ASPI

SPECIAL REPORT

Accelerating autonomy

Autonomous systems and the Tiger helicopter replacement

Marcus Hellyer

December 2019

ASPI
AUSTRALIAN STRATEGIC POLICY INSTITUTE

About the author

Dr Marcus Hellyer is a senior analyst focusing on defence economics and military capability. Previously he was a senior public servant in the Department of Defence, responsible for ensuring that the government was provided with the best possible advice and recommendations on major capital investments such as the Joint Strike Fighter, Future Frigate and Future Submarine. He also developed and administered Defence's capital investment program. Marcus has also worked in Australia's intelligence community as a terrorism analyst. Before joining the public service, Marcus had a career as an academic historian in the United States.

Acknowledgements

The author would like to acknowledge the following colleagues who reviewed drafts of this report and offered valuable comments and suggestions: Mr Michael Shoebridge (ASPI), Dr Malcolm Davis (ASPI), Lieutenant Colonel Ned Holt (US Army/ASPI), Dr Natalia Jevglevskaja (Values in Defence and Security Technology, UNSW Canberra), Professor Jason Scholz (Trusted Autonomous Systems Defence Cooperative Research Centre).

Cover image: An Australian Army Tiger armed reconaissance helicopter on its inaugural flight with the Australian Defence Force. Photo: Defence image library, online.

Accelerating autonomy

Autonomous systems and the Tiger helicopter replacement

Marcus Hellyer

December 2019

ASPI
AUSTRALIAN STRATEGIC POLICY INSTITUTE

First published December 2019

Published in Australia by the Australian Strategic Policy Institute

ASPI
Level 2
40 Macquarie Street
Barton ACT 2600
Australia

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

Facebook.com/ASPI.org

@ASPI_org

CONTENTS

We always overestimate the change that will occur in the next two years
and underestimate the change that will occur in the next ten.
Don’t let yourself be lulled into inaction.

—Bill Gates

EXECUTIVE SUMMARY

The many benefits that autonomous weapons systems can potentially provide to militaries have been widely recognised. They include:

- removing humans from high-threat environments
- breaking out of manned platforms' vicious cost cycle
- achieving greater mass on the battlefield
- exploiting asymmetric advantage
- leveraging the civil sector's massive R&D spend on autonomous systems
- accelerating capability development timelines.

Despite this potential, the ADF has largely adopted a gradual, incremental, almost 'wait and see', approach to autonomous weapons systems. Its autonomous systems 'literacy' is growing, and the ADF has started to bring some systems into service. Nevertheless, it has focused mainly on remotely operated systems, which tend to be at the lower end of the autonomy spectrum. And some of these systems simply replicate the exquisite capability/exquisite cost dilemma presented by manned systems and consequently can be acquired only in small numbers.

The key challenge facing the ADF is in generating trust in systems with higher levels of autonomy. Defence needs to do more to develop its people and its culture so that they're more disposed to trust and therefore adopt autonomous systems rather than defaulting to manned platforms. Despite our cognitive biases towards manned systems, they can present more risk than autonomous systems, such as in the form of capability risk when we're unwilling to deploy manned systems because of the threat to their crews. Defence also needs to invest substantially more in developing technologies that support autonomy—or adapting existing commercial technologies to the military environment—so its autonomous systems are more reliable and capable and therefore more trustworthy.

The recent decision to retire Defence's manned minehunting ships and replace them with autonomous systems to clear sea mines shows that compelling arguments for autonomous systems can be made inside the Australian Defence organisation—provided sufficient trust in autonomous systems has been generated and Defence is willing to adjust its approach to risk.

However, securing greater investment in autonomous systems will be difficult, considering Defence's continued heavy investment in traditional manned platforms, which is unlikely to be moderated in the near term. Yet, autonomous systems offer the potential for Defence to hedge its capability risk, particularly if they can come at reduced cost and can relieve pressure on Defence's investment program. To do that, Defence will need to jump-start its approach to autonomous systems and accelerate the approach it has taken to date.

One way to achieve this is to not replace manned platforms with other manned platforms where there's no compelling need to do so. This frees up funding not only for autonomous systems but for other emerging priorities. Another way is to not seek to replace manned platforms with an autonomous solution that essentially does the same job. Rather, Defence could think disruptively and explore new roles that autonomous systems can perform that are quite different from those of current manned platforms.

The Tiger armed reconnaissance helicopter (ARH) provides a clear example of a case in which it's possible for Defence to avoid an expensive 'like for like' replacement of a manned platform. While the Tiger has had a troubled history, the Army has publicly stated that it now provides a high level of capability, including operating from the Navy's landing helicopter docks in amphibious roles. Moreover, Defence's Integrated Investment Program is delivering other systems that provide many of the effects sought from an ARH.

This an area where Defence can experiment safely with the accelerated adoption of autonomous systems without extreme capability risk should that experiment not succeed. It's an ideal area to explore human–machine teaming. It's also an area where accelerated experimentation can produce positive lessons for Defence more broadly.

Therefore, as part of the strategic and capability review that Defence is currently conducting, it should avoid investment of the roughly $3 billion needed to acquire a new ARH. Rather, it should keep the Tiger in service while investing around $1 billion of the funds saved in the development and acquisition of autonomous systems.

While these systems could deliver some of the effects sought from an ARH, Defence shouldn't seek primarily to develop an unmanned version of an ARH. Rather, it should actively explore in an open-ended way the disruptive potential of armed unmanned and autonomous systems for battlefield aviation. This would initially complement the Tiger to create greater effects than the Tiger alone. Eventually, this pathway would allow Defence to remove the Tiger and its human crews from the battlespace.

To accelerate this development, Defence could establish an interdisciplinary team, including a broad range of the Army's trades as well as industry and academia, whose sole function is to identify and experiment with disruptive autonomous innovations in battlefield aviation.

RECOMMENDATIONS

Recommendation 1

Defence should double the budget of the Next Generation Technologies Fund and Defence Innovation Hub to boost investment in the development of emergent technologies—around $850 million over the next six years.

- Defence should continue to prioritise autonomous systems and their enablers such as artificial intelligence for innovation funding.

Recommendation 2

Defence should retain the Tiger in service beyond 2025 to at least 2030.

- Defence should invest a portion of the funding saved in obsolescence management and limited upgrade of the Tiger—indicatively around $750 million.

Recommendation 3

Defence should redefine the Tiger ARH replacement in its Integrated Investment Program to:

- clearly set out the effects sought rather than specify a particular platform solution, so that a full range of options to deliver those effects can be identified and explored
- avoid acquiring another manned ARH to replace the Tiger unless it can be clearly demonstrated after extended development, acquisition and evaluation that autonomous systems combined with other capabilities in or entering service won't be able to deliver the effects sought by 2030.

Recommendation 4

The ADF should acquire a broader range of small, precision-guided munitions for the Reaper and other capabilities already in service or currently being acquired, at an indicative cost of around $300 million.

- Defence should explore whether those munitions can be domestically designed and manufactured.
- The munitions could also be integrated onto the autonomous battlefield aviation systems to be developed and acquired under recommendations 5 and 6.

Recommendation 5

Defence should establish a funding line of $1 billion over the next decade in its investment plan to support experimentation with and acquisition of unmanned autonomous systems to deliver disruptive battlefield aviation effects.

Recommendation 6

The Army should establish an interdisciplinary team that can draw on the funding line established under Recommendation 5 to experiment with and acquire unmanned autonomous technologies to deliver disruptive battlefield aviation effects, including:

- intelligence, surveillance and reconnaissance
- close air support
- electronic warfare
- cyberwarfare.

Recommendation 7

As part of its current review of the force structure plan, Defence should identify areas of capability across all six of its capability streams where it can make greater use of unmanned autonomous systems.

- Those systems would be developed to deliver disruptive effects, rather than simply replace manned platforms with similar but unmanned platforms.
- The systems should also reduce the scale of investment required for new manned platforms, or avoid the need to acquire them at all.

THE CASE FOR AUTONOMOUS SYSTEMS

What do we mean by 'autonomous systems'?

To avoid lengthy debates about autonomy, I'll try to keep this very simple. Robotic systems sit on a spectrum from lower to higher levels of autonomy. Automatic systems that provide the same response to the same kind of input are at the lowest level. Remotely operated systems generally have more autonomy, but of a limited or narrow nature, such as an automatic pilot functionality.

As systems gain the ability to learn and respond to new inputs in new ways, autonomy increases. Enabled by artificial intelligence (AI), systems with higher levels of autonomy can learn from their environment and make decisions in response to new inputs. They can essentially solve new problems in ways that are meaningful and therefore useful to humans.

Higher levels of autonomy allow humans to focus on command (that is, decision-making), rather than control (that is, directly operating the systems). This potentially allows organisations to move from having many humans remotely operating one vehicle to having one human commanding many vehicles. This study uses the term 'unmanned' to cover this spectrum.

