

**An Enterprise Model of Rising Ship Costs:
Loss of Learning Due to Time between Ships and Labor Force Instability**

Authors:

Jessica R. Summerville, 15036 Conference Center Dr., Chantilly, VA 20151-3801

(703) 961-3407 (work), (703) 961-0939 (fax), jessica.summerville@ngc.com

Richard L. Coleman, 15036 Conference Center Dr., Chantilly, VA 20151-3801,

(703) 615-4482 (work), (703) 961-0939 (fax), richard.coleman@ngc.com

Bethia L. Cullis, 15036 Conference Center Dr., Chantilly, VA 20151-3801,

(703) 961-3432 (work), (703) 961-0939 (fax), bethia.cullis@ngc.com

Eric R. Druker, 15036 Conference Center Dr., Chantilly, VA 20151-3801,

(703) 227-8540 (work), (703) 961-0939 (fax), eric.druker@ngc.com

Gabriel B. Rutledge, 15036 Conference Center Dr., Chantilly, VA 20151-3801,

(703) 961-3468 (work), (703) 961-0939 (fax), gabriel.rutledge@ngc.com

Peter J. Braxton, 15036 Conference Center Dr., Chantilly, VA 20151-3801,

(703) 961-3411 (work), (703) 961-0939 (fax), peter.braxton@ngc.com

Presented at:

Naval Postgraduate School's 4th Annual Acquisition Research Symposium, 16 - 17 May
2007

**An Enterprise Model of Rising Ship Costs:
Loss of Learning Due to Time between Ships and Labor Force Instability**

Introduction

Since the end of the Cold War, the perceived need for Navy ships has dropped, and so the shipbuilding budget has dropped, as has most of the Defense Department's budget. The Global War On Terror (GWOT) has further stretched budgets. Seemingly coincidental with this budgetary pressure, and perversely aggravating the problem, ship costs began to rise steeply. This paper will demonstrate that a significant portion of the increase in ship costs is not a coincidence at all but is a direct result of decreased budgets in an unavoidable way. We will set aside that ships have grown in weight by about three percent per year since World War II, and that ever-more weapon systems are being put into them, and confine ourselves to discussions of costs rising for ships beyond the increase in "content." Conventional wisdom holds that fewer ships, divided among a fixed industrial base, causes reduced base for overhead and reduced opportunity for the effects of quantity-driven learning; these effects are fairly well understood, and yet cost growth in ships exceeds the cost growth explained by conventional effects. This paper will show that two additional effects are at play, each of which causes individual ship direct labor to rise in a way that existing models and cost estimating tools do not predict. The paper will explain these two effects and conclusively show their impact.

The analysis that follows will demonstrate, via a statistically significant model, two effects: the effect of less demand and the effect of less-steady demand. The former causes stretched-out ship-class acquisitions, increasing time between ship starts; the latter causes workforce instability. The analysis was conducted first on the *Arleigh Burke* (DDG 51) Aegis guided missile destroyer program and demonstrated to be statistically significant (25 data points). In order to be absolutely certain of the analysis, as is the custom on important studies, it was conducted a second, independent time, on a wholly different ship type, the *Wasp* (LHD 1) class of amphibious assault ships (7 data points). These ships are nearly four times as big as a DDG; have a different mission; and were built in the same era, but on a significantly different schedule. The analysis revealed the same result and was again statistically significant. The import of this second trial cannot be overstated; it results in dramatic increase in the significance of the results, from the customary 95% significance level to a 99.75% level. The resultant model is called the Advanced Learning Model. (The three progressively more comprehensive versions, ALM 1, ALM 2, and ALM 3, are summarized in this paper.)

Graphics would greatly simplify and clarify this summary, but the necessity of removing values and avoiding the compromise of proprietary data are delaying this process. The paper as presented will show these very illuminating graphics. The authors apologize for the lack of graphics herein and hope the reader will attend the briefing!

