SPECIAL REPORT

ASPI

November 2011 — Issue 42

ASPI AUSTRALIAN STRATEGIC POLICY INSTITUTE

Three views of risk

Selecting and acquiring military equipment

by Mark Thomson, Andrew Davies and Chris Jenkins

Selecting and acquiring equipment for the Australian Defence Force is difficult. Military technologies are often expensive and complex, and almost inevitably there are tricky problems balancing capability, price and market availability.

No simple rule exists to assure value for money in every situation. Sometimes it will be appropriate to maximise effectiveness—measured against criteria that are themselves somewhat subjective—while at other times it will be best to pursue the least costly or most readily deliverable option that meets minimum performance criteria. And sometimes it's hard to know in advance exactly what trade-offs are even possible. The three papers in this report examine different perspectives of the problems faced by the decision makers in Defence and government.

The first paper, by Mark Thomson, looks at the application of commercial techniques for evaluating and managing risk in defence projects. Some of those tools readily transfer to the military capability business, while others need careful application when dealing with a public good like defence, and especially with the frequently unquantifiable benefits that form one side of the 'value for money' equation. The paper concludes that, while qualitative judgments are inherent in all but the most trivial of defence acquisition decisions, the techniques of commercial investment analysis can assist decision makers to sort through the options.

The second paper, by Andrew Davies, concerns itself more with the management of risk within defence projects. Given the technical complexity of much modern military equipment, development projects are necessarily risky ventures. Because technical risks aren't always obvious when projects are conceived, it shouldn't be a surprise when a portfolio of developmental projects contains some that experience setbacks or even fail. The best chance of success requires a rigorous project management approach implemented by a cadre of suitably qualified and experienced project managers and engineers—all of whom are in short supply in Defence and much sought after in other parts of the economy. The paper concludes that the government should be judicious about the number and size of developmental defence projects it takes on.

The final paper, by Chris Jenkins of Thales Australia, provides an industry perspective. Defence industry is an important component of Australia's defence capability, and maintaining its health should be a consideration when formulating defence policy. This paper explains that a long-term view is necessary when evaluating the pay-off from investment in industry capability. In particular, a 'boom and bust' model of procurement makes it very difficult for industry to maintain the capabilities required to manage complex projects.

1 Balancing cost, benefit and risk in defence capability acquisitions

Mark Thomson

Decisions about large defence acquisitions can be fiendishly hard. Not only are hundreds of millions of dollars often at stake—along with the nation's defence—but the competing options rarely allow a straightforward comparison. Instead, it's usually a case of comparing apples with oranges in the context of differing risks, costs and schedules. And unlike purely financial investments, there's no single parameter such as rate of return to help isolate the best option. As a result, qualitative judgments are inherent in all but the most trivial of defence acquisition decisions.

Nevertheless, the tools and concepts of commercial investment analysis can be useful in decisions about defence acquisitions. If nothing else, a commercial perspective highlights the importance of managing rather than minimising risk. What follows explores the application of commercial techniques to defence acquisition decision-making in three parts. First, the problem of acquisition decision-making is described. Second, the basic building blocks of investment analysis—cost, benefit and risk—are examined in some detail; each has special characteristics in a defence context. Finally, the application of investment analysis techniques to defence decision-making is discussed.

Decisions

Commercial investment analysis involves comparing alternative options for the allocation of capital. In the resources sector, it might involve the comparison of alternative mining projects in different countries. In the civil construction sector, it would usually involve deciding which tenders to bid for from those on offer. And in the finance sector, it would come down to comparing alternative investment options—equities, bonds, managed funds and so on. In each case, the fundamental question is the same: how best to allocate financial and human capital among the options.

A similar problem arises constantly within Defence. The relentless ageing of military equipment coupled with the emergence of new technologies demands a continuous program of modernisation. Because national defence is capital intensive, the acquisition of new equipment is a major activity. Each year, around 20% of the defence budget is spent on new equipment, and around $150 billion in defence projects is scheduled for approval over the next decade.

In many ways, Defence's portfolio of acquisition projects is analogous to the portfolio of investments that a commercial firm might have. One key difference is that the defence portfolio aims to shape the evolving structure of the defence force rather than simply deliver a return on investment. For this reason, commercial techniques have limited application in deciding the balance of defence investments: aircraft versus ships, say, or tanks versus submarines. That's not to say that economic thinking has nothing to add in this area—it certainly does—but that's a story for another day.

Where commercial techniques come to the fore is in deciding between different options to fulfil a single capability aspiration. For example, the choice might be between several different ships or between several different aircraft, or perhaps even between different equipment types that aim to fulfil the same broad purpose, such as artillery versus mortars or bombs versus missiles. In such cases, there are usually a range of options being considered, each with different costs, benefits and attendant risks.

Costs

The initial purchase of an item of defence equipment invariably comes at a substantial cost to taxpayers. Ensuing personnel and operating costs can then amount to an additional two or three times the purchase price of the asset over its life. For this reason, sound decisions about defence acquisitions must take into account the total through-life costs of the competing options.

Every dollar spent on defence has an opportunity cost—forgone private consumption (via lower taxes) or alternative public services such as health, education and welfare. Each dollar can only be spent once. In principle, the same logic applies within Defence because every dollar spent on one item can't be used for another defence purpose. In practice, however, the internal opportunity cost to Defence is softened by the prospect of going back to the government for more money. Because the government is ultimately responsible for the nation's defence, it can't avoid bearing the consequences of Defence planning or spending beyond Defence's allocated means. Moral hazard abounds.

Spending on projects rarely occurs in a single payment. More usually, costs are spread over a number of years and often extend beyond a decade. And because of opportunity costs, a dollar spent today isn't the same as a dollar spent in a year's time, even after accounting for inflation. All other things being equal, people have a preference for benefits sooner rather than later. Consistent with that preference, individuals usually demand compensation for deferring benefits to a later date. This takes the form of interest payments when people defer consumption by saving their money. Moreover, because a dollar invested in capital allows greater output to be derived from other inputs, deferred consumption can be rewarded through that greater output. Similarly, many individuals are willing to pay interest to borrow money in order to bring forward a benefit—be it a car or a holiday. As it is for individuals, so it is, too, for society. Opportunity costs incurred by the community today are felt more acutely than those deferred into the future.

One way that economists take society's time preference into account is by discounting the value of future spending relative to today's by calculating what's called the 'present value' of a multi-year expenditure profile. The details needn't concern us, but it's basically compound interest working in reverse. Although various discount rates can be used, they tend to be based on the long-term bond rate. Indeed, so long as the Commonwealth remains in debt, the bond rate effectively captures the cost of further borrowing. That cost is lower than the equivalent cost to the private sector, because the government can fund repayments through taxation, which makes it a better credit risk.

An example is useful. A present value calculation using a discount rate of 6% would place the same value on $1,000 spent today as it would on $1,281 spent in five years' time or $1,745 in ten years.

Present value estimates are not routinely used to compare competing defence acquisition options. In part, this is because it's usually not possible to quantify the benefits of military acquisitions in monetary terms (more on this below). In at least equal measure, present value calculations are rarely used because defence planners are usually more concerned with fitting a portfolio of projects into an overall predetermined multi-year spending envelope—a non-trivial packing problem—than with reducing the present value of a single acquisition.

Nonetheless, if the prospective outcomes of two or more options are similar in terms of capability levels and schedule, a comparison of the present value of costs is useful. At the very least, such an analysis takes account of the cost of forgone alternative investments elsewhere in the portfolio of projects. Put simply, options with lower net present values allow the earlier commencement of other projects within a given overall funding allocation. Alternatively, the overall funding provision can be adjusted so as to reduce the economic opportunity cost to society.