I focus primarily on weapons systems; that is, things on the battlefield. They may carry advanced sensors or even lethal weapons. They may simply carry things useful to humans from one place to another. They may be part of a larger manned system, such as self-defence weapons enabled by AI. Primarily, this study is about autonomous weapons systems that allow Defence to remove at least some of its people from dangerous environments.

There are many other areas where Defence is exploring the use of AI for greater efficiency and performance, such as in logistics or in cybersecurity, but, ultimately, they're not what this study is looking at.[1]

The vicious cost cycle

There are good reasons for militaries to make greater use of unmanned and autonomous systems. The first is that the modern battlespace is becoming increasingly dangerous for humans. The Western states have now lost their monopoly on precision-guided weapons, which have proliferated widely. China has developed guided weapons that are extremely difficult to counter.

And it's not just great powers that have changed the operational landscape. Iran has just demonstrated the ability to deliver distant precision effects at a level akin to what the US demonstrated during Gulf War II, and Islamic State as a non-state actor proved adept at creating new problems for coalition forces through novel employments of commercial drones.

The threat spectrum also includes low-cost systems that can simply overwhelm platforms through numbers—swarming boats, smart sea mines, kamikaze drones. And even low-technology threats such as improvised explosive devices can cause high numbers of casualties and require major change to operational activities, as we've seen in operations in the Middle East and Afghanistan.

As threats have become more lethal and numerous, so has the complexity of manned platforms, because they need to be able to survive by defeating those threats in order to protect their human crew. Manned systems also need to be designed with a greater safety margin than unmanned platforms to protect their human cargo from environmental extremes, whether ocean depths or G-forces. But that complexity comes at great cost. Manned systems are often larger than unmanned systems, partly because of the need to provide life-support systems for the humans they carry, and this consumes capacity that could be used for other things and is another cost driver. And large, complex systems also take longer to design, build and test than simple ones.

The result is that the cost of conventional manned platforms is increasing exponentially.[2] Consequently, Western militaries can afford fewer and fewer of them, as can be seen in the trajectory of the US and Royal navies. Australia is bucking this trend only through an unprecedented scale of investment in naval shipbuilding. However, our future frigates and submarines are costing around $3 billion per vessel. Moreover, all going well, it will be around 20 years from the initiation of the Future Submarine Program until the first submarine is operational.[3] These time frames impose capability risk, as the threat environment will evolve even during the design and build of the first vessel, let alone over its service life.

Similarly, the Army's desire to protect its soldiers is driving it towards a $15–20 billion investment in a new generation of armoured vehicles that are several times heavier than the vehicles they're replacing, have many more subsystems such as sensors, and consequently will cost $25 million each when total project costs are taken into account.

Since our manned systems, such as ships, are increasingly optimised to defend themselves against threats and therefore use a considerable amount of their power and weight to provide defensive capabilities, they don't provide large amounts of offensive power. A future conflict could be a very one-sided affair if we face an adversary who has rigorously pursued the cost benefits provided by low-cost systems to deliver great lethal effect.

Breaking out of the cycle

Removing the human from the platform potentially allows us to break out of the capability–cost death spiral. First, it reduces the system's complexity, which reduces cost and the time needed to design and build the system. Since valuable space, weight and power aren't taken up keeping the crew alive, either the system can be smaller or the space, weight and power can be devoted to improving performance.

Simpler systems don't take as long to design, build and test, particularly if the lives of human operators aren't at stake. Unlike new designs for manned submarines, the large unmanned underwater vessels currently being developed by the US Navy and others won't take 20 years to enter service. And they'll be able to be further developed and enhanced more rapidly, as there won't be the same need to conduct the rigorous 'safety case' testing required for manned systems.

Because they're cheaper, they don't need to be designed for a 25- or 30-year service life, further reducing cost. In fact, some can be designed essentially to be disposable. This has the potential to increase mass: rather than small numbers of exquisitely capable systems, militaries will be able to field large numbers of systems potentially optimised for a small number of functions. This holds significant promise for the ADF, which has historically faced the challenge of covering huge operating areas with small numbers of people and systems.

The cost benefits and agility of the development cycle can be maximised if the development and production of unmanned systems exploits the components of the 'fourth industrial revolution', such as:

- AI enabling greater autonomy
- cheaper and more accessible space resources
- the use of digital twins in design, build and sustainment
- advanced manufacturing, including 3D printing.[4]

Another benefit provided by autonomous systems is that their development and sophistication are greatly facilitated by advances in the civilian sector, which can help to reduce the overall cost of military systems. The military has a monopoly on most weapons, but not on autonomous systems. In many ways, the civilian world has recognised the opportunities presented by autonomous systems and is ahead of the military. Civilian sector investment in R&D in AI and autonomous systems far outstrips the military's. Australia is a world leader in the use of autonomous systems in the mining industry. This potentially provides economies of scale, which, when coupled with the fourth industrial revolution, enable the creation of bespoke solutions on small budgets.

One illustrative example is the Queensland University of Technology's RangerBot, which is designed to kill crown-of-thorns starfish that destroy coral on the Great Barrier Reef. RangerBot employs the same technologies as autonomous military systems: it has a propulsion system, a navigation system, a sensor to detect potential targets, AI to confirm that the something detected is a legitimate target, and a lethal payload in the form of a toxin that's injected into the starfish.

Of course, the design task was somewhat simpler, since the starfish don't shoot back or seek to interfere in RangerBot's activities. But if a military drone is designed to be disposable, it doesn't necessarily need to incorporate a full suite of counter-counter-drone technologies.

Reasons to prefer unmanned systems to manned ones

It's often been said that autonomous systems are well suited to do the dull, dirty, dangerous and dear in place of manned systems. We can add the difficult to that list. Designing, building and operating manned combat helicopters involves all five of those challenges, and there's substantial scope for unmanned systems to address them.

Dull

Intelligence, surveillance and reconnaissance (ISR) often requires endurance—for example, in providing enough time on station to detect patterns of life or in escorting a convoy or covering a patrol. Much longer endurance is often needed than manned rotary-wing assets can provide without numerous sorties. Uninhabited systems can potentially provide longer endurance with better value for money. The fixed-wing unmanned aerial system (UAS) already being acquired by Defence—notably the Reaper and the successor to the Shadow 200 tactical UAS—have greater range and endurance than manned helicopters. It's likely that future rotary-wing UASs will also have greater endurance than manned helicopters, as they don't need to devote space, weight and services to support humans, allowing them to be optimised for other characteristics, such as endurance.

Dirty

Many environments are harsh for helicopters. Both the Tiger and the MRH-90 have encountered difficulties in dusty environments. The maritime environment is notoriously harsh on helicopters, causing accelerated corrosion. This can be addressed by the 'marinisation' of key components in the design and construction of the aircraft, as well as rigorous maintenance regimes in service. Both, however, involve additional cost and effort. UASs that have short service lives as part of the cycle of rapid development and replacement don't need to be designed to last for decades in such demanding environments.

Dangerous

Operating military helicopters is inherently dangerous. The two accidents with the largest number of fatalities in recent ADF history both involved helicopters: the SAS Black Hawk accident in Queensland in 1996 caused 18 deaths, and the Sea King crash on Nias Island in Indonesia in 2005 caused 11 deaths.[5]

Operating helicopters in combat is even more dangerous. According to one source, the US lost 5,607 helicopters in Vietnam—nearly half of those that served. Nearly 5,000 crew members were killed in action. Most of the losses were caused by anti-aircraft artillery and small arms, not sophisticated surface-to-air missiles.[6]

More recently, the coalition lost 129 helicopters in Iraq from 2003 to 2009, including 46 to hostile fire. Again, those losses were largely caused by unsophisticated threats. They included 29 AH-64 Apaches and 7 AH-1W Super Cobras (19 and 3, respectively, to hostile action), both of which are possible replacements for the Tiger.[7]

Australia was spared substantial helicopter losses in Iraq and Afghanistan because Defence judged, correctly, that its Black Hawk, MRH-90 and Tiger fleets weren't suitable to be deployed in the threat environment. Nevertheless, Australians were killed in allied helicopters. The ADF did rotate a flight of two CH-47D Chinooks through Afghanistan. One was lost through accident, resulting in the death of an Australian Army officer.

In Vietnam and the Middle East wars, the US enjoyed air dominance. In an environment of proliferating conventional capabilities, whether operated by or merely supplied by a technological near-peer adversary, losses are likely. Missiles aren't the only threat. Direct-fire weapons enabled by sensors and processing power are increasingly capable of accurately hitting fast-moving targets such as helicopters. The threat will only increase.

Dear

According to Defence's Integrated Investment Program, Defence has provisioned $5–6 billion for the ARH replacement.[8] Anecdotally, industry seems to think it can deliver a new fleet for significantly less than that, which is plausible, as my estimate is around $3 billion.[9] And even if the sustainment cost of a replacement is less than the Tiger's current $162 million annual cost, it will still be a lot. As a comparator, the Tiger's sustainment cost isn't an order of magnitude different from that of the Navy's combat helicopter, the MH-60R Romeo, which costs $120 million a year for a fleet of 24 aircraft.