ALM 1 Analysis: The DDG Class

The analysis begins with an examination of Vessel Labor (VL), recurring direct labor, in manhours (MH), by ship. The problem that presents itself, in ship class after

ship class, is that there is a fairly obvious learning curve for the beginning of the ship class, but at some point this curve fails to predict future ships, and VL begins to rise seemingly without cause. If we simply run a learning curve through the initial ships, and then run subsequent curves as the number of ships rises, we will get an answer that seems to show an ever-flattening learning curve. (As we note later, some parties even claim that learning ceases at some point in a ship class; this is a misinterpretation of the perturbing – and, as we will show, perturbed – data.)

This is a learning curve only in a manner of speaking, because, at some point, the “learning curve” does not pass through the data points in any satisfactory way. There is no “learning curve” other than in the sense that any set of points can have a best fit line run through them. What is going on here?

One of the problems is that traditional learning curve theory requires a steady work force building the same product multiple times with no significant interruptions or pauses. Modern ship programs aren’t like that. If these conditions are not fulfilled, there is discussion in the literature of loss of learning, but no closed-form statistically-based method to predict how much learning is lost. One model, the Anderlohr Break-in-Production Model, quantifies the effects of production breaks, but requires expert opinion and is therefore not defensible – it is only useful when mutually agreed to by, e. g., the contractor and the contracting officer. In practice, this agreement is rare. In competition, or in cases of outside scrutiny, for example by the OSD CAIG, this method will almost certainly not stand up.

This problem seemed intractable, with no defensible methodology, until one graph changed everything; that single graph, one of many scatter plots, led to the breakthrough that started the ALM. In that graph, Vessel Labor and Time Between Ship Deliveries (in the final version, start fabrication was used as the marker for ship construction interval) were jointly plotted by ship (this necessitated use of a dual-axis graph). From this one “Rosetta Stone” graphic, it became clear that MH rise and fall as time between ship starts rises and falls. We were not yet ready to assert a closed-form learning model, but we began to suspect that learning was continuing throughout the ship class, and perhaps being lost as well. The driving variable was ultimately expressed in terms of percent of build duration instead of days, as will be explained later. Many painstaking analytic steps remained, but the basic answer had begun to emerge.

The analysis took a number of steps, as we have said. It should be noted that nothing was easy in this analysis, since the data went back many years and, more often than not, was not available in the form in which it was needed. Accordingly, practically every step required data normalization using analytical constructs or some form of “decoding” or deconstructing. These steps were all done after careful consultation with people who were present and in responsible positions at the time of production, and are in accord with conventional wisdom. Most significantly, in the end the analysis works, so the steps are, to a point, justified. We say this knowing that “the end does not justify the means”; nevertheless, when steps that make sense are taken, and the result is both in accord with our ideas as to how it should be, and statistically verifiable, the result acts to

shore up our confidence in the steps. With that warning, we will proceed with the description of the analysis.

1. First, we subtracted Type I change orders as listed in the contracts for each ship. These are changes requested by the Navy for new “content” on the ship.
2. Second, we removed change orders that were unseen originally because they were embedded in the base work of later ships. These change orders were no longer tracked in the tracking system, as the work they contained was “rolled into” the base contract.
3. Third, we took into account “Green Labor,” adjusting the MH to what they would have been had the labor mix been a nominal mix of experienced and inexperienced labor. This third step removes ship-by-ship variability of the labor force mix.
4. Fourth, we adjusted for time between ship starts, as represented by the Start Fabrication step. This fourth step is the keystone to the analysis. The statistics for this step will be shown.

The first three steps were difficult to judge on a one-at-a-time basis, since they were not statistically testable. As we have said, however, we assert that the final step, having been performed after the first three and tested as valid, acts to indicate (and we use this word advisedly) that the first three steps were valid. Had the fourth step failed, we would not know which step(s) to blame, but since it worked, and since all four steps conform with conventional wisdom on how these effects should work, we feel justified in claiming they all were valid. We did not rest on this claim, however, but conducted analysis on a second ship class and ascertained that our findings were neither a coincidence, nor a result of “wishful analysis.”

These steps as described result in a “peeled onion”; by “backing out” the various compounding effects to arrive at what turns out to be the smooth and continuous underlying learning curve. It should be noted that all steps can be reversed to constitute a cost estimation model that relies on accurate plans for change orders, yard manning, and ship intervals, at least the last two of which can be derived from shipbuilding plans.