In the absence of a present value analysis, the costs of competing defence options are not being weighed in line with society's valuation. As a result, there's a risk that society will bear higher opportunity costs than it should.

Benefits

In the financial world, the benefits (returns) from investments are denominated in dollars. This allows the overall outcome of investments to be calculated by subtracting the present value of costs from the present value of benefits. Such a process is usually termed cost–benefit analysis, the result of which is a 'net present value' for each of the competing options.

Defence projects don't lead to benefits that can be measured in dollars. Instead, they deliver military capabilities that may or may not be used at some point in the future. What's more, it's usually not possible to assign a single measure, or even a set of measures, to accurately capture the benefits sought. Consider, for example, a naval ship. The efficacy of a ship depends on myriad characteristics, such as range, speed, sensor reach, weapons effectiveness and habitability for the crew. The relative importance of each of those and other characteristics depends on how, where, when and against whom the vessel will be employed.

Of course, it's possible to assign a weighting to each of the factors and calculate an aggregate single measure of effectiveness, or perhaps a set of aggregated measures of effectiveness. Such approaches are routinely used in tender evaluations when comparing offers from suppliers. Sometimes the measures are quantitative from the start—speed and range, for example—while at other times there's no alternative but to assign a quantitative measure to a qualitative assessment. In all but the simplest of cases (such as comparing the effectiveness of two similar munitions), there are in general many subjective judgments underpinning a measure or set of measures of effectiveness.

In some cases, measures of benefit can be derived by examining trade-offs. For instance, if the aim is to maintain defence capability, reducing the availability of an asset by x units might require a matching increase of y in some other asset or assets—producing an implied valuation. Even when such methods aren't useful, there are techniques, known as multi-criterion analyses, that can guide the overall valuation of different options. While such approaches are open to manipulation through the selective formulation of measures, if properly used they can impose useful discipline on the comparison of options.

If a measure or set of measures of effectiveness is available, the relative cost-effectiveness of options can be compared by dividing the measure(s) of effectiveness for each of the options by their respective present value costs (although this needs to be done with care, especially if a fixed budget is being allocated, as allocation on the basis of the resulting ratios may prove suboptimal). In practice, such an analysis is at best only an adjunct to broader considerations. For example, the relative ability of competing options to deliver equipment prior to the retirement of existing assets, or before the expected development of new capabilities by potential adversaries, needs to be taken into account. In general, factors like these do not readily translate into measures of effectiveness. Nonetheless, at least ranking options in terms of different facets of effectiveness is a vital discipline, as it can be used both as a point of comparison to the rankings of relative costs and as a way of assessing just how much gain there is in moving from less expensive, less capable options to more expensive, more capable options.

Similarly, supposed ancillary benefits from defence acquisitions, such as employment generation and industry development, are usually best considered in parallel with any cost-effectiveness analysis. Except when it's necessary to nurture industry capacity to deliver through-life support as part of an acquisition project, these sorts of factors tend to be driven more by the demands of vested interests than by rational economics. This strengthens the argument for quarantining such factors from the rigorous comparison of options—except that the additional cost of providing such benefits should be made transparent.

Risks

Defence projects tend to be risky. Costs often increase far beyond initial estimates, schedules routinely extend years beyond planned delivery dates, and technical performance is too often below the original level sought. While similar problems routinely occur elsewhere in the public and private sectors, they tend to be particularly acute in the defence arena. This is true not just in Australia but around the world. The underlying reason is that defence equipment often has to push against the leading edge of technology to provide an advantage over potential adversaries. As a result, costs and risks are driven higher in the pursuit of ever-diminishing performance improvements.

The inherent risks in defence projects are exacerbated by a systematic tendency to underestimate costs and lead-times, at least in the early stages of projects. In part, this reflects the intrinsic optimistic bias of defence planners—something observed elsewhere in the private and public sectors. And when the overrun and delay chickens come home to roost, there's a tendency to respond by delaying the project rather than cancelling it, thereby rewarding the optimistic bias. Moreover, there's no doubt at times an element of unrealistic optimism on the part of project advocates (in both industry and Defence) who are eager to lock the government into their preferred course of action. Unavoidably, it's the government rather than Defence that ultimately bears the consequences of poor decisions.

Whatever the causes, risk is a constant factor in every acquisition proposal. Sometimes relatively low-risk options are available, as with the recent off-the-shelf purchases of C-17 transporters and F/A-18 *Super Hornet* fighters. At other times, some degree of developmental (and therefore risky) work is inescapable—for example, when backward compatibility with legacy systems is needed or when the market only offers options that are still under development. ASPI has long argued that greater use should be made of proven off-the-shelf equipment than currently occurs. Nonetheless, the question arises as to how risk should be factored into acquisition decision-making.

A risk is an adverse event that has a likelihood and consequence. In general, there's a distribution of such events, each with a specific probability and consequence. In financial analysis, actual mathematical distributions are used to represent the likelihood of a parameter such as an interest rate or commodity price having a certain value at some point in the future. When that level of detail is available, average returns can be estimated. In effect, the returns from various parameter values are added up, weighted according to their likelihood. Taking the full distribution into account, rather than focusing on the 'most likely', is especially important when the distribution is skewed—for instance, when there's a small probability of a very large loss that could wipe out the entire portfolio. Failure to take account of that probability—say, because the analyst was concentrating on the average or mean of the distribution—could seriously underestimate the risk.

Unfortunately, this sort of detailed approach is rarely employed for managing defence acquisitions. That's a pity, since there's undoubtedly more that could be done. Just as it's possible to use historical data to construct the probability distribution for future interest rates, the inherent risk in many technology areas is well documented—for example, industry benchmarks for software development are widely available. Equally, Defence's own experience with hundreds of projects provides ample data for anticipating the distribution of project schedule and cost performance.

For some years now, Defence has used a 'project maturity' score to measure how ready projects are to proceed. Roughly speaking, project maturity measures the extent to which risks have been retired in a project. The project maturity score covers seven areas: schedule, cost, requirement, technical understanding, technical difficulty, confidence in the commercial world's ability to deliver and operational/support. Each area is scored out of 10, resulting in an overall maximum score of 70. A score of 1 in an area reflects high risk, and a score of 10 represents low risk. As a rule of thumb, a score of about 20 is expected at first pass and a score of 35 or so at second pass. The implicit judgment is that projects with lower scores carry too much risk to proceed and therefore require further work to retire risk.

The project maturity benchmarks set for projects at the milestones of first and second pass reflect the 'risk appetite' of Defence and hopefully that of the government. Many commercial firms make similarly deliberate choices about the extent of risk they'll allow themselves to be exposed to. The setting of such thresholds is one of the tools that can be used to manage risk. Another is the provision of contingency in project funding and scheduling, both of which are routine in defence projects and elsewhere. The options of hedging and insurance are also available in the commercial world, but are rarely useful or available in a defence context.

Defence's current two-pass process aims to actively manage risk by investing money to better understand and where possible to retire risk between first and second pass. In that process, it's inevitable that estimates of cost and schedule will change as more is learned.