If a new fleet of helicopters costs around $3 billion, that's effectively a capital cost of around $100 million per year. Combined with an annual operating cost likely to be well over $100 million, that's an opportunity cost of well over $200 million per year.

Moreover, investment in manned helicopters presents a huge opportunity cost. The 29 airframes specified in the Tiger replacement request for information will deliver an extremely small number of deployed aircraft. In a world where the numbers and capabilities of sensors and weapons are increasing rapidly, does it make sense to invest billions to get a small number of sensors and weapons in the air through a new ARH?

Difficult

Combat helicopters projects are notoriously difficult, as Defence's recent record and those of other militaries show. The nadir for Australia was the Super Seasprite project, which was cancelled after around $1 billion was spent due to intractable problems with the aircraft's automatic flight control system. Both the Tiger and the MRH-90 experienced significant delays and have underdelivered against capability requirements.

Upgrading helicopters throughout their service lives can also be difficult. Any modification needs to be well understood to avoid negative second-order effects. The long saga of putting floatation devices on the Tiger in case it crashes in water is a case in point.[10] With an unmanned aerial vehicle (UAV), you don't need to bother, because there's no crew whose safety must be your priority. The design flaws can be addressed in the next iterative design and production cycle, or you can live with the risk.

Disruptive versus sustaining innovation

Despite these advantages, the greatest benefit of autonomous systems lies not in their potential to perform the same tasks as manned platforms better or more cheaply. In his seminal study of innovation, Clayton Christensen distinguished between 'sustaining innovation' and 'disruptive innovation'.[11] Sustaining innovation seeks to conduct the current business model a little better. Disruptive innovation blows the old business model away—as digital cameras did to film-based cameras, or Netflix did to bricks-and-mortar video rental stores. The challenge, or dilemma, to use Christensen's term, for all organisations is recognising the disruptive potential of an innovation and not rejecting it outright because it doesn't support the current business model.

So, we shouldn't start by asking whether an autonomous system can do the same task as a manned platform better or cheaper. If the expectation is that unmanned platforms will provide better performance than manned platforms doing the same roles, we're likely to be disappointed. That measure is almost perfectly designed to perpetuate the age of the exquisitely expensive, exquisitely complex manned platform.

Rather, we could look at how autonomous systems can potentially make it impossible or at least very costly or difficult for potential adversaries to use their capabilities to achieve their desired ends. Even a highly autonomous underwater vessel is not going to perform all the missions of a manned submarine in the near term. However, saturating the maritime battlespace with cheap, disposable unmanned underwater vehicles (UUVs) and surface vessels with many sensors and even weapons could raise the risk for traditional manned submarines to unacceptable levels. And in the nearer term, operating large numbers of UUVs along with manned submarines is almost certain to provide a much more effective combat force than the most elegant manned submarine alone.

The potential has been recognised

Numerous militaries have recognised the potential benefits of autonomous systems and stated their intent to push for greater use of them. Autonomous systems are at the heart of the future operating concept of the US Marine Corps, which is generally regarded as one of the most innovative Western military formations:

> By exploiting the technical revolution in autonomy, advanced manufacturing, and artificial intelligence, the naval forces can create many new risk-worthy unmanned and minimally-manned platforms that can be employed in stand-in engagements to create tactical dilemmas that adversaries will confront when attacking our allies and forces forward …[12]

While the US Navy at one level seems to be aiming to return to a fleet of 355 manned warships, at another it appears to have accepted that the cost of doing that (around twice as much as it has spent on average on shipbuilding over the past 30 years) is unachievable.[13] Consequently, it's also investing in the development of a range of UUVs and unmanned surface vessels, some with significant levels of autonomy, in order to generate mass.[14] This includes platforms with high levels of capability. For example, Boeing's very long range Orca UUV has successfully navigated autonomously from San Diego to Hawaii. It also has a large payload bay for deploying sensors, sea mines and, potentially, other weapons.[15]

Since China has been actively seeking asymmetric advantage, it isn't surprising that it, too, is developing an impressive array of unmanned systems for the land, air, surface and subsurface domains. Many of them are optimised to defeat conventional manned platforms. Moreover, it's enthusiastically exporting them around the world.[16]

Even non-state actors have exploited remotely operated systems. The Islamic State of Iraq and Syria beat the ADF in getting armed drones into service. It integrated grenades with cheap commercial drones, which it deployed effectively against the Iraqi Army.

Progress in the Australian Defence Force

The ADF is also exploring autonomous systems. For example, the Australian Army has published its *Robotic and Autonomous Systems Strategy*, which recognises that robotic and autonomous systems (RASs) can generate mass as well as 'improve firepower, force protection and manoeuvre, enable sustained missions and identify threat and targets on the battlefield.' It flags 'moving away from the current human centric methods of operation'. The strategy states that 'we must start thinking about how Army can best use RAS capabilities, determine what human–machine teaming could look and operate like, and consider how we could operate with and alongside machines.'[17]

But Defence needs to do more than 'start thinking', as progress to date has been mixed. The Army has acquired micro military UAVs and introduced them into service throughout its units. It's also developing its personnel's UAV literacy by distributing cheap, off-the-shelf commercial drones to them. The Army has also experimented with a range of unmanned ground vehicles, but so far it hasn't retired any manned platforms and turned to unmanned systems to fulfil related missions.

The RAAF currently doesn't operate UASs. It did lease the Israeli Heron UAS for use in Afghanistan, but that capability has since been retired. It's acquiring the Triton high-altitude long-endurance UAS, which is developed from the Global Hawk but optimised for maritime surveillance. The acquisition of a variant of the Reaper armed medium-altitude long-endurance UAS is progressing through the capability development process. Yet neither has entered service with the ADF despite earlier versions of both systems having been operated for well over a decade by other countries.[18]

While both the Triton and the Reaper offer substantial improvements over traditional manned ISR platforms in terms of endurance, they're both analogous to traditional platforms in being extremely expensive. Defence's investment plan gives a cost range of $3–4 billion for a system of six or seven Tritons, so they're certainly not disposable. This exquisite cost didn't prevent Iran from shooting down a US Global Hawk earlier this year using much cheaper systems.[19]

The Navy has been following a crawl, walk, run strategy in its implementation of UASs. It has established 822X Squadron specifically to experiment with UASs for future use on its ships (Figure 1). It has trialled both fixed-wing and rotary-wing platforms at sea in preparation for the large-scale acquisition of aircraft for the Navy's major surface combatants and future Arafura-class offshore patrol vessels through a major project in Defence's Integrated Investment Program.

Figure 1: 822X Squadron launches a ScanEagle from HMAS *Leeuwin* at sea

Source: Defence image library, online.

As I discuss in the next chapter, the Navy has been experimenting with autonomous mine-clearance systems that, in contrast to traditional minehunting approaches, allow humans to remain outside the minefield.

The Prime Minister announced one encouraging development in this area during the 2019 federal election campaign. Although the details weren't explicit, it appears that, instead of conducting an expensive upgrade of its manned minehunting ships, Defence will move more quickly to replace them with autonomous systems.[20] Those systems will be operated from dedicated support vessels or 'vessels of opportunity' outside the minefield. No doubt a factor here is the difficulty and cost of upgrading the Huon-class minehunters, as each component on board needs to be able to operate with a very low acoustic and electronic signature, meaning that many items must be manufactured specifically for this purpose, which drives costs considerably. Removing the human from the minefield removes those additional costs. Despite the lack of fanfare, this is a significant development, as it's the first time that Defence will be entirely dependent on autonomous systems for a capability previously provided by manned platforms.

DEVELOPING TRUST IN AUTONOMOUS SYSTEMS

Trust

The ADF's mixed record on autonomous systems is largely typical of Western militaries. Despite the benefits that unmanned and autonomous systems seem to offer, their promise has often exceeded the reality.

There are many reasons for that, but a common element is the issue of trust. Virtually all studies indicate that trust is critical in the development and adoption of autonomous systems. Generally, analysts identify three groups of trust factors:

- the human related—the propensity of individuals to trust autonomous systems
- the technology related—the ability of the autonomous system to reliably perform the tasks expected of it
- the culture or environment related—whether a particular organisation or culture is open to the use of autonomous systems.[21]

None of those factors is fixed; that is, the level of trust that individuals or groups have in particular autonomous systems and AI is context specific and can increase or decrease.

All three trust factors are important; it isn't enough for the technology itself to be effective and reliable. For example, for individuals or groups who fear that machines may displace them from their jobs, the more effective an autonomous system is, the less incentive there is to adopt it. Militaries are no exception to this.

The following discussion shows that all three trust factors have played a role in inhibiting the take-up of autonomous systems in the ADF, and there are trust deficits in all three areas that need to be addressed.

Insufficient R&D investment in autonomous technologies

Defence isn't investing adequately in understanding and adopting the technologies that underpin autonomous systems. Without investment, those systems won't deliver a trustworthy level of performance. Lack of trust prevents investment, but investment is indispensable for the generation of trust.