Analytical Details

Analysis for Step 1, Change Orders. Change orders were available for all ships, but only at the aggregate level. When these aggregated changes were removed, it was clear further analysis was needed. We knew from testimony that, at a certain point in the life of any class, change orders from prior ship contracts are “rolled into” the base contract; from this point forward, these change orders are not tracked separately from the base work. This led us directly to Step 2, embedded change orders.

Before we depart from step 1, however, we should note that *actual* VL MH for change orders are not tracked separately and so only *estimated* change orders are available. There is considerable doubt among analysts as to the value of change orders, some believing that change orders are significantly overestimated (and hence a great source of revenue), others believing they are a significantly underestimated (and hence a

loss). When this analysis was done, we emerged with the conviction that change orders are estimated quite well, else the analysis would never have held up; we urge analysts and decision makers to consider change order estimates to be more accurate than they have in the past.

Analysis for Step 2, Embedded Change Orders. The removal of embedded changes required a great deal of decryption. Using the aggregate changes as a base, we segregated out one-time allowances for rip out (removing completed or in-progress work that the change order replaces), disruption, and one-time change orders. We also had to apply learning to the change orders, because like any work, as units progress, MH decrease. Applying the normal percentage for each of these one-time effects, and applying learning, we “reverse engineered” to decode the change orders for follow-on ships.

With the ship-by-ship deconstructed change orders in hand, we were able to “see” what the change orders looked like and thus could see when the change orders were rolled into the basic contract. We were able to spot a distinct rise in base work and an offsetting drop in change orders on the same ship; we matched this detective work with the memories of senior shipyard personnel and with the contractual picture (changes in the base work could only happen at the start of a contract of multiple ships). The resulting embedded change orders were then subtracted from the appropriate ships; total changes were now removed.

Analysis for Step 3, Green Labor. “Green Labor” is defined as workers with less than 5 years experience; after 5 years, Green Labor becomes “Dry” or “Seasoned” Labor. Green Labor is less effective than Seasoned Labor, based upon a study conducted at Northrop Grumman Ship Systems at their Ingalls Operations, and as indicated by a percentage. (For example, if a worker is 80% effective, then it will take him or her 25% more hours to complete a task, since $1.0/0.8 = 1.25$.) This old study was not much in use, because it was perceived to be less than defensible since it was based upon expert (inside yard) opinion. One of the benefits of our analysis was that we revitalized this model and gave it credibility, as we will discuss. The percentages by year probably differ with type of ship, yard, training programs, etc., as we know that the percentages are different at Northrop Grumman Newport News, based upon a study conducted there. For this analysis, we used a weighted average of effectiveness for workers in the first 5 years.

As in earlier steps, however, it wasn’t just that easy. Data did not exist, for any unit, tracking the portion of the labor force that was green vs. the portion of the labor force that was dry. Again, we needed a construct. By interviewing managers and HR personnel, we ascertained the key variables in the experience-level problem, and built a Markov Chain model for labor experience. We know, for example, that layoffs were LIFO (last-in-first-out), we know the attrition rates for Seasoned and Green labor and we know that most hired labor was Green. We built a Markov Chain model, and it produced the then-current labor mix within about 2 percentage points; the model was seeded with a reasonable start point, so we were confident that the Markov Chain labor mix estimates were close. This model gave us the year-by-year labor mix, and we used this mix to

adjust to a notional (most common) labor mix, thereby nearly (because this is an approximation, we can only say “nearly”) eliminating the effects of labor force mix.

As a side note, in past briefings we often said that the effects of labor mix are not dominant in this problem, because labor mix is fairly stable; after Hurricane Katrina, however, when we looked at reconstituting the labor force with this model, we saw that when labor swings considerably, the effect becomes quite big. In a small-business-base case, considerable cost impact is in fact caused by labor swings. While average Green Labor is less efficient than Seasoned Labor, first-year labor is *quite a bit* less efficient, so an unstable work force that is constantly ramping up and down will result in considerable loss of efficiency. This all said, however, our model demonstrates that the potential for cost impacts are greatest in the last variable: loss of learning due to interval effect.