In some cases, risk can be transferred to another party. For example, a fixed price contract can in principle transfer the risk of cost overruns from the customer to the supplier. But this can be effective only up to the point where the supplier runs out of capacity to bear the financial consequences. In defence acquisitions, the sheer size of projects means that the risk associated with cost increases will often revert to the government because of the inability of a supplier to carry the burden. What's more, suppliers invariably charge higher prices if they're asked to carry risks associated with a project. And their cost of bearing risk may be higher than the Commonwealth's, because the government is running a large book of programs over which it can spread its risks. To the extent that risk can be transferred, the transfer comes at a cost. On the positive side, assigning risk to a supplier increases their incentive to actively manage any potential problems. As a general rule, risk (and the reward for managing it or the penalty for failing to do so) should be carried by the party best able to manage it—although achieving that contractually without introducing perverse incentives can be difficult.

For any individual project, the range of available risk mitigation measures will vary, but in general a mixture of risk retirement, risk transfer and contingency will be employed. In making decisions about defence projects, it's important to keep in mind that the goal is to *maximise the delivery of benefits over costs*, rather than to minimise the risks. Indeed, it makes no sense whatsoever to incur greater costs (or forgone benefits) to reduce the risks in a project beyond the costs (or forgone benefits) that the risks might give rise to. More simply, you wouldn't pay an insurance premium that cost more than your car was worth.

There exists a point in the development of defence projects where attempts to retire risk will begin to impose higher costs (or greater forgone benefits) than can potentially be avoided. In practice, the danger's that an unreasonably low risk appetite will lead to mounting delays as additional time is taken to try to minimise risk. Such delays can be highly costly and disruptive if they reach a point where it isn't possible to meet the planned withdrawal date of an existing fleet of assets or, worse still, if the fighting capability of the Australian Defence Force slips below that of prospective adversaries.

Making decisions

Many people believe that the men and women of the defence force deserve the best equipment that money can buy; some politicians might even try to claim that this is what they get. The facts speak otherwise. Almost every acquisition undertaken by Australia since at least World War II has involved a compromise between performance, platform numbers and cost.

So by what criteria should the inevitable compromises be made?

Under the *Commonwealth procurement guidelines* issued in 2008, the core principle upon which procurement decisions are based is 'value for money'. Fortunately, the definition of value for money offered in the guidelines is sensibly vague. In reality, decisions need to be tailored to the outcome sought and the circumstances in which the decision is taken.

A recent example from the United States is illustrative. In 2007, the US Air Force invited tenders for a fleet of air-to-air refuelling aircraft. In 2008, the European Airbus A330 MRTT (multi-role tanker transport) was announced as the winner over the competing Boeing KC-767. Following a protest by Boeing, the contract was reoffered, and in early 2011 the Boeing proposal was announced as the winner. Two contests, two aircraft, and two different outcomes.

While the machinations surrounding this $35 billion battle could fill a sizeable tome, the simplest interpretation (politics aside) is that the larger and more capable Airbus won the contest at a time when money was relatively easy to find, prior to the US recession, but that when the proposals were reconsidered in the tight post-financial crisis era, the less capable but more affordable Boeing option came out on top. In each case, it can be argued that best value for money was achieved in the circumstances that the decision makers found themselves in—circumstances in which the values of opportunity costs were weighed differently.

As this example shows, no simple mechanical rule exists to assure value for money in every situation. Sometimes it will be appropriate to maximise effectiveness, and at other times it will be best to pursue the least costly option that meets minimum performance criteria. Whatever approach is used, the techniques of commercial investment analysis can assist decision makers to sort through the options.

2 Walking the tightrope: understanding and managing risk in defence projects

Andrew Davies

Defence projects have a distressing tendency to hit the headlines for the wrong reasons. To some extent, this is unfair—many projects managed by the Defence Materiel Organisation (DMO) run pretty well. Unfortunately it's often the bigger, high-profile projects that stumble. That combines with the 'good news is no news' leanings of the press to ensure that there are more headlines along the lines of 'Billion dollar project goes off rails' than 'Ten million dollar project on time and within budget'.

Legacies, reforms and MOTS

When such headlines appear, there's an inevitable defensive response from Defence, invariably with the disclaimer that the project in question is a 'legacy' project—one that dates from before the latest set of acquisition reforms—and therefore can't be expected to perform as well as its recently approved counterparts. Meanwhile, the minister's office, keen to be seen to be tackling the problems it inherited from a previous government—after all, it's a legacy project—announces that it will be taking steps to review and overhaul the acquisition process to avoid these failings in the future. And so the cycle begins anew.

There are several competing timescales at work here. Australian governments come and go in at most three years, and recent defence ministers have turned over more frequently than in earlier times. Major reviews of defence acquisition have occurred at roughly five-year intervals—most recently Kinnaird in 2003 and Mortimer in 2008. But major defence projects operate on a longer timescale still. The *Defence Capability Development Handbook* notes that it's not uncommon for a project to take up to ten years to progress from government approval for inclusion in the Defence Capability Plan to final approval—after which acquisition and delivery can *begin*.[1]

So changes in the approach to project management and governance potentially prevent changes from being bedded in long enough for their end-to-end effect on projects to be made clear. Most major defence projects therefore harken back to a previous school of thought on acquisition policy. And very long projects (such as the 1983–2000 *Collins* class submarine project) might span several incarnations of governance models, with instability and a lack of clarity in governance as predictable outcomes. The net result is that today's reformation almost inevitably becomes tomorrow's legacy.

But it would be a mistake to focus exclusively on bureaucratic processes and governance arrangements. While they're important factors, they're only part of the story. At least as important over the past two decades has been the readiness of successive governments to approve *intrinsically* risky developmental projects. By definition, a significant share of risky projects will encounter problems. Indeed, there's only one way to avoid the cycle of poor project outcomes—by avoiding development projects altogether.

That's not always possible or desirable. Sometimes the capability return from a development project warrants the concomitant risk. Just as the commercial world accepts higher risk in the pursuit of higher commercial return, in the defence world the pursuit of a capability advantage sometimes necessitates the acceptance of higher technical risk in the procurement/development phase. However, it's a balancing act—sometimes the pursuit of higher performance results in a great outcome that saves lives or greatly increases combat capability. But sometimes it produces gains that are outweighed by disproportionately higher costs, with opportunity costs in either the numbers acquired or to the rest of the acquisition program, or it results in delays that mean that increasingly obsolescent platforms must stay in service longer.

ASPI has been an advocate of military off-the-shelf (MOTS) procurements, preferably from Australia's major ally, the US, whenever they provide acceptable levels of performance. Where possible, equipment should be acquired off established production lines with little or no modification. After delivery, its configuration should be managed so that upgrades performed for the parent services can be adopted by its Australian Defence Force users. Of course, there are pros and cons to such an approach, which were expanded upon in an earlier ASPI publication.[2]

Even conceptually straightforward modifications of equipment, or the integration of two or more fully functional subsystems, can result in poor outcomes. The lengthy process of integration of the AGM-142 missile onto the Royal Australian Air Force's F-111 fleet and the much delayed FFG frigate upgrade program are but two examples. The *Super Seasprite* helicopter required the integration of flight and combat systems and weapons and sensors sourced from multiple suppliers—the net result of which was a $1.4 billion investment for no capability return.

The situation can be summed up in a couple of rules:

- Off-the-shelf products that are technically mature are the only reliable way of avoiding cost and schedule risks.
- Because of the systems integration that's required, combining two or more MOTS systems is not MOTS.

Historical data supports these assertions, and the appendix to this paper expands on them. The conclusion that MOTS should be the preferred solution wherever possible is sufficiently important to reprise it here.

That said, the main aim of this paper isn't to remake the case for MOTS. Rather, it's to make some suggestions for managing those projects for which development work is judged to be the best path to a capability outcome.