Defence has two main innovation funds. The Next Generation Technologies Fund is funded at $730 million over 10 years (2016–17 – 2025–26) and is intended to promote research in future technologies that 'have the potential to deliver game-changing capabilities'. The other is the Defence Innovation Hub, which has a more immediate focus on turning innovation into capability. It's funded at $680 million over those 10 years.

Between them, the two funds have around $140 million per year. That's only slightly more than one-third of 1% of the Defence budget. This is spread over a wide range of areas; the Next Generation Technologies Fund invests in 13 priority areas; the Innovation Hub covers all six of Defence's capability streams, but the bulk of its funding goes into two ('ISREW, Space and Cyber' and 'Key Enablers'). Autonomous systems are just one part of their remit.[22]

One key measure that's been funded is the Trusted Autonomous Systems Defence Cooperative Research Centre.[23] However, its funding is $50 million over seven years.

The Next Generation Technologies Fund and Innovation Hub fund good work, but they just don't fund enough of it. Compared with high-tech companies that see R&D as essential to their competitive edge, Defence's level of investment in innovation is simply insufficient.[24] No company that saw itself at the leading edge of technology would stay in business if it invested only one-third of 1% of its budget in R&D. That's not enough to be a 'fast follower', let alone a leader.[25]

There are other areas of funding within Defence than can be described as R&D, but whether they're being used to develop autonomous systems is another matter. Under the Future Submarine Project (SEA 1000), Defence Science and Technology has received $397 million for studies. Yet the program has said that the first batch of submarines will use only existing technologies in order to reduce risk.[26] There seems to be little prospect, then, of that funding being used to develop disruptive technologies in undersea warfare.

This picture may be partially explained by the fact that, unlike the private sector, defence organisations are artificially protected monopolies during peacetime. They're exposed to their competitors only when conflict breaks out. It will be only at that point that they'll learn whether they have engaged sufficiently in innovative technologies and practices. But that doesn't justify inaction.

Moreover, the vast bulk of Defence's capital investment is locked up in the acquisition of traditional platforms. In 2019–20, there are 14 acquisition projects that each have a larger spend than the innovation funds combined.[27] The key challenge facing Australian R&D, whether in defence or in other sectors, has always been turning innovation into products. Defence has not yet found a way to embed that in its capability development processes. It must find a way to use the big dollars—the government's $200 billion investment commitment—to promote R&D that generates technological innovation that can be rapidly turned into military capability.

Like-for-like replacement rather than disruption

Many analysts have lamented Defence's propensity to replace like with like. For example, the next frigate needs to be bigger, faster, better protected and better equipped than the one we have now, as in the shift from the Anzac class to the future Hunter-class frigates. The result is that Defence's force structure now is very similar to what it was 30, 40 or even 50 years ago, with roughly the same number of roughly the same kinds of manned platforms.

This is a deeply rooted cultural factor inhibiting the development of trust in autonomous systems. Christensen has pointed out that the difficulty of identifying disruptive innovations arises precisely because they generally don't look like what the organisation is currently doing. Unfortunately, Defence collectively tends to see the right capability solution as something that looks a lot like the thing it's replacing. This inevitably limits Defence's openness to the disruptive potential of autonomous systems.

Trust in crewed platforms obscures risk

Another cultural factor that potentially inhibits the development of trust in autonomous systems is the value that the ADF (quite rightly) places on the ability of its people. According to this view, at least with a crewed platform in a worst-case scenario, the well-trained, imaginative humans on it can be trusted to display initiative and achieve something. Even if communication links are disrupted, we can rely on the initiative and training of the crew to carry out their mission anyway, and even adjust it as necessary. They can still process sensor data and launch weapons. In contrast, if communication with its human controllers is interrupted, an unmanned system can only do that if it's highly autonomous, and that level of autonomy may exceed our legal or ethical standards.

Crewed systems' subsystems are integrated to a degree by the crew. Or, if automated integration fails, the crew can do it manually (such as by entering coordinate data from a sensor into a weapon). Therefore, such systems can be multirole or even perform roles for which they weren't designed. In contrast, autonomous systems, at least for now, are each optimised for a single role (and since that's what provides their cost and schedule advantages, it's likely that they'll continue to be).

This reinforces humans' inherent bias against autonomous systems. Our tendency to judge machines more harshly has been repeatedly demonstrated. For example, one study showed that humans were much more forgiving of errors they themselves made than they were of the same errors made by machines in automatic mode.[28]

This bias obscures the fact that often manned platforms don't provide any useful capability at all. Because of their vulnerabilities, we're unwilling to deploy them into situations in which their crews could be at risk. This was the case with Australia's M113 armoured personnel carriers in operations in Iraq and Afghanistan (despite the M113 having completed a lengthy and expensive upgrade). The same was true of several classes of helicopters.

That may well be the case in future, even with capable platforms; because we have so few of them, the loss of even one or two could pose unacceptable risk in all but the most extreme contingencies. Think of any decision-maker's dilemma: should we deploy two of the Navy's nine frigates, or one of its three air warfare destroyers, each with 150 sailors on board, into an environment exposed to ballistic missile, cruise missile, hypersonic missile, rail gun, sea mine and torpedo threats?

Moreover, legacy manned platforms are extremely vulnerable to cyberattack, as a recent US Government Accountability Office report highlighted.[29] They may even be more vulnerable due to the sheer number of legacy-software-driven systems on them. They might not be able to provide any more capability than unmanned systems if key systems are shut down by cyberattack.

The right level of autonomy

RASs sit on a spectrum from lower to higher levels of autonomy. Automatic systems that provide the same response to an input are at the lowest level. Remotely operated systems have more autonomy, but generally of a limited nature, such as an automatic pilot. As systems gain the ability to learn and respond to new inputs in new ways, autonomy increases.

In order to develop greater trust, it's important to seek an appropriate level of autonomy. If users expect a level of autonomy that can't be achieved, they'll be disappointed and the trust deficit will be reinforced. In the near term, robotic systems will probably be best suited for tasks that are narrow but complex, such as landing an aircraft on an aircraft carrier.[30] The human tells the non-human system what to do and the system works out the best way to do it.

But too much autonomy can also undermine trust. As systems move along the spectrum of autonomy, from automatic to autonomous, or from narrow to more general autonomy, their behaviour becomes less predictable. This is challenging for military organisations that expect their test and evaluation processes to determine exactly how a weapon will perform—although this is probably resolvable by using testing and experimentation to define operational boundaries and characteristics, rather than seeking to rigorously predefine them in system requirements or design. But, currently, the more autonomous a system is, the harder it is for the military to trust it.[31]

Therefore, for the foreseeable future, the most likely and beneficial uses of autonomous systems will be to contribute to human–machine teaming or, in other words, complementary manned and unmanned systems operating together. Stephen E Davis writes:

> [A]utonomous systems are therefore better viewed as potential force multipliers that broaden human operated capabilities by allowing the role of personnel to involve more command and less control. For example, the introduction of autonomy in UAVs would allow human operators to manage multiple coordinated platforms.[32]

It's highly likely that human-crewed platforms will be with us for a long time to come, but they'll increasingly make use of human–machine teaming in which machines contribute the form of autonomy that they're best suited for. For example, ships may become minimally manned, with a small crew maintaining key propulsion systems, but with sensors and weapons functioning largely autonomously under a remote human in or on the loop. Those maintainers might not even need to be on the ship all the time. And that's consistent with the historical reduction in warships' crew sizes as machines perform roles previously done by humans.

In the area of ARH capability, we're already seeing militaries experimenting with human–machine teaming in which the human crew of a helicopter operates complementary unmanned systems.[33]

This may also mean that in some cases AI may 'run' the team and direct manned platforms. During Exercise Autonomous Warrior in Jervis Bay in 2018, Defence experimented with AI generating optimal search patterns for unmanned vessels and aircraft. But there's no reason why AI can't do the same in order to make the best use of scarce crewed ships and aircraft. Of course, the humans can still overrule the AI but, in the 'narrow but complex' task of optimising the use of resources, the AI can probably do a better job.[34]

Familiarity builds trust—but can we go faster?

When engagement with autonomous systems is conducted well, increased familiarity leads to increased trust and that results in what we're ultimately seeking, which is enhanced military capability. Defence's engagement to date has been measured, and there's evidence that, overall, it's slowly generating trust.

An example of this is the Navy's ongoing program of work with autonomous mine clearance systems. It's been working with such systems for some time, assessing their strengths and weaknesses (Figure 2). It's already acquiring deployable minehunting systems that can be operated from any vessel of opportunity. But its intent was not to move solely to autonomous systems yet; Defence had planned on conducting an expensive upgrade of its manned minehunting fleet. Yet, as noted in the previous chapter, in this year's election campaign the government announced that it would instead adopt autonomous systems operated from new support ships or vessels of opportunity from outside the minefield.

Figure 2: Navy personnel launch the REMUS, an autonomous underwater vehicle for locating ordnance on the seabed, off Talaghi Island as part of Operation Render Safe 2013

Source: Defence image library, online.