Analysis for Step 4, Loss of Learning due to Interval Effect. The final step is the most important step, both because the prior steps were either understood or in practice, and because this step acts to verify, in the way we have described, the earlier steps. In this step, we conducted a linear regression of the difference between learning as observed and learning as expected. Though this was difficult in derivation, and had many false steps, the result is intuitive.

Two major assumptions underlie our method. These two assumptions are in accord with current learning curve theory, but deserve mention. First, we assumed that learning is constant and incremental, and proceeds from one ship to the next. This means that as other effects occur, the learning curve factor is still at play. We have been surprised that others have asserted that learning stops after nine or ten ships¹. This is not in consonance with anything in the literature of learning, and we believe that the results of this analysis bear out our belief that learning continues through all units. Getting ahead of ourselves for a moment, we believe that at some point in a ship class there is surely a considerable change order effect, and the attendant slow-down in production as the changes are engineered, which causes loss of learning. This is not to say learning stopped, but rather to say that other effects are causing a rise in ship costs. Some of these effects are controllable, and all are, to some degree, predictable. Second, we believe that should there be a loss of learning, (1) the learning begins anew, arguably at the same rate; and (2) it continues from the new point with later ships. Both things occurred in the two classes we analyzed. Said another way, loss of learning acts as a reset of the curve; later units will never revert to the *old curve*, but rather proceed from the reset unit following a *new curve with the old slope*. This last belief means that we must measure loss of learning from the effective unit of the prior ship, moved ahead one unit. Thus we must correct out all other effects, since those effects (or at least those we encountered) are

¹ Why Has the Cost of Navy Ships Risen? A Macroscopic Examination of the Trends in U.S. Naval Ship Costs Over the Past Several Decades, Mark V. Arena, Irv Blickstein, Obaid Younossi, Clifford A. Grammich, RAND, 2006

either transient, or fade on their own timescale. Green labor appears to be one such effect; as the labor force matures, workers approach a basic efficiency.

This discussion should by now have alerted the reader as to how complex this analysis is, and how dependent it is on each step. Get one step wrong – we know, we did it many times as we groped through it – and the whole thing will not work. One can “see” the effects; to prove them analytically, however, requires considerable precision and care. It was particularly confounding that for many of the ships in the DDG class, the Green Labor and interval effects tended to counteract each other. This was because when time between ships stretched out – causing loss of learning due to interval effect, which tended to drive VL MH up – there was less work in the yard, so manning was reduced (largely via the firing of Green Labor), resulting in a more efficient labor force, which tended to drive VL MH down. At other times, such as at the onset of major changes in a class, the effects can reinforce each other, because there has been a gap, but work then suddenly builds up in the yard. It was not until *both* effects were quantified that this complex interaction could be fully understood. This interplay has doubtless been a major complication in understanding these effects.

When conducting our analysis, we first determined to what point on the learning curve each ship should have progressed, based on the previous ship. This was conducted iteratively from the second ship to the final ship. We then regressed the difference between the expectation and the realization, expressed as numbers of units (or fractions thereof) lost, against days between ships. The regression was significant at the 95% significance level (the “p-value” was 0.0106). The coefficient of determination (called r^2) was 0.4324. While this is not an overwhelming r^2 , it is a respectable one, and in any event, the determinant of validity is significance, not r^2 .

This regression shows that loss of learning is linear with days between ship starts. As ships are built less and less frequently, learning seems to “flatten”. The underlying learning curve slope, however, continues to remain in force. The ships do not cease to learn, they lose learning. This is an important distinction. If the interval shortens, as it occasionally does, the MH drop faster, reverting to the original learning curve slope if they are close enough together.