Systems engineering

If there's no viable MOTS option that meets an identified operational requirement, then either the requirement needs to be scaled back (that is, we accept that capability will be less than 'ideal') or a developmental project will be required. In the latter case, a structured approach that allows dispassionate assessments to be made at each critical step of the process is required. This is essentially the reasoning behind the discipline of *systems engineering*, an approach designed to manage projects where the end result isn't a simple 'sum of its parts'—which is almost invariably the case with major defence capabilities.

NASA's 1995 definition of systems engineering is useful:

> System[s] engineering is a robust approach to the design, creation, and operation of systems. In simple terms, the approach consists of identification and quantification of system goals, creation of alternative system design concepts, performance of design trades, selection and implementation of the best design, verification that the design is properly built and integrated, and post-implementation assessment of how well the system meets (or met) the goals.[3]

As well as being a fairly succinct description of a complex idea, this definition allows us to categorise the points at which defence projects frequently go awry.

System goals

Pursuing the very top end of performance has a predictable effect on costs and schedules. As Norman Augustine, former president of Lockheed Martin and US Under Secretary of the Army, put it; 'the last ten percent of the performance sought generates one third of the cost and two thirds of the problems' (Augustine's VIIth Law).[4] Nonetheless, it's not unusual to find that user requirements are overreaching in the first instance—they're almost always skewed towards the top end of what's achievable. The net effect is an ambitious set of specifications being locked in before the associated risks are properly identified and understood. As pointed out by ASPI recently, the description of the future submarine in the 2009 Defence White Paper is an exemplar of this phenomenon.[5]

System design concepts and design trades

Exacerbating the problems caused by overambitious goals is the tendency to adhere to them even after the resulting difficulties are appreciated, rather than allowing what NASA calls 'design trades' to balance performance against project risk. Ideally, only those performance specifications that are truly necessary should be non-negotiable 'must haves'. Everything else should be regarded as 'nice to have' and be pursued provided that it doesn't have too much impact on overall project performance.

While there's an argument for each individual capability having the highest possible performance, the resources consumed pursuing that outcome can have an opportunity cost across the rest of the procurement portfolio. In the case of the future submarine, must-haves include sea denial and (probably) intelligence, surveillance and reconnaissance capabilities, while land strike and the deployment of Special Forces should be in the nice-to-have category. The overspecification of 'essentials' in projects is something that the services have raised to an art form. This has the effect of reducing the potential for competition at all levels in the project—between technologies, between suppliers and, at the most fundamental level, between capability solutions.

Design trades are especially important if they allow mature technologies to be selected. Figure 1 shows the cost growth within Pentagon R&D programs based on mature and immature technologies. It's often the case that getting the 'last ten percent' of

Figure 1: The impact of immature technologies on project performance

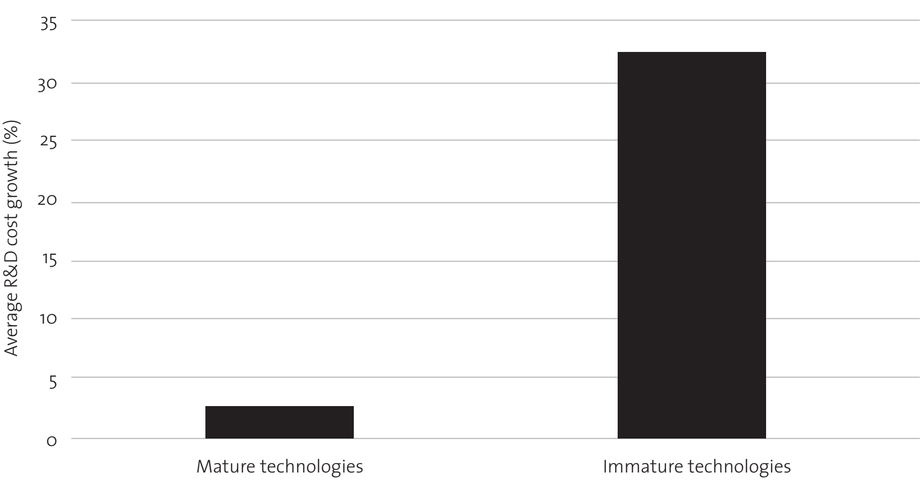

Source: US Government Accountability Office, Major Weapon Systems, *2007.*

performance requires new technologies, with predictable results for cost and schedule.

Selection and implementation of the best design

Before commencing production, it's important to bed down the design. The US Government Accountability Office has a dataset which shows that cost, schedule and performance estimates will be much more reliable *after* a structured sequence of design and requirement reviews (standard practice in systems engineering; see Figure 2). In the Australian Kinnaird process (see box), that should occur between first and second pass, but it's important to note that any estimates made before the review steps are finished are little better than educated guesses. That's how the price of Australia's three air warfare destroyers rose from a first-pass estimate of $4.5–6 billion in the 2004 Defence Capability Plan to the actual contracted figure of $8 billion at second pass.

The Kinnaird process

The Kinnaird process, introduced in 2003, is a two-stage decision-making process:

First-pass approval—at which the government considers alternatives and approves one or more capability development options to proceed to more detailed analysis and costing, with a view to subsequent approval of a specific capability.

Second-pass approval—at which the government agrees to fund the acquisition of a specific capability system with a well-defined budget and schedule, and to allocate future provision for through-life support costs.

Source: Defence capability development manual, 2006.

Figure 2: The value of performing systems engineering reviews before estimating costs

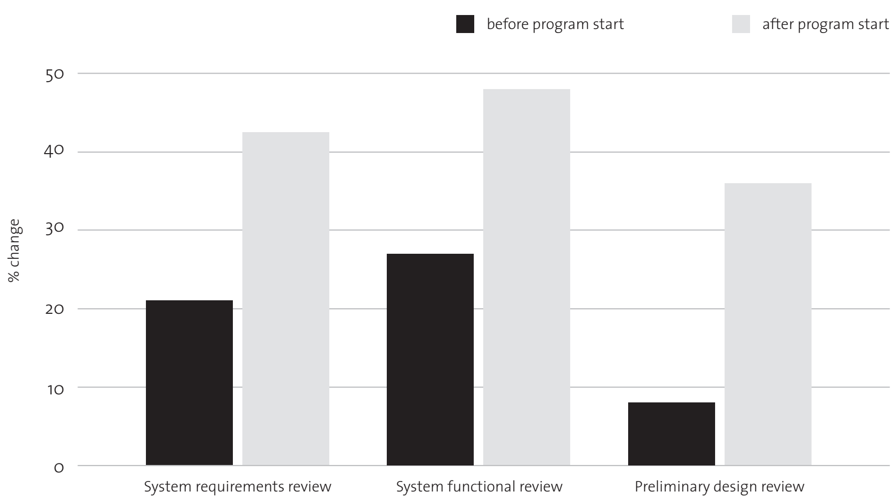

Source: US Government Accountability Office analysis of US Department of Defense program data.

Verification of design and build

Verifying the design of a major system necessarily involves verifying not only the design of subsystems but—perhaps most importantly—verifying that the subsystems can be made to work together. It's at this 'systems integration' step that many projects run into difficulties. Even if all of the subsystems are based on mature technologies—or are even off-the-shelf—it doesn't necessarily follow that they'll work smoothly together. It's not hard to find local examples where this has proven troublesome (the *Collins* submarine drive train; fitting the AGM-142 missile to the F-111) or has even been abandoned as impractical or prohibitively difficult (integrating the European MU-90 torpedo onto American aircraft).