It's not clear what prompted this. Certainly, the long program of experimentation helped generate trust (it's hard to see Defence advising the government to commit to an autonomous mine clearance system if it truly thought the system was unviable), but it's also possible that other factors played a role. Committing a billion dollars to keep a small number of manned ships in service for only another decade probably represented limited value for money. It's also possible that the government saw this as an opportunity to announce an expansion to the scope of the minor war vessel stream of its continuous naval shipbuilding program. Building new ships in Western Australian shipyards certainly gets better election campaign headlines.

In short, the extended process of experimentation increased the Navy's trust in an autonomous future, although other factors probably helped propel it into making that leap of faith.

However, while Defence's incremental approach is slowly developing trust, others in the civilian world and other militaries are moving faster. Incremental approaches seldom bring disruptive change, so we need to be actively looking for areas of capability in which we can accelerate the process, even if it might be outside Defence's comfort zone. The key is to find areas where the confluence of opportunity substantially outweighs the risk. One such area is the replacement of the Army's Tiger ARH.

A CASE STUDY: THE TIGER ARMED RECONNAISSANCE HELICOPTER REPLACEMENT

The replacement of the Tiger ARH offers a safe space for the ADF to experiment with the introduction of armed autonomous systems. This is not because the Tiger itself operates in a safe space, but because any risks associated with experimental autonomous systems can be mitigated. Rather than replacing the Tiger with another manned helicopter, Defence should retain it in service while developing autonomous systems that will initially operate in conjunction with the Tiger. Ultimately, those systems may remove the need for a manned helicopter in this role, as well as potentially evolving entirely new disruptive capabilities.

Figure 3: Tiger operating from a landing helicopter dock during Indo-Pacific Endeavour 2019

Source: Defence image library, online.

The case against the Tiger

Certainly, the Tiger hasn't been an unqualified success story. In retrospect, it was probably not the best choice at the time. Australia is the only country that's acquired it, other than the three members of the Tiger consortium (France, Germany and Spain), and the aircraft is no longer in production. Every other country opted for something else. Sixteen countries have chosen the Apache, including Singapore and Indonesia in our near region and the US, South Korea, India, Japan and Taiwan in the wider Indo-Pacific.

There have been two Australian National Audit Office (ANAO) reports into the Tiger project, 10 years apart in 2006 and 2016.[35] Both make for depressing reading. The project is a sad tale of delays and missed performance targets. At the heart of it, the Tiger was evaluated by Defence—and sold to the government—as an off-the-shelf capability. In reality it was far from that. Instead of acquiring a mature, in-service capability, the Australian Defence Department was in many regards the lead customer and operator. Support and training systems were immature.

The result was that, even once aircraft were delivered, achieved flying hours were in some years only a third of what was planned, and Defence never came close to its original requirement of 7,147 hours per year. Due to its proprietary Eurogrid datalinks, the Tiger couldn't be fully integrated with other ADF platforms or the Army's battle management system.

Eventually, the Chief of Army declared final operational capability in 2016, but the ANAO noted that 'this was seven years later than planned, and was accompanied by nine operational caveats.' The Australian Tiger was never deployed to Iraq or Afghanistan (nor were our Black Hawks and MRH-90s). As I noted above, sometimes manned platforms provide no deployable operational capability due to the risks they embody.

Because the Tiger delivered few operating hours, its operating cost was extremely high. For many years it averaged over $40,000 per hour, or around three to four times more than fast jets such as the F/A-18 Hornet. Due to the nature of the sustainment contract, Defence was paying a lot for little capability. It's worth noting that the annual sustainment cost of the Tiger at $162 million this year is an amount comparable to Defence's two main innovation funds.

Without stating it explicitly, the ANAO recommended that Defence consider cutting its losses:

> The *2016 Defence White Paper* allocated $500–$750 million to address the current capability requirements of the Tiger platform with a view to replacing the platform mid next decade, at a cost of some $5–$6 billion. In effect, an upgrade is scheduled for consideration less than 12 months after the Tiger achieved Final Operational Capability. Defence should conduct a thorough analysis of the value-for-money of investing further in the Tiger, pending the introduction of a replacement capability.

Defence has issued a request for information to industry seeking information on potential replacements. Three responses were received, including one from Airbus that involves retaining and upgrading the Tiger and two proposing a new helicopter.[36] So, is replacement with a new manned ARH the best way to go?

The case for keeping the Tiger

Keeping the Tiger in service and investing in an upgrade would seem to be a classic case of the sunk cost fallacy. So why shouldn't Defence take the ANAO's advice and cut its losses?

One important reason to not follow that advice is that, since the most recent ANAO audit, the Army has developed the Tiger to the point that it provides a capability that the Army's happy with and is willing to deploy on operations. Lieutenant General Angus Campbell, the then Chief of Army, stated at Senate estimates hearings on 30 May 2017 that the Tiger is 'a very capable aircraft. In an operational environment I would welcome seeing a Tiger above me in the air.'[37]

Brigadier Stephen Jobson, the commander of the Army's 16th Aviation Brigade, has written that 'what you read and the way you think about Tiger is probably wrong now.' He has stated that the Tiger is now:

> a truly world-class platform … It is no longer the system we acquired. It is now modernized with next generation weapons, digitized connectivity, revolutionary tactics with Unmanned Aerial Vehicles, and interoperability with the Royal Australian Navy, Royal Australian Air Force and our allies.[38]

Lieutenant Colonel David Lynch, the commander of the School of Army Aviation, stated in May 2017:

> The Tiger aircraft is a potent weapon system that we have worked hard to mature to a point where it delivers the required effects on time and on target ... The aircraft itself has best-in-class agility and manoeuvrability which enhances our ability to employ the precise weapon effects to fight and win in either the land or amphibious environment. This agility also makes it an absolute dream to fly, as it has crisp control response and truly carefree handling characteristics, enabling greater focus on the mission.[39]

The Tiger has completed first-of-class flight trials with the Canberra-class landing helicopter docks (LHDs), so it can now be deployed on amphibious operations. Tiger helicopters have also been deployed on Indo-Pacific Endeavour,[40] during which they were operated for an extended period from an LHD, including for night flying.[41]

The Tiger's availability has progressively improved, except in 2017–18, when the international fleet was grounded for an extended period due to the crash of a German Tiger on operations in Mali (Figure 4). The original target of 7,147 annual flying hours was excessively ambitious for a fleet of 22 aircraft because the average hours per aircraft would have exceeded those of any of the ADF's other combat aircraft. The Army is now aiming for an enduring rate of effort of 5,300 hours a year. At 240 hours per year per aircraft, that would be consistent with other advanced combat aircraft.

Figure 4: Tiger ARH annual flying hours, planned (2008–09 to 2022–23) versus achieved (2008–09 to 2018–19)

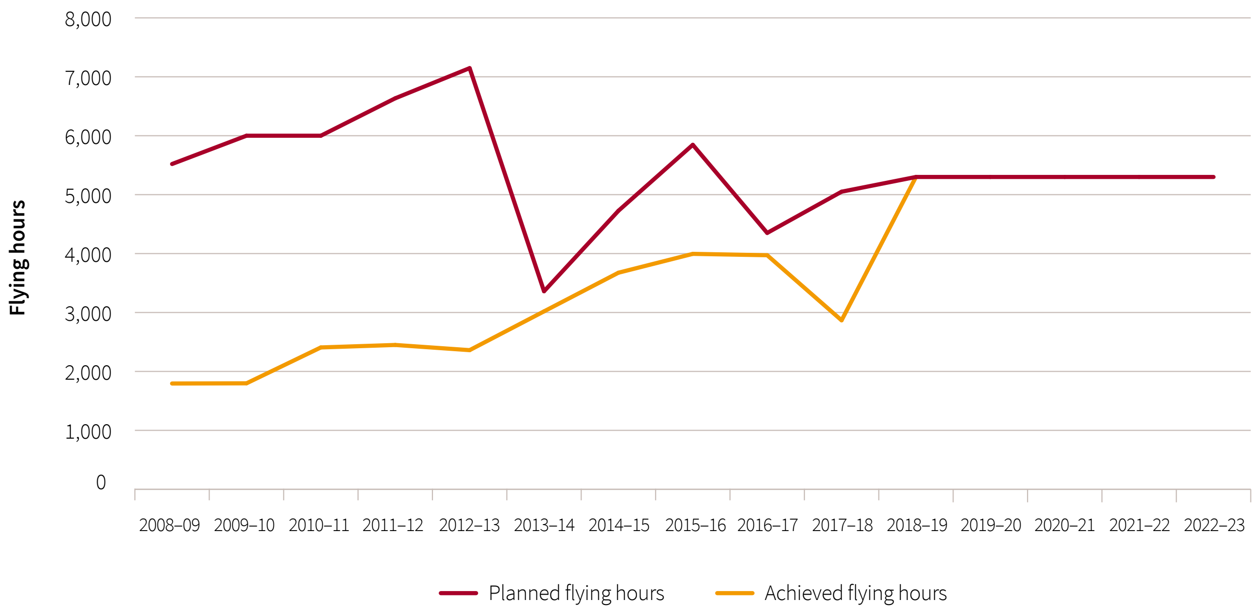

Source: Defence annual reports.