There is much to discuss about this finding, but first we should observe one very important fact. The underlying learning that emerged for a ship class of 25 ships was the same learning realized by the first nine undisturbed ships. After the ninth ship, cost began to rise, so much so that the cost of the thirteenth ship was nearly equal to the cost of the fourth ship. This may not be surprising, after all our descriptions of the effects that change costs. The point is that most of these effects do not happen in the early ships because they tend to be produced at steady, close, intervals; this proved true for both DDGs and LHDs. The rise that began in DDGs after nine ships (and in LHDs after four ships, though, like DDG, about five years into the program) was caused by the combined effects of change orders, embedded change orders, labor inefficiencies caused by the coincidence of a new class of large ship in the shipyard, and loss of learning due to the delays caused by the incorporation of a major change (DDG 51 Flight IIA). We are

confident that this is why so many analysts in shipbuilding cite cessation of learning, but we wish to say this again for emphasis: learning continued in this ship class, as far as we had data; out past the 25th ship. The reader will see that this was true in LHDs as well. Costs rose due to predictable (or at least expected) effects, and, though learning was lost due to slowdowns, it continued on from each point “as regular as sunrise.” To say learning stopped is to miss the point, and to forsake the opportunity of prediction.

We also found that there appeared to be an “ideal” interval between ships that causes consistent learning. But as our understanding deepened, we began to appreciate that a ship class can seem to take any learning curve, depending on the predominant interval between ships. We will also show the days between ships was a poor way to display the interval, and we later changed to percent of construction duration; this change promises to open a whole new possibility. We will we discuss both things further below.

The Need for Confirmation: The LHD Class

The DDG model has been laid out. In many important analyses, it is customary to do second trials (e.g., in medical studies). The second test guards against spurious results. We will describe the meta-analysis later. In this case, though, a confirming analysis was needed more than in usual cost analysis cases because of the many data normalization steps that preceded the regression. We have asserted that the ultimate regression sustains the earlier steps, but we cannot escape being nervous that we may have driven towards a conclusion, inadvertently forcing the conclusion. This is natural and must be guarded against. Accordingly, we immediately set about repeating the analysis for a second ship class.

We did exactly the same analytical steps as described for the DDG and arrived at exactly the same conclusions. We found that the LHD class had done just what the DDGs did. The LHD step 4 regression for the Interval Effect was significant at the 95% significance level (the “p-value” was 0.017). The coefficient of determination (called r^2) was 0.887. This r^2 was higher, due simply to the fact that the LHDs fell into two groups, 3 units lost no learning and two lost all learning; once again, the determinant of validity is significance, not r^2 . The only difference was that the rate at which the LHDs lost learning was slower (the slope of the line was flatter); and the “cross-over” point, the interval at which there was no loss of learning at all, was longer in days (the x-intercept was greater). In a manner of speaking, the “time constants” were longer. What was more interesting, however, was that LHDs had on two occasions lost learning very dramatically; the class had reverted to a first unit (lost all learning) twice, at LHD 5 and LHD 7. The underlying learning for the LHD class was quite similar to the DDGs, only a few percent flatter. As a last caution, we know that two LHDs lost all learning, and reverted to the first unit. We thus cannot know whether the interval effect might be stronger than we have evidence for, because these two ships were capped at the first-unit level. Accordingly, we know that the LHD interval effect is at least as strong as we found, and quite likely stronger.

This brings us to an important point: learning due to interval, according to our model, occurs in units with time. For mature classes, say after nine ships, like the DDG,

the loss of a few units is not a large number of MH, because the curve has flattened out by this point to a few percentage points per unit. For less mature classes, like the LHD, however, learning is still occurring at a steep rate, and the loss of a unit may be more on the order of 10% or more. In either case, if the Navy or the Coast Guard have negotiated an FPI contract, because “follow ship costs should be well known,” and there is a delay, due to negotiations, design activity, or lessened SCN budgets, the rise can erase all profit. This can be even more of a problem if the method of EAC computation does not catch an insidious rise in preceding ships, perhaps also due to factors that do not lie at the feet of the builder, and which will cause the next ship to rise for additional (additive) reasons.

