiral William S. Parsons Award for Scientific and Technical Progress. He holds a doctorate in Agricultural and Resource Economics from Ohio State University.

(E-mail address: birchled@cna.org)



Mr. Gary Christle retired from the federal civilian service in October 2000 as the deputy for acquisition management, Office of the Under Secretary of Defense for Acquisition, Technology and Logistics, and is currently with the CNA. In his prior position, he was responsible for DoD acquisition policy as embodied in the DoD 5000-series documents. This responsibility included the role of Defense Acquisition Board executive secretary, establishment of Acquisition Program Baselines, and supervision of the monthly Defense Acquisition

Executive Summary process for monitoring the cost, schedule, and technical status of major acquisition programs. He holds an MBA in Finance from The George Washington University and a BS in Mechanical Engineering, Northeastern University.

(E-mail address: christlg@cna.org)



Mr. Eric Groo joined the CNA in 2008. As a research applications programmer, he has significant technical computing experience, and has contributed to numerous studies supporting the DoD. He is working toward an MS in Statistics at The George Washington University.

(E-mail address: grooe@cna.org)

REFERENCES

- Arena, M. V., Leonard, R. S., Murray, S. E., & Younossi, O. (2006). *Historical cost growth of completed weapon system program* [RAND PROJECT AIR FORCE]. Santa Monica, CA: RAND.
- Congressional Budget Office. (1988). *Concurrent weapons development and production*. Washington, DC: Author.
- Department of Defense. (2006). *Design/build concurrency* [Memorandum]. Washington, DC: Office of the Assistant Secretary of the Navy for Research, Development and Acquisition.
- Department of Defense. (2009). *DoD fiscal year 2009 budget request summary justification* (Green Book). Washington, DC: Office of the Under Secretary of Defense (Comptroller).
- Efron, B., & Tibshirani, R. J. (1998). *An introduction to the bootstrap*. Washington, DC: Chapman & Hall.
- Goldberg, M. S., & Touw, A. E. (2003). *Statistical methods for learning curves and cost analysis*. Hanover, MD: Institute for Operations Research and the Management Sciences.
- Government Accountability Office. (2008). *Defense acquisitions: Assessments of selected major weapons programs* (Report No. GAO-08-467SP). Washington, DC: Author.

ENDNOTES

1. Quoting from the 2009 "Greenbook":

DoD arrives at the figures in this book using inflation rates published by the Office of Management and Budget (OMB) as a baseline. OMB typically bases their rates on Gross Domestic Product (GDP) composite rates, accounting for non-pay factors only. DoD, however, includes pay, fuel, and medical accrual factors in its composite rates. In addition, outlay rates are factored into the final DoD inflation rates. (DoD, 2009)

2. The LOESS bootstrap method is nonparametric, implying that we make no assumptions about the structure of the error term. However, we measure tightness by creating an interval around the original curve that includes 90 percent of the bootstrapped curves. These curves approximate the 5th and 95th percentiles of the true underlying distribution.