The Tiger's cost per flying hour has stabilised, although the total cost is rising due to more hours being flown. Defenders of the Tiger argue that its sustainment costs as presented in the Portfolio Budget Statements are misleadingly high, as many of the functions performed by contractors for the Tiger are performed by uniformed personnel on other aircraft, and so are an invisible cost that isn't attributed against those other aircraft in the budget statements (Figure 5).

Figure 5: Tiger ARH sustainment costs, 2009–10 to 2019–20

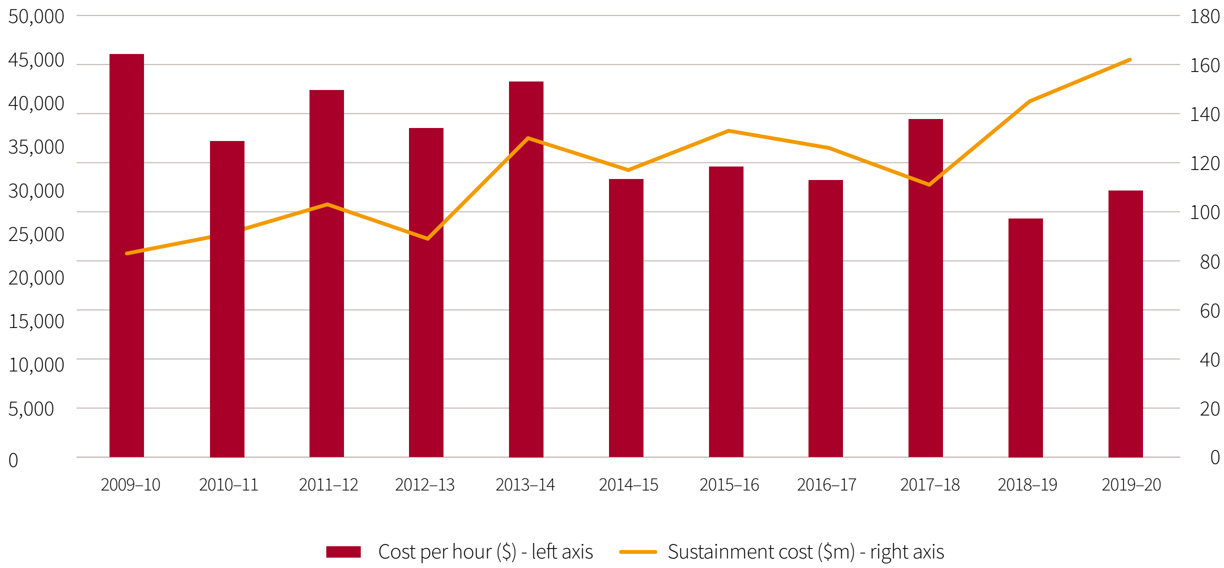

Source: Defence annual reports.

Due to the aircraft's prolonged journey into service, the airframes themselves are young. The first one reached 2,000 flying hours only in May 2017. It's reasonable to assume that they have at least half their life remaining. This, of course, is a separate issue from the obsolescence management of subsystems.[42]

There's certainly an element of the 'Collins effect' at work here. No matter how good the Tiger is, its reputation will find it hard to recover from the aircraft's history of underdelivery and underperformance. However, while it might never achieve the original requirement, in the judgement of the Army's leaders and aviation specialists it has reached a level of capability that's deployable and useful in the current threat environment. That might be surprising, but it's certainly welcome news that needs to be taken into account.

What does Defence want an ARH *capability* to do?

Before we can assess whether an option to replace the Tiger is optimal or even viable, we need to look at what Defence wants it to do. That question's difficult to answer. The request for information (RFI) that Defence issued to industry earlier this year didn't say. Rather than being a solution-agnostic RFI that listed the effects sought and asked industry how they could best be delivered, the RFI specified that the Army sought 'a proven and mature, off-the-shelf manned armed helicopter to deliver armed reconnaissance effects'.

That's probably the most effective and succinct statement of Defence's historical approach to force design: seeking to deal with the end of life of most capabilities by taking the 'like with like, but better' approach discussed above. The RFI didn't state what those effects were.

The RFI did at least describe three primary missions:

> The three primary missions for the LAND4503 platform are reconnaissance, attack and security operations. Reconnaissance operations are conducted to obtain information about an area, route or battlespace actors, and share the common tactical picture to enhance the awareness of battlefield commanders. Attack operations focus on the destruction and disruption of enemy forces. Security operations include the close support to air and ground movements and provision of protection to friendly elements.[43]

If Defence wants to give industry the latitude to develop innovative solutions, specifying that the answer must be a manned, off-the shelf helicopter isn't the way to do it. This suggests that Defence isn't interested in innovative solutions in this space—which is surprising, given the technological developments related to those effects, but also not surprising, given the slow speed of adoption of autonomous systems across the Australian Defence Department.

The effects sought seem to be ISR to feed into the common operating picture; a kinetic strike ability at some depth behind the front line; enough responsiveness to provide close air support; and enough persistence to be able to deliver those effects when needed.

The bigger picture

Beyond those mission requirements, consideration of whether and how the Tiger should be replaced also needs to take other, broader factors into account.

The funding context

Despite the incremental growth of the defence budget under the 2016 White Paper's funding model, it's clear that the Defence Department is squeezed for cash, particularly in its capital acquisition budget. That squeeze will be particularly acute in the delivery time frame for the Tiger replacement. Construction of the Attack-class submarines and Hunter-class frigates will have commenced. Domestic shipbuilding programs are likely to be consuming around $3.5–4 billion per year. Production of the LAND 400 infantry fighting vehicles will also be underway, requiring around another billion per year.

The acquisition of a replacement ARH will produce a requirement for several additional billions of dollars of investment funding in the middle of the 2020s. While a new helicopter may be cheaper to operate than the Tiger, those potential longer-term sustainment savings wouldn't ameliorate Defence's cash flow pressures in the mid-2020s (and the empirical evidence for new systems being cheaper than previous systems is very thin in the land of defence capabilities). Moreover, to deliver a return on investment on the acquisition funding, the new helicopter would need to operate well into the 2050s. It's difficult to predict the exact operating environment at that time, but there's no doubt it will be immensely more dangerous for manned helicopters.

In essence, replacing the Tiger with another manned ARH is a more egregious example of the sunk cost fallacy than retaining the Tiger. Investing in another manned ARH exacerbates rather than ameliorates a key challenge that Defence is facing: at a time when it needs to be investing in emergent technologies (and, to be frank, technologies that have already emerged and are in widespread use outside of Defence), much of its investment budget is locked up in traditional manned platforms.

The capability context

According to Defence's Integrated Investment Program, at the point when a traditional ARH is scheduled to be delivered, Defence will have already added to its increasingly long list of new capabilities several that can perform the missions outlined in the RFI. They mightn't be able to perform all of them individually, or in the same way as an ARH. Those capabilities include the following.

The Reaper armed medium-altitude long-endurance UAV

The Reaper is capable of much greater range and endurance than a manned helicopter. With a sophisticated sensor suite, it's designed to perform ISR. It's likely to be the first armed UAV in the ADF's inventory, with the ability to deliver precise, kinetic munitions including Hellfire missiles and both laser-guided and GPS-guided bombs. The small diameter bomb (already being acquired by Defence) is also being integrated onto the Reaper, as is the small glide munition, which weighs even less than the Hellfire, increasing the number of weapons it can carry.[44] The Reaper can't do everything that an ARH can (and vice versa) and it can't operate off an LHD, but it can deliver similar effects, although Defence will have Reapers in only very small numbers owing to their expense.

A new tactical UAS

Under project LAND 129 Phase 3, Defence will acquire a replacement for the Shadow 200 tactical UAS. By 2025–26, Defence aims to have six 'capability bricks', each of which can support 24/7 surveillance. The requirements for the future tactical UAS state that it won't be armed (why it can't have this option when small commercial drones operated by Islamic State have carried weapons is unclear), but it will have an extensive suite of sensors as well as the ability to laser-designate targets for other assets. The Shadow 200 already has a 125-kilometre range and nine hours' endurance, and it's reasonable to assume that its successor will be at least as capable. According to the project's tender documents, the capability should support amphibious operations.[45]

Figure 6: The Shadow 200 UAS

Source: Defence image library, online.

Longer range land-based fires

For kinetic fire options, the Army will have its current M777 towed artillery and, according to the government's election campaign promises, self-propelled howitzers, which have a greater range and rate of fire than the towed guns. Both can use extended-range munitions. While the range of the self-propelled howitzers would not be likely to exceed 40 kilometres even when using extended-range munitions, the Army is also acquiring a new system with substantially greater range. According to the 2016 Integrated Investment Program, 'a long-range rocket system will be acquired in the mid-2020s to complement the ADF's existing artillery capability, providing an option for long-range fire support (up to around 300 kilometres) to joint operations.'[46] That's further than a manned ARH can strike.