During the build part of the project, unforeseen problems that arise require some agility in response. The trick is to respond in a way that doesn't derail the overall objectives or jeopardise the overall fidelity of the design. Any changes to the approved initial design need to be revalidated, including integration aspects.

Post-implementation assessment

Post-implementation assessment is the final step in the process for verifying and validating that equipment fully meets contractual requirements and is 'fit for service'—that is, it satisfies the government-approved requirements. While seemingly relatively straightforward, this part of the process requires as much rigour as any other step and has proven difficult in a raft of Australian programs (the *Tiger* armed reconnaissance helicopter and multi-role tanker transport aircraft being two recent examples).

That's the theory. How are we doing?

A recent Australian National Audit Office (ANAO) audit of twenty Navy acquisition projects provides a snapshot of the current state of health of systems engineering and the end-to-end management of major projects within the Department of Defence. At first blush, the result isn't pretty. While acknowledging that Defence has been making steps towards 'seamless and well-developed processes and systems', including the adoption of systems engineering practices, the audit found that the benefits expected from the discipline (as shown in the US Government Accountability Office's data) are not being realised:

> ... the overall picture is of a capability development system that has not consistently identified and responded, in a timely and comprehensive way, to conditions that adversely affected Navy capability acquisition and support. Opportunities to identify and mitigate cost, schedule and technical risks have been missed, resulting in chronic delays in Navy Mission Systems achieving Final Operational Capability. [6]

The audit identified shortcomings in each of the systems engineering steps described above. There were instances of requirements not being consolidated and agreed before project commencement (or even at the completion of the project for verification purposes); some of the projects examined lacked clear expositions of government-authorised scope, cost and schedule parameters; record keeping of testing and evaluation and configuration management—critical data for best-practice project management—was subpar.

As per tradition, the ANAO noted that fifteen of the projects were pre-2003, and were thus legacy projects. Happily, it seems that the 2003 Kinnaird reforms have had a positive impact; the five post-Kinnaird projects audited fared better than their predecessors, and improved practices, such as more tightly defining the capability to be delivered before contracts are signed, have produced better project outcomes.

Nonetheless, the ANAO and (to its credit) Defence both note that there are still improvements to be made. The trouble is that one of the root causes will prove difficult to eradicate. Successful implementation of systems engineering requires both technical and managerial expertise. The DMO has made a concerted effort to improve the latter in the past few years, and there's some evidence that good progress has been made—and has contributed to the improved outcomes of recent projects.

However, the other half of the equation is harder to improve. The in-house technical expertise of the services has been declining for some time, in part due to the dual trends of outsourcing all naval design, construction and deep-level maintenance to the private sector and of sourcing sophisticated equipment from overseas suppliers. But the Navy has also failed to sustain its residual capability. The ANAO observes that the Navy has filled only two-thirds of its own engineering positions, 72% of the Navy engineer positions in DMO, and only about one-third of the Navy engineer positions in the Capability Development Group. Hence ANAO's recommendation:

> In some essential systems engineering, technical regulatory elements and capability integration management areas, there are insufficient numbers of qualified staff, and this needs to be addressed as a priority.[7]

However, that's not necessarily easy. The external reasons for the downsizing of the Navy's engineering capability remain extant. And while the Navy has announced the appointment of a two-star officer with responsibility for reinvigorating the remaining in-house capability, that's likely to be an expensive and lengthy process because many of the skills are in high demand in other parts of the economy. While there are doubtless some things that can be done to improve the career paths of service engineers and thus the retention rates of skilled staff, market forces will continue to lure people away.

Finally, it's fair to observe that not all of the shortcomings that have been identified by the ANAO or by other studies can be laid at the feet of Defence. As ASPI showed earlier this year, underperformance by defence industry is responsible for many of the woes that have beset the portfolio of Defence projects.[8] The government has two plausible responses. The first is to tighten up its own contract management processes—making sure that incentives for good performance and disincentives for poor performance are in place and are enforced. The second is to give greater weight to past and potential performance shortfalls in industry when weighing up development versus MOTS options.

When things go wrong

No matter how good the processes used for managing projects are, whenever cutting-edge technologies are involved things will sometimes go wrong. The role of systems engineering is to help manage risk and to reduce the likelihood of poor outcomes. But even low-probability events happen sometimes—risk in projects is retired gradually as progress is made, but is never reduced to zero. (See Rule 3 in the appendix for a quantified explanation of this point.)

Even the best-managed projects will suffer setbacks or even failures from time to time.

When the inevitable review occurs, it's important to distinguish between the avoidable and the unfortunate. One of the most beguiling notions is that mistakes made in previous projects could have been avoided if only more, and more detailed, information had been available at key decision points. And to an extent that's true. As the data presented above shows, there are key points in a project where an expert look at the available data greatly improves the predictability of the outcome.

Indeed, the two-pass Kinnaird process was designed along similar lines, and it's important to provide decision makers with the appropriate level of detail at each step. But that's not synonymous with gathering ever more information. In fact, sometimes gathering more information leads to *worse* outcomes.

This somewhat counterintuitive phenomenon is explained in a classic study by Central Intelligence Agency psychologist Richard Heuer of the causes of errors in intelligence work. Heuer doesn't argue against the value of information. Indeed, he says that:

> [A]nalysis of information to gain a better understanding of current developments and to estimate future outcomes is an essential component of decision-making in any field. In fact, the psychological experiments that are most relevant have been conducted with experts in such diverse fields as medical and psychological diagnosis, stock market analysis, weather forecasting, and horserace handicapping.[9]

Note that the emphasis in that passage is on the analysis of information. Heuer's central thesis is that good analysis of the information at hand often trumps the gathering of further information.

So why do most reviews of project failures (and of the intelligence failures that are the focus of Heuer's work) focus on providing decision makers with more data? The answer lies in a subtlety of human psychology—we're inclined to be more confident about predictions based on large amounts of information. We weigh the quantity of data more highly than the quality of analysis.

An experimental example illustrates the point. Experienced horse handicappers were presented with datasets of variable sizes for a set of horse races (with the race and horse identities concealed). They were asked to predict the outcomes based on variables such as the weight to be carried, the percentage of races in which the horse finished first, second or third during the previous year, the jockey's record and the number of days since the horse's last race. Their performance at the task is shown in Figure 3. With a small number of variables, their predictions were accurate between 15% and 20% of the time. Importantly, their assessment of their accuracy almost matched their performance. However, as more data was provided, their confidence in their predictions grew steadily—in contrast to their performance, which was resolutely unaffected by the extra data.

The effect of this phenomenon on both decision-making during the approval process and subsequent project performance is bound to be negative. Once an artificial degree of confidence has been established, other aspects of human psychology do the rest. After a faulty assessment has been made, it tends to take on a life of its own. Established perceptions—in this case, 'this project is well understood and the predictions are reliable'—are hard to shake. Problems that

Figure 3: Gathering more information doesn't necessarily improve the accuracy of predictions—but it can work wonders for misplaced confidence