Meta-analysis and Mutual Confirmation

To summarize, the DDG model is, as far as statistics can take us, valid. The LHD model is also valid, and we can use either. The significance of the entire analysis is one minus the square of the significance of each: $1 - 0.05^2 = 1 - 0.0025$ (significance is 99.75 %.) This well-known technique is called meta-analysis. Most importantly though, the LHD model represents a second ship class and was undertaken to confirm the DDG model. Taken together, the LHD and DDG models are much stronger than either is alone because the models are mutually confirming. Taken alone, the DDG analysis has the weaknesses that “first models” commonly have: although reasonable, the adjustments were taken to drive towards a smooth learning curve. Statistically, this amounts to an uncredited loss of “degrees of freedom.” Alternatively, a hostile view can arise that the data was “cooked.” The LHD model alleviates this concern. Since the steps taken with DDG were replicated in LHD and the same result was obtained, our analysis was neither dumb luck nor manipulation. Further, the LHD class had only 7 ships completed, so taken alone, the LHD analysis lacks data across the full spectrum of interval length. The DDG model alleviates this concern with 25 ships and a wide range of intervals.

ALM 2: The Enterprise Model

We built an enterprise level model that instantiates the Advanced Learning Model for an entire shipyard complex, namely Northrop Grumman Ship Systems. We call this a forward-looking model, because it projects effects given the results of the prior analysis, which we call ALM 1. This model, called ALM 2, brought about three significant new capabilities. First, the model allows for the quick determination of the impacts of the effects described above on either a future ship class or the remainder of a class. It also allows the input of all ships in the enterprise. Second, it computes the labor profile, at the gross level, for the yard, a capability which formerly took weeks. The detailed trade-by-trade labor profiles are still done by hand, and much mischief can arise at the trade level that the enterprise model is blind to, but large what-ifs can nevertheless be done quickly as a first-order approximation. Third, this enterprise model has the ability to model cross-ship effects and, since it is in closed form, feedback effects on labor.

When a ship cost estimate is done, the apparent labor needs are determined for a notional labor profile, the profile that was in effect at the time of the historical data. The model first needs to compute the efficiency of the labor force in the yard where it is being built based on the mix of Green Labor. The model then will compute interval effects based on the preceding ship and the interval before the ship in question. At this point, the model will adjust the need for labor. This adjustment changes the labor mix in the yard,

because new labor must be hired above (or below) what was expected, due to the effects of these two variables. This produces a feedback loop which happily converges fairly quickly. Likewise, the model, having all ships in the yard loaded in, quickly computes the effects of the new labor on all the other ships in the yard as the labor force swells or shrinks, and the associated feedback loops, and does so quickly.

The enterprise model brought out an understanding that is worthy of some discussion. As we began to use the model to predict ships, we needed to develop a good prediction for change orders. When we did the analysis, using the carefully deconstructed change order data for these two ship classes, we were taken aback by the size and the implications of what we found. We found that change orders were large on the first ship, which we had known, and averaged 12.30% of base costs (with a wide variability). We knew that by definition change orders would be large at Flight change or at major modifications, so when we found that changes were 8.62% of cost it was sobering but not surprising. What surprised us was that change orders on the rest of the ships, ships that were neither first ships nor major modifications, was 2.97%, with a fairly clear pattern of rising over the class but virtually always exceeding 1.2%. This change order pattern meant that by the 25th ship, with two major modifications, the labor cost of the DDG change orders had become equal to about half that of the base ship. On the LHD change orders have risen to be about one-quarter of the cost of the base work by the 7th ship. This finding showed us that change orders are a much larger part of the cost of follow ships than we had realized, even though we knew the conventional wisdom that changes are a big factor in ship costs.

It should not be concluded that change orders are inherently bad. Change orders are at least partly a response to lessons learned, changes in technology, and in the threat. That said, they bear careful monitoring due to their size, which begins to rival initial cost in their claim on scarce SCN. Change orders also have a second, hitherto-less-well understood effect on ship costs. It is quite likely that change orders act both as chicken and egg in the ship-cost scenario. As budgets drop, interval increases, allowing an ever increasing pressure for change, be it due to an industrious engineering community, technology changes, or requirements changes. Conversely, changes cause design delay, increasing intervals, sometimes considerably, as they do in major modifications. Thus, changes are already a considerable factor, and have the ability to become more of a factor in tight budget years as well as to drive the other effects to become larger as they delay ships and contribute to cost rises which destabilize labor costs.