The F-35A Joint Strike Fighter

Also, by the mid-2020s the F-35A will have achieved final operational capability. With its extensive, integrated sensor suite and extensive complement of weapons, it can deliver many of the effects sought from an ARH capability—but from substantially greater range and altitude, making it immune to the small arms and anti-aircraft artillery that's responsible for most combat helicopter losses. Again, the F-35A can't operate from ships.

A key weakness in that mix of capabilities is the inability to deliver kinetic effects from ships. Naval gunfire support could provide some, but not to the same depth inland as the Tiger. Nevertheless, it's not clear which essential effects sought by Defence couldn't be delivered by a combination of those capabilities, particularly if the Tiger remains in service. And that's before we consider future autonomous capabilities.

A SAFE SPACE TO EXPERIMENT

Don’t replace the Tiger with a manned ARH

The Minister for Defence has announced that Defence is currently conducting a ‘re-assessment of the strategic underpinnings of the 2016 White Paper’, to be completed in early 2020. The review offers Defence the opportunity not only to reassess geopolitical factors such as the changing power balance between the US and China, but also to reassess its approach to autonomous systems. Just as the White Paper underestimated the speed of change in geopolitical factors, it’s clear that it also underestimated the potential of autonomous systems and the need to invest in them to achieve that potential.[47]

The Army’s *Robotic and Autonomous Systems Strategy* states that ‘force design implications of the rollout of RAS technology must be explored in detail, without being inhibited by perceptions, bias and human limitations—moving away from the current human centric methods of operation.’[48] As we’ve seen, however, a number of factors have impaired the development of the trust necessary for the ADF to embrace autonomous systems more fully. Defence’s approach to the adoption of autonomous systems has been incremental and measured.

This could be accelerated by placing bets on areas of capability where the odds of success are high and, equally importantly, the consequences of failure are low. The ‘replacement’ of the Tiger ARH appears to be an area where Defence can move away from its ‘human centric’ methods of operation with minimal risk.

Why bet on something when success is not assured? But is that any different from what Defence is doing with its heavy investment in manned platforms? In fact, by moving from six to 12 submarines Defence is doubling down on a strategy fraught with risk. Our inherent bias towards manned platforms can blind us to their risks—and substantial capability risk has already been realised because we haven’t been willing to deploy vulnerable manned platforms on operations. Investing in autonomous systems is a strategy of hedging risk by broadening Defence’s investment portfolio.

There’s no pressing need to replace the Tiger with another manned ARH. The Tiger itself provides a reasonable capability, and many of the effects sought from an ARH can be delivered by other systems already in or planned to enter service. It’s even more apparent that retaining the Tiger and complementing its combat power with a variety of autonomous systems would probably be substantially cheaper and provide a more effective overall combat capability than replacing the Tiger with the most reliable and cheapest new ARH.

There’s time to use the ARH replacement funding to experiment with and acquire real autonomous capabilities. This is an area where Defence can experiment safely with unmanned and autonomous systems. Initially, they may work in conjunction with the Tiger, but, ultimately, they’re likely to remove the need for a manned combat helicopter.

Industry aspects

In many ways, Australian industry and academia are better suited to contribute to the development and production of autonomous systems than to traditional manned platforms. There's significantly greater potential for Australia to develop sovereign capability in armed UAVs, for example, than in traditional manned helicopters. If Queensland University of Technology can develop the RangerBot, Australian industry and academia between them can also develop autonomous systems that can perform not just tactical ISR functions but also strike, close air support and electronic warfare.

The Defence Innovation Hub has already funded several UAV projects, and the Trusted Autonomous Systems Defence Cooperative Research Centre already provides a local centre of excellence. Australian industry has already begun to explore the development of autonomous weapons systems, including some with kinetic capabilities. For example, the Australian company DefendTex has a range of systems, such as the Drone-40, which can be launched by hand or from grenade launchers. The Drone-40 doesn't have the range of an ARH, but if it can be combined with a drone 'mothership' such as DefendTex's Tempest fire-support drone (Figure 7), the resulting combination could have comparable or even greater range and endurance than a manned ARH. Defence has yet to field systems of this kind.

Figure 7: DefendTex displays the Tempest fire-support drone at the Army Innovation Day 2017 at the Australian Defence Force Academy in Canberra on 6 November 2017

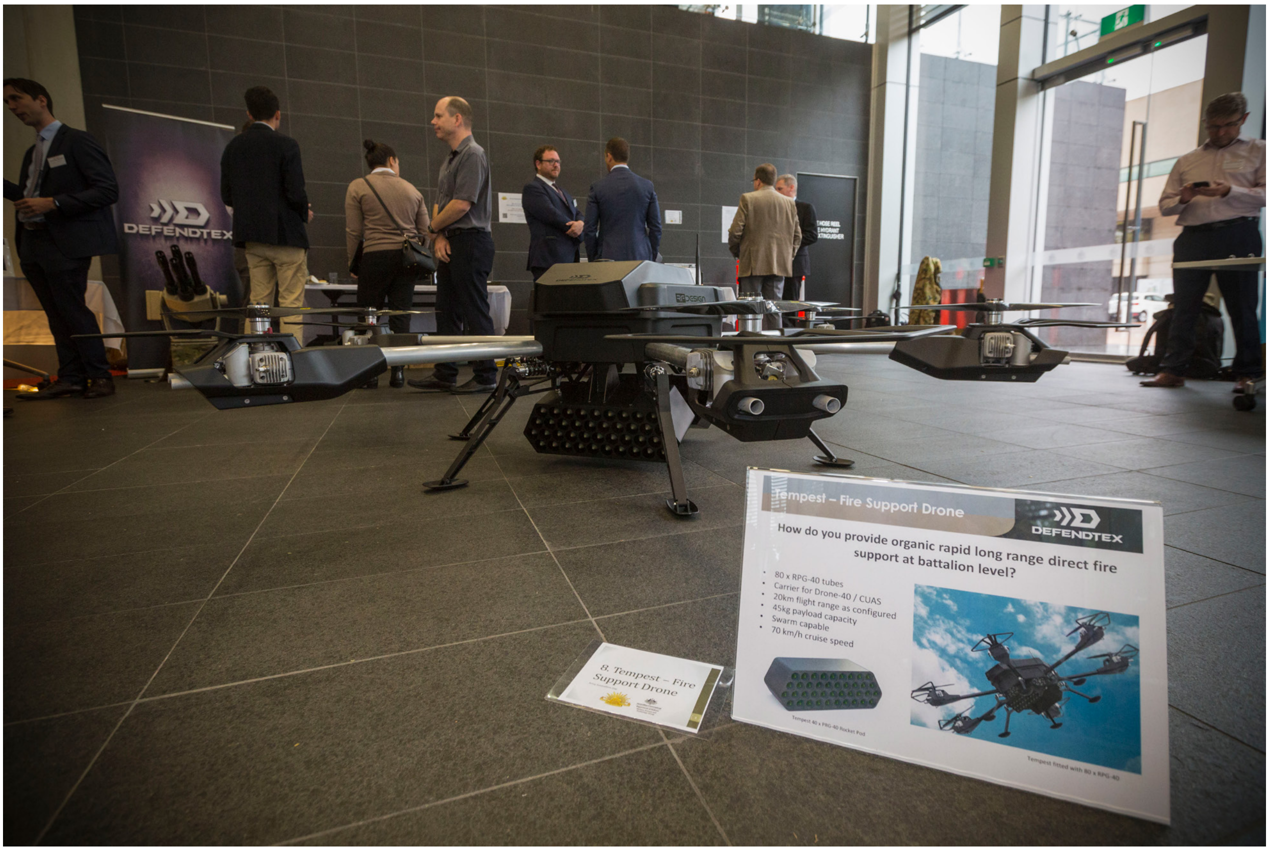

Source: Defence image library, online.

Australian industry also has the capability to develop a range of modular payloads that could be integrated into a standard drone mothership. In addition to kinetic weapons, the payloads could include ISR packages such as synthetic aperture radars and passive radars, laser target designators, communications relays, electronic warfare equipment such as sensors and jammers, and so on. By drawing on the attributes of the fourth industrial revolution, such drones could be repaired and modified in the field through 3D printing. Because the drones wouldn't be subject to the same airworthiness standards as manned platforms, and could essentially be disposable, modifications could be quickly made, tested and implemented.

Ultimately, pursuing innovation in autonomous systems would be likely to result in better outcomes for Australia's science and technology base than would the acquisition of an off-the-shelf foreign ARH.

Focus on effects

Of course, if the goal is simply to deliver an unmanned version of a manned ARH, any autonomous system will come up short. It's important not to see the UAS as a poor man's ARH. It's crucial to start with the effects sought, particularly those that the constellation of capabilities outlined in the previous chapter can't deliver.