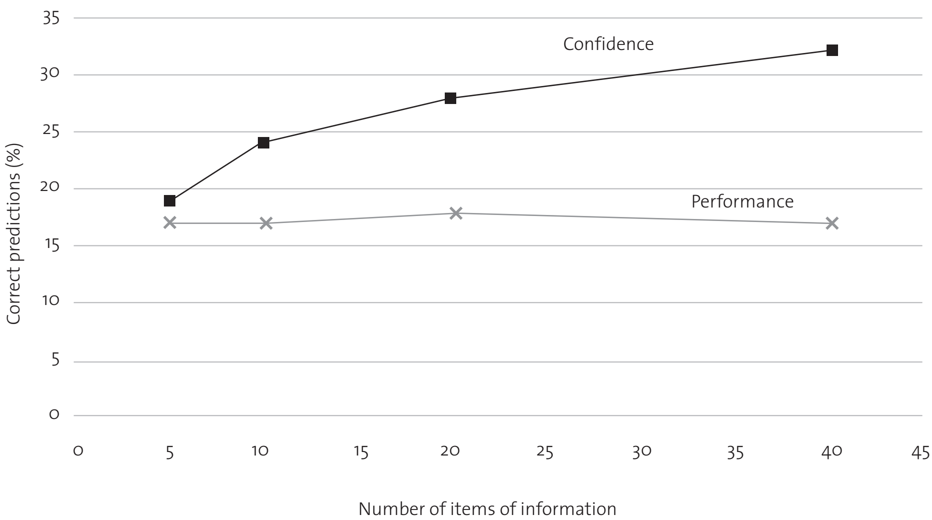

Source: RJ Heuer, Psychology of intelligence analysis*, Center for the Study of Intelligence, Central Intelligence Agency, Washington DC, 1999.*

arise in the early stages of the project tend to be viewed through a lens of overconfidence and their significance is downplayed. This effect is probably responsible for the degree of confidence expressed in the revised timelines and estimated costs—each subsequently proven to be overoptimistic—at each of the several resets in the F-35 Joint Strike Fighter program over the past ten years. It may also have reared its head recently in the almost perfunctory reset of the timeline for Australia's Air Warfare Destroyer Project when problems in the early stages manifested themselves—but time will tell on that one.

That's not to say that gathering information is pointless. Many risks can be mitigated through up-front analysis of *appropriate* information. The trick is to identify and gather the data that will help to identify risk and qualitatively improve the accuracy of predictions—and not overburden the capability development and project management process with reviews, meetings and documentary requirements that give an illusion of precision but little of the reality.

Conclusions

If reducing the level of risk in defence projects is assigned high priority, then the only sure-fire approach is to buy off the shelf, even if some capability compromises are required. If that's not possible, then the best chance of success will be through a rigorous approach to project management, implemented by people with the right skills and experience—preferably having successfully delivered big projects—and empowered to make the resource allocation decisions required. But all of the evidence suggests that such folk are in short supply in Defence, and the ones the organisation does have are spread thinly over a wide portfolio of projects and encumbered by a system that diffuses responsibility and manages by committee.

In the current economic circumstances, Defence will be competing with other employers seeking engineering and project management skills. Rebuilding the skill base in Defence will be, at best, a slow process, so the limited cadre of technically skilled personnel needs to be managed strategically. Their expertise will need to be reserved for those projects—which will necessarily have to be few in number—where systems integration and R&D are unavoidable.

It's always tempting to try to improve outcomes by applying more process, but mandating more review steps in the project process needs to be done judiciously. It may seem counterintuitive, but gathering more information can actually be counterproductive. Misplaced confidence is no substitute for a rigorous understanding of risk.

Appendix: Gumley's rules of defence acquisition

ASPI asked former DMO CEO Dr Stephen Gumley for his thoughts on this paper and his own 'lessons learned'. The following four rules are his response. Not surprisingly, the rules agree with much of the analysis of Australian and other project data that ASPI has done over the years.

Rule 1. Off-the-shelf acquisition of products that are technically mature is the only reliable way of avoiding cost and schedule risk.

Corollary: Deviation from MOTS increases risk, cost and schedule and should only be pursued when there are no MOTS systems with acceptable performance.

Rule 2. Combining two MOTS systems is not MOTS.

Corollary 1: Systems integration is an integral part of many defence projects.

Corollary 2: 'Australianised MOTS' is an oxymoron.

Rule 3. Risk is retired throughout a project but is never zero.

Corollary 1: It isn't possible—no matter how much information is gathered in advance—to eliminate all risk from a project.

Corollary 2: Cost and schedule estimates are refined progressively as R&D progresses. Signing a contract before R&D is completed will result in bad outcomes for at least one of the parties involved, and most often both.

Rule 4. Schedule delay is a non-symmetrical distribution (work can always be deferred to a later date but can never be advanced earlier than today).

Corollary: A portfolio of projects should always be overprogrammed (15% is a good rule of thumb).

Figure A1: The more a project deviates from MOTS, including the integration of other MOTS systems, the longer the delay in delivery, illustrating rules 1 and 2

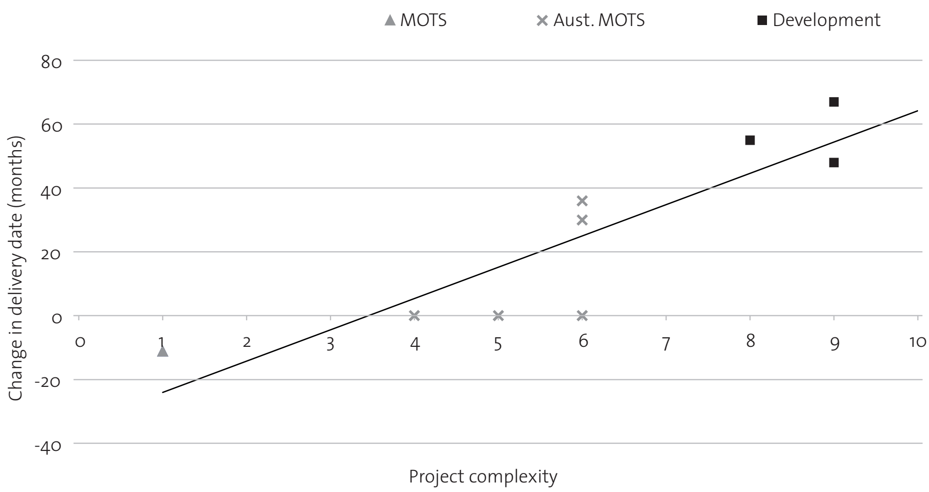

Source: ANAO data on major defence projects.

Figure A2: The accuracy of cost predictions for major US defence projects as a function of the fraction of R&D completed

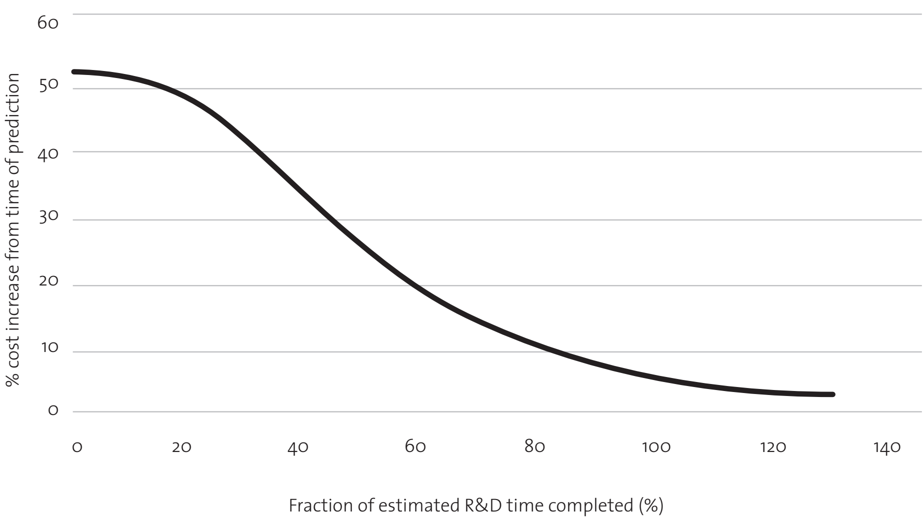

Source: NR Augustine, Augustine's laws *(revised and enlarged), American Institute of Aeronautics and Astronautics, New York, 1983*

Figure A3: Illustrating Rule 4—the distribution of annual project expenditure performance for 411 project years of defence projects from 1990 to 2010.