Implications

At the outset of this paper we described how ship costs have suddenly seemed to go out of control just as budgets have dropped, exacerbating the problem. There are no doubt many reasons for such a problem, some of which are outlined in the RAND study already referenced. Many of the reasons cited in the RAND paper, such as inflation, and commodity volatility, have conventional means of prediction, however flawed the results may be. Our paper, on the other hand, puts forward three new effects that were hitherto generally ignored analytically, or discounted by decision makers. As we have already said, change order estimates in general turned out to be quite accurate; without this

accuracy, the analysis would have been impossible. The “lost” change orders are unarguable in principle, and although the exact values estimated could be quibbled with, they are certainly “generally right.” The illumination we have brought to this issue will help explain past cost increases, and closer attention to tracking change orders going forward should be a big help to future cost estimation. The Green Labor model we described was in use at one Northrop Grumman shipyard, and is now in use at them all, although its impact is debated and occasionally somewhat discounted. We believe that Green Labor’s part in this analysis will go far to bolster the basis for the computation of this effect. The interval effect has always been part of shipyard lore, but has never been quantified. Even now, shipyards are struggling to understand the model, and to overcome past “explanations” for rising costs, such as labor inefficiencies (true, but overstated for the lack of an agreed-upon model); one-time events (again, true, but arguably overstated for lack of sufficiently reliable discrete cost tracking); and disruption effects such as were alleged to have happened twice on DDG and once on LHD (we found no significant differences that needed explanation, so the supposed disruption, although appealing intellectually, was absent if you believe our analysis).

The net of the effects in this paper is extremely large and almost all was cost growth. The newly explained cost growth in labor, resulting from embedded change orders, Green Labor and loss of learning due to interval for all ships from 12 through 25 of the DDG is a total of 21%, and for all ships from 3 through 7 of the LHD is a total of 19%. This is a staggering amount of cost to have been hitherto unexplained. Of the newly explained cost growth in labor, 8.8% (almost half) of the growth in DDGs was from loss of learning due to the interval effect; for the LHD class, 14% (almost ¾) was from loss of learning due to the interval effect. Most of the rest of the unexplained cost growth was due to “lost” change orders. The DDG interval effect was less of the total partly because the class was more mature, and further down the learning curve, thus units lost were smaller in MH than for LHDs; conversely, for LHDs, change orders were less of the total because there had been fewer units for them to have accumulated. Green labor was the least of the effects, and tended to be pluses or minuses, as labor waxed and waned. We thus believe that a large mystery has been solved, and that we have provided the ability to avoid similar large errors in the future. As with any cost estimating technique, of course, the ability to project will depend upon how well we can determine the inputs.

This model has the capability to account for many large and unpleasant cost growth effects, especially for modified ships that are characterized by long intervals, large change order packages, and large swings in labor as the class goes back into production. We thus believe that the model will be invaluable in better predicting costs. We hope it will be helpful in reducing mistrust and tension between the Navy and the Coast Guard on the one hand, and the shipyards on the other, as well as between the Congress and the services. We believe the model goes a long way towards explaining cost rises in terms the various parties can better understand as unavoidable in the past, and more predictable in the future.

There are policy implications to these findings as well. We believe costs can be better estimated now, but beyond that, they can be better controlled as well. One simple exercise involved changing the acquisition profile for Large Deck Amphibious ships. By optimizing the timing of ship acquisitions, we found that labor costs could be reduced by almost 6% in constant dollars – and this was with only a single ship class to work with. In addition, in order to be successful, the Navy’s current cost reduction strategy of re-using existing hull forms for new classes of ships must obey the constraints illuminated here. Change orders must be suppressed to the degree that they can, intervals between ships must be thoughtfully timed, and demand variance that results in workforce instability must be avoided. We note that this is consistent with an enterprise view of shipbuilding: the Navy making acquisition decisions that support its private industry partners in their efforts to contain costs. On the other hand, for the Navy to pursue acquisition strategies that ignore the demonstrated effects of change orders, loss of learning due to interval, and Green Labor -- and at the same time to press for cost reduction – is inherently problematic. It is further setting the enterprise up for failure.

Next Steps

We hope to take this analysis to other ship classes. We further hope to take it beyond ships to other commodities. As strange as it may sound, application to satellites is quite promising, as they are, like ships, complex objects with long development spans involving integration of many components; operated in demanding environments; and procured in small “batches” (usually one or two). There is no clear reason why this methodology should not work for other sorts of production. Indeed, the problem may be considerably easier since labor force changes should be less of an effect in items of shorter duration and higher production rate. Change orders may also be easier, because data will not be so old in calendar terms and since the shorter times lend themselves to fewer changes.

We would like to point out that so-called “rate effects” in learning have been elusive. We know of no demonstration of a statistically significant rate effect. We suggest that perhaps the “rate effect,” which is commonly introduced as a second term in the learning equation, may simply be an incorrect portrayal of the interval effect we have demonstrated. Rate effect uses the number of units to be produced in a given year and lowers costs as the quantity rises and raises costs as it drops. This acts to reduce cost when units are built closer together (more production units per unit time), thus mirroring the sense of the interval effect we found, but using different mathematical expression to model it. Statistics have the unfortunate limitation that a hypothesis must be formulated for the test. If the hypothesis is structured in an incorrect form, the model is unlikely to be significant, so the hypothesis that the model is invalid cannot be rejected. Many good ideas founder on this shoal, and this failure of an “almost right model” or a model that acts in a sensible way, but is formulated a bit incorrectly, frustrates the analyst and leads to unreasonable distrust of statistics. We look forward to investigating the rate effect relative to the interval effect, and believe that the interval effect will replace it in practice.

As we have alluded to, we have already begun a re-formulation of the model. We have changed the expression of the model for loss of learning from days to percent of

construction duration. This is merely a change in scale and the results do not change for either model. What does change, however, is that when the model is expressed in the new way, we discover that the LHD and DDG graphs of loss of learning overlay each other. In the “days form,” we had no way to move to another class, because we had no way to interpolate save linearly, and we were wary of this assumption. Worse yet, when applied to a much smaller ship, like a frigate, or a much larger ship, like an aircraft carrier, the extrapolation became irrational. This amounted to *reductio ad absurdum* of linear extrapolation and so cast doubt upon interpolation. When expressed as percent of interval, however, the result, although not yet tested, is appealing and intuitive. This version of the model is called ALM 3.

We are not shipbuilders, and further, we are always reluctant to speculate beyond what we can prove. Nonetheless, one of the reasons we are inclined to believe this “percent form” is that the nature of work on a large ship changes less quickly during production than for small ships. For example, shell plating and other large-object welding lasts a similar percentage of construction time as it does for smaller ships, hence we speculate that the opportunity to “learn” lasts longer, and time to lose it takes longer in days, but not in percent.

This change in scale has taken us to another point. We are in the midst of trying to determine whether the percent model will lead us to a long-sought capability: the capability to predict the underlying learning curve for any ship knowing only the expected construction duration and the planned schedule for the class. We anticipate having a preliminary version of this in time for the symposium, or to have learned that this is un-doable. If we can achieve this, we feel the results will be every bit as important as the loss of learning model. It would also increase the importance of both accurate schedule estimating and stable program schedules, as these are key inputs to the learning curve determination, which in turn is the single biggest driver of labor cost. It is premature to make a claim of broader application, especially since the ship portion is as yet undone, but we cannot stop our minds from racing ahead. We hold out hope for this model being useful in other commodities. We have, on an informal basis, noted that the interval for Global Hawk is similar to that for DDGs in percent, and that the learning curve is similar. We raise this point in order to excite the reader’s interest, and to indicate that the possibility of further application is real.

Conclusions

This paper is short, and so our conclusions should be short as well. We have shown that a careful treatment of change orders, labor force mix, and time between ship starts has enabled us to demonstrate – we believe for the first time – a closed-form and statistically rigorous treatment of loss of learning. The analysis has validated an old labor mix model and shown it to have considerably credibility; and has unified change orders, labor profiles, and loss of learning in a way that substantially changes ship cost estimation. We have already created an enterprise model and have hopes of being able to extend the model to other commodities. We will perhaps even be able to predict learning curves for ship classes, a capability that does not yet exist.