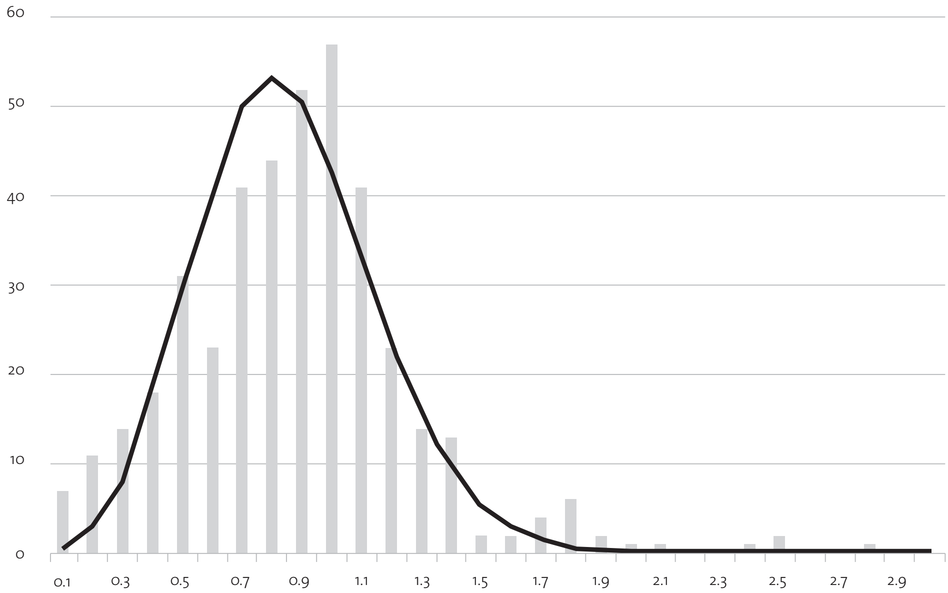

Note: The horizontal axis is the proportion of expected funds actually spent in-year. The average is 85%. The vertical axis is the number of project years that spent the corresponding fraction.

Source: Defence Budget papers and annual reports.

Notes

1 *Defence capability development manual 2006*, para 3.12, available from http://www.defence.gov.au/publications/dcdm.pdf.

2 Andrew Davies and Peter Layton, *We'll have six of them and four of those: off-the-shelf procurement and its strategic implications*, ASPI, November 2009, available from http://www.aspi.org.au/publications/publication_details.aspx?ContentID=231&pubtype=10.

3 *NASA systems engineering handbook SP-610S*, 1995, available from http://snebulos.mit.edu/projects/reference/NASA-Generic/NASA-STD-8739-8.pdf. Alas, the 2007 revision of this handbook is twice the length of its predecessor and is not as crisp in its definitions.

4 Norman R Augustine, *Augustine's laws (revised and enlarged)*, American Institute of Aeronautics and Astronautics, New York, 1983. This edition is the one to get—it was later rewritten as a business school book, which is much less fun.

5 Andrew Davies and Mark Thomson, *The once and future submarine—raising and sustaining Australia's underwater capability*, ASPI, April 2011, available from http://www.aspi.org.au/publications/publication_details.aspx?ContentID=291.

6 *Acceptance into service of Navy capability*, ANAO Report 57, 2010–11

7. ibid.

8 *The cost of Defence: ASPI Defence budget brief 2011–12*, ASPI, May 2011, available from http://www.aspi.org.au/publications/publication_details.aspx?ContentID=294&pubtype=3.

9. RJ Heuer, *Psychology of intelligence analysis*, Center for the Study of Intelligence, Central Intelligence Agency, Washington DC, 1999.

3 Risk and value: a CEO's perspective

Chris Jenkins

Australia is an advanced nation with deep science, engineering and manufacturing capabilities. It makes perfect sense that those deep technical capabilities should be harnessed for national security and defence.

Indeed, the changing nature of the global economy and our strategic position demand that we *must* harness our national research and industrial resources as a key element of national security policy. The unique nature of Australia's position, in terms of both our large continental scale and our geographic position, means that we'll always need to balance being self-reliant in key areas of defence capability against being able to rely on long-term support from our strategic partners.

Despite the depth of skilled capability in Australia, as a nation we have perennial bouts of self-doubt about our ability to do complex things in the defence arena. Governments begin to doubt whether Defence and defence industry can deliver the capabilities they aspire to; Defence begins to doubt whether local industry is up to the job or whether governments will fund it; and the companies investing real money in defence projects begin to doubt whether governments can drive the decisions or whether Defence can efficiently implement them.

The crisis of confidence amplifies risk dramatically. It slows decisions, impedes investment and discourages expansion. This is a negative spiral which, unless arrested, will have significant negative impacts on the national security of our country.

After nearly thirty years working in defence industry, I believe the catalyst for this spiral is the all too frequent demonstration of poorly managed risk in the delivery and support of defence capabilities. There are two key structural elements that together have the greatest impact on effectively managing this risk:

- the level of skill and, more importantly, experience we're able to deploy when projects are planned and launched
- the level of organisational maturity within Defence and defence industry to work collaboratively to identify risk, adapt, and respond quickly when project circumstances change.

The fundamental driver for both is closely linked to how effective we—Defence and defence industry—are at planning and implementing a continuous and achievable sequence of programs. Continuity is vital, and allows us to systematically build the skills, experience and organisational maturity needed to meet the capability requirements that will deliver future national security.

Risk management is integral to what we do

Working successfully in the business of supplying critical mission systems and equipment over an extended period necessarily requires being constantly attuned to risk. The systems and services we supply must work. If they don't, the consequences can be catastrophic because all programs have an effect on the capability available to our defence forces. Ultimately, increased capability delivers a greater probability of mission success and a higher level of safety in executing missions.

In defence industry, we're continually reminded of the human element to this risk, particularly at times of high operational tempo of our armed forces. This awareness runs deep in most defence companies. In my own company, we live and breathe some of that operational tempo and risk with the

Bushmaster vehicle, which in many ways provides a clear example of Australia's ability to manage risk and deliver exceptional, unique capability.

Workers on the *Bushmaster* production line in Bendigo take great pride in the lifesaving record of the vehicles they build. When a returned soldier comes to the factory and tells them the story of how the vehicle saved his life and the lives of his mates, it reinforces their commitment. Or, as happened recently, when the wife of a soldier approaches one of our *Bushmaster* support staff at a social function and tells him that thanks to the *Bushmaster* her children still have a father and she still has a husband, it's a powerful reminder of why we need to stay focused on delivering critical capability.

Ultimately our job, in concert with the various arms of capability procurement in Defence, is to deliver a capability advantage to our armed forces. Creating advantage involves risk in development and delivery. The risk shouldn't be hidden, but nor should we hide from it—it should be minimised when possible, but not avoided as a prerequisite to the procurement of capability. A refusal to take on risk in the delivery process is self-defeating and will ultimately reduce the capability advantage delivered to the Australian Defence Force.

Project continuity

Ten years ago, the *Bushmaster* vehicle program was very nearly cancelled. Had it been abandoned, we'd now be counting the cost not in dollars or schedule delays but in soldiers' lives. It's acknowledged at the highest levels in Defence that many Australian lives have been saved because of the protection provided by the *Bushmaster.* Two other forces operating in Afghanistan have purchased it, and others, notably the US, have adopted many of its lifesaving features into their vehicles after suffering terrible casualties in vehicles less suited to the Afghan war environment.

The *Bushmaster*'s success is due to a number of factors, including design, construction and off-road mobility. Less well known is the story behind the steel that the vehicle is constructed from, and its importance in the success of the vehicle. That steel is Bluescope steel. But it's not your average steel. It's made to a particular composition and then goes to a company called Bisalloy, also near Wollongong, where it undergoes heat treatment to produce the particular grade of high-hardness steel for the *Bushmaster.*

Bisalloy didn't suddenly come up with this steel overnight. In the 1980s, it began producing hull plates for the FFG frigates—producing steel that could deform without fracturing under blast testing. It further refined its hardened steel on the *Collins* submarine program, supplying about 8,000 tonnes of steel over fifteen years of the submarine build. As a result, when the *Bushmaster* program came along in the late 1990s Bisalloy already had a depth of knowledge on high-hardness steels. It worked closely with the Defence Science and Technology Organisation to further develop that steel, which now helps to make the *Bushmaster* such a lifesaving vehicle.

In short, decades of work by Australian companies under successive defence contracts created a continuous bridge of industrial skill and experience that have brought about a genuine capability advantage. On the flip side, there are other examples, some current, where a lack of continuity in defence projects has resulted in critical skills being lost from key industry sectors. In a tight labour market, those skills can't be turned off and on at will. The result is project delays and ultimately delays in capability being delivered.

Heightened risk and higher costs in these circumstances are predictable outcomes of a discontinuous process. In effect, risk is reintroduced into programs through the reduction in the base of skilled and, most importantly, experienced resources.

To minimise this risk, we need procurement planning that joins the Defence Capability Plan with, at the very least, the 'priority industry capabilities' in a coherent fashion. Gaining input from defence industry expertise in the pre-first-pass planning stage is essential. Combined with timely decision-making, the outcome should be greater continuity in projects within the plan and a climate more conducive to investment by defence companies. In turn, that will maintain and build the critical skills and capabilities needed for future programs.

Organisational maturity

Ending the stop–start cycle that kills experience and skills is an integral part of reducing risk. But equally important is building what I call 'organisational maturity'—which in my view is the key to effectively managing risk in defence projects. Identifying risk early is one thing, but having customer and supplier organisations that can collaborate to jointly and effectively respond to risk is another. By organisational maturity, I mean resilient organisations that have experience in dealing effectively with risk and have retained, in key roles, people who know what to do when the unexpected arises and who can interact constructively across organisations, customers and suppliers.

Defence procurement projects often take a long time, and their complexity means that almost by definition things won't always go according to plan. Retaining skilled people who have experienced project difficulties and solved them is critical. This is important throughout the chain of organisations involved in Defence projects—the Defence Materiel Organisation, the Capability Development Group and defence contractors all need to be building organisational maturity so that when risks are identified the organisations involved act quickly and jointly.

One thing you learn from experience is that while the problems on these projects may be very diverse, the solutions are often remarkably similar. They almost always involve closer collaboration between the Commonwealth, the end user and the contractors. The key benefit of organisational maturity is that it can bring the right resources to bear in the right way at the right time to resolve the inevitable unknowns when they arise.

For the most part, the 27,000 or so people who work in defence industry in Australia are proud of what they do and see it as a noble, nation-building vocation. We're here to create long-term value, and many of us have deep experience in managing risk. The aspirations of the latest Defence White Paper, with its future submarines, surface ships, vehicle fleets and other technically advanced equipment, mean that there's never been a more important time to build the expertise and organisational maturity that Australian taxpayers have the right to expect is being built to ensure future national security.

About the Authors

Dr Mark Thomson is the Director of the Budget and Management Program at ASPI. Prior to joining ASPI he held a number of positions in Defence working in the areas of capability development and resource management. He is the author of the annual *Cost of Defence* review of the Defence budget.

Dr Andrew Davies is the Director of the Operations and Capability Program at ASPI. He has written on the impact of Asian military modernisation programs, nuclear proliferation, defence acquisition projects and major Australian capability decisions.

Chris Jenkins is the CEO of Thales Australia, a leading company in aerospace, defence and security. Thales Australia is a corporate partner of ASPI.

Other Special Reports

Special Report Issue 41 – June 2011
Serving Australia: Control and administration of the Department of Defence
by Mark Thomson

Special Report Issue 40 – May 2011
Jihadists in jail: Radicalisation and the Indonesian prison experience
by Carl Ungerer

Special Report Issue 39 – April 2011
A reliable partner: Strengthening Australia – Timor-Leste relations
Various contributors

Special Report Issue 38 – March 2011
A better fit: National security and Australia's aid program
Report of an Independent Task Force

Special Report Issue 37 – February 2011
Sharing risk: Financing Australia's disaster resilience
by Edward Mortimer, Anthony Bergin and Rachel Carter

Important disclaimer

About Special Reports

Generally written by ASPI experts, Special Reports are intended to deepen understanding on critical questions facing key strategic decision-makers and, where appropriate, provide policy recommendations. In some instances, material of a more technical nature may appear in this series, where it adds to the understanding of the issue at hand. Special Reports reflect the personal views of the author(s), and do not in any way express or reflect the views of the Australian Government or represent the formal position of ASPI on any particular issue.

ASPI
Tel +61 2 6270 5100
Fax + 61 2 6273 9566
Email enquiries@aspi.org.au
Web www.aspi.org.au

TELL A FRIEND ABOUT ASPI

Join Australia's liveliest minds writing today on defence and strategic issues. Each year the Australian Strategic Policy Institute (ASPI) will produce up to six issues of ***Strategy***, six shorter ***Strategic Insights*** and a number of ***Special Reports*** on issues of critical importance to Australia and the Asia–Pacific.

Thoughtful, ground-breaking and often controversial, ASPI leads the public debate on defence and security issues.

JOIN ASPI

Name

Position

Company/Organisation

❏ **Government** ❏ **Non-Government**

Address

City **State** **Postcode**

Country

Telephone

Email

SELECT 3 FREE PUBLICATIONS

❏ Forks in the river: Australia's strategic options in a transformational Asia
❏ Changing pace: ASPI's strategic assessment 2011
❏ Regionalism and community: Australia's options in the Asia–Pacific
❏ Southeast Asia: Patterns of security cooperation
❏ A natural power: Challenges for Australia's resources diplomacy in Asia
❏ Our western front: Australia and the Indian Ocean

INDIVIDUAL	❏ 1 year *$199*	❏ 2 years *$378*	❏ 3 years *$537*
STUDENT*	❏ 1 year *$99*	❏ 2 years *$188*	❏ 3 years *$263*
CORPORATE (Oct 06+)	❏ 1 year *$649*	❏ 2 years *$1233*	❏ 3 years *$1752*

* (STUDENT ID ________________________________)

To join
1) Subscribe online www.aspi.org.au
2) Mail to Level 2, Arts House, 40 Macquarie St, Barton ACT 2600, or
3) Phone (02) 6270 5100 or fax (02) 6273 9566

❏ **Cheque** ❏ **Money Order** ❏ **Visa** ❏ **MasterCard** ❏ **AMEX** ❏ **Diners**

Payable to Australian Strategic Policy Institute ABN 77 097 369 045

Name on card

Card no.

Expiry Date / **Total Amount $**

Signature

This will be a **TAX INVOICE** for GST purposes when fully completed and payment is made.
Please note specialist publications are not included.