One gap appears to be the ability to deliver lethal effects in support of amphibious operations, since neither the F-35A nor the Reaper can operate from the LHDs, and the Shadow 200 replacement will be unarmed. But it's by no means certain that a manned ARH is the only platform that can deliver the necessary effects in an amphibious context, so a particular focus of innovation, experimentation and acquisition could be on amphibious operations. There's no reason why an armed drone mothership couldn't operate from an LHD. In fact, it offers greater operational flexibility than an ARH, as it could operate from a range of naval platforms.

Use interdisciplinary teams to seek disruptive innovation

Viewing autonomous systems merely as direct replacements for manned systems is an exercise in what Christensen termed 'sustaining' innovation; that is, limited, incremental improvement to the current business model. That will most likely miss the disruptive potential of autonomous systems, and Defence will miss the real opportunities they present.

Studies of successful innovation have identified the need for a critical mass of supporters to experiment and advocate for innovation. Communities, constituencies and organisations are just as important as the technologies. In the case of autonomous systems, those communities need to be ones that are well disposed to trust their potential. Fortunately, the Army is on the path to developing a critical mass of RAS operators and general RAS literacy.

But most innovation is generated from interdisciplinary teams. They bring different perspectives to problems and can generate solutions that are not 'known' in advance by narrow specialists.[49] Furthermore, teams exploring potentially disruptive technologies must usually be situated outside of and quarantined from the larger organisation's traditional core business. This protects their resources, which would otherwise be sucked into addressing the larger organisation's problem of the hour.

More importantly, sitting outside the business-as-usual structure also allows them to pursue lines of experiment that potentially undercut the broader organisation's core business. This allows the innovation to head in directions that weren't (and couldn't be) predicted, which is often where the value of new approaches is found. That's why the US Department of Defense behemoth has the stand-alone experimentation agency DARPA.

In the Army context, this means that experimentation with armed UAVs shouldn't be conducted solely or even predominantly by the 1st Aviation Regiment or by the 20th Surveillance and Target Acquisition Regiment, which currently operates the Shadow UAS. Rather, it needs to draw on a wide range of expertise from across Defence—ISR, surveillance and target acquisition, fire support, combat arms such as infantry and cavalry, targeting analysts, and of course UAV operators. Why not include cyber experts who can consider how drones can deliver tactical cyber effects on the battlefield?

Disruptive innovation can be as much about how technology is used and who uses it as about the technology. History is replete with examples of how the true disruptive potential of technological innovation wasn't realised because the technology was shoe-horned into existing operating models or organisations.[50] The use of aircraft at the start of World War I is an example—leadership determined that aircraft should be limited to reconnaissance roles. It was operational experimentation by pilots that resulted in armed fighter aircraft.

Experimentation shouldn't start with preconceptions about where the capability should sit, for example. It may be that every infantry company could have its own armed UAV capability, distributing lethal effects widely over the battlespace, contributing to the disruption of the adversary as well as dispersing the high-value targets that the adversary is trying to hit.

What should Defence do?

The $5–6 billion currently programmed for the ARH replacement provides a substantial fund to enable Defence to pursue autonomous systems far more ambitiously. Even if a manned ARH replacement is 'only' $3 billion, avoiding that huge opportunity cost by retaining the Tiger will provide funds for much more imaginative investment.

Granted, some funding will need to be reserved for obsolescence management and limited upgrades to the Tiger. The current investment plan budgets $500–750 million for the Tiger assurance program. Since we're recommending that the Tiger stay in service longer, the higher end of the sum should be preserved in the investment plan.

A key focus of the funding freed up by retaining the Tiger should be to increase Defence's investment in R&D and innovation. Defence should double the budget of the Next Generation Technologies Fund and the Defence Innovation Hub. Around $850 million would cover the remaining six years of the White Paper decade up to 2025–26. While not all of that funding would be devoted to autonomous systems, Defence should ensure that such systems remain high priorities for both innovation funds.

Beyond simply funding R&D activities, Defence should also use this opportunity to acquire autonomous capabilities complementary to the ARH and put them in the hands of Army operators who will work with the aviation regiment and ground forces.

This would accelerate what the Army has begun and move it beyond the limited, incremental experimentation seen to date. It would also create a much more informed user base within the Army, which will be vital for understanding the combat potential of other autonomous systems. In this way, the Army can achieve the vision in its RAS strategy and position its people as the 'thought leaders' and early adopters of autonomous systems that are complementary to the predominantly manned platform force structure across the ADF.

As we move forward and Defence becomes increasingly reliant on autonomous systems with short lifespans and rapid refresh cycles, it will need to redesign its investment program and move away from funding models that rely on platforms serving 30-year life spans. Rather, it will need to program ongoing funding lines so that there's assured funding for experimentation and the acquisition of rapidly evolving technologies. We've seen this approach already in areas such as special forces capability under Project Greyfin.[51]

There's no reason why an evolving Army drone capability couldn't be funded in the same way. We suggest that $1 billion of the funding saved by not acquiring a manned ARH should be programmed in the investment plan to acquire unmanned autonomous systems for battlefield aviation.

As a capability risk mitigation strategy should these systems not be able to deliver all the effects sought, Defence could also invest some funding (say, $300 million) in acquiring a broader range of small, precision-guided munitions for the Reaper and other capabilities already in service or currently being acquired. Defence should explore whether those munitions can be domestically designed and manufactured. The munitions could also potentially be integrated into the autonomous battlefield aviation systems.

These measures total under $3 billion. That's substantially less than the funding currently programmed for a manned Tiger replacement, but it will drive much greater innovation in autonomous systems and other areas of emergent technology than the acquisition of an off-the-shelf manned helicopter.

Other opportunities beyond Tiger

The strategy and capability reviews currently underway in Defence present the opportunity to redefine programs other than the ARH in the investment plan, even if the Defence Minister's public caveat that no changes are expected to the major programs remains in place. There are potentially many other programs in which Defence could acquire rather than observe others acquiring autonomous systems, some more mature and some less so.

If Defence sees itself as a truly innovative organisation and wants to act accordingly, the default approach shouldn't be to automatically replace manned systems with improved manned systems. It should seek to obtain the best capability mix that balances operational risk with operational success, and that will mean seriously considering the contributions autonomous systems can make without imposing the constraint that they do what previous manned systems did.

In most capability areas, the least-risk path will be to introduce autonomous systems to complement legacy manned platforms, although, as those new types of capabilities mature, the case for continuing numerous manned platforms will get harder and harder to maintain.

The key is to view lines of funding in the investment plan as capability programs rather than projects to acquire particular platforms. For example, both the Collins-class submarine life-of-type extension and the Future Submarine Program should be regarded as undersea warfare programs, not simply platform projects. One way to ensure that the Collins remains a relevant undersea warfare capability is to disaggregate and enhance some of its individual capabilities by investing in complementary unmanned underwater vessels that can act as sensors for it.

Similarly, rather than wringing out the last increments of performance from the manned conventional submarine concept at enormous cost, Defence could use some of the Future Submarine Program's massive budget (or even just part of its $397 million in science and technology funding) to progress unmanned systems where vast technological improvements can still be made at limited cost.

Under either approach, Defence could work cooperatively with the US Navy and acquire some of its Orca long-range underwater vessels to see how they can be used disruptively. A small number of Orcas penetrating deep into adversaries' waters to lay smart sea mines, for example, would impose significant operational risk on their ships and submarines. We can be sure that they're thinking about doing that to us. Or Defence could acquire large numbers of small surface or subsurface vessels to flood the battlespace with sensors.

There are likely to be many other areas where Defence can make leaps towards the fielding of autonomous weapons systems. The Tiger replacement offers an opportunity to jump-start this process, build momentum, help Defence's people develop trust in autonomy, and prepare themselves for a conceptual shift in war fighting to Australia's advantage.

NOTES

1 On the broader use of AI in military contexts, see Christian H Heller, ‘The future Navy—near-term applications of artificial intelligence’, *Naval War College Review*, Autumn 2019, 72(4), article 7, online.

2 For example, Mark V Arena, Irv Blickstein, Obaid Younossi, Clifford A Grammich, *Why has the cost of navy ships risen? A macroscopic examination of the trends in US naval ship costs over the past several decades*, RAND Corporation, Santa Monica, 2006, 15, online.

3 On the cost of the future submarine and frigates, see Marcus Hellyer, *The cost of Defence: ASPI defence budget brief 2018–2019*, ASPI, Canberra, 2018, online, in particular, Chapter 6; Andrew Davies, Marcus Hellyer, ‘The very hungry future submarine’, *The Strategist*, 5 November 2019, online.

4 See, in particular, Peter Layton, *Prototype warfare, innovation and the fourth industrial age*, Air Power Development Centre, Department of Defence, Canberra, 2019, online.

5 During the drafting of this study, 13 French soldiers were killed in an accident on operations in Africa when a Tiger helicopter collided with a Cougar transport helicopter; AFP/Reuters, ‘Thirteen French soldiers killed in helicopter collision in Mali’, *ABC News*, 26 November 2019, online.

6 Vietnam Helicopter Pilots Association, *Helicopter losses during the Vietnam War*, 31 December 2018, online.

7 ‘List of aviation shootdowns and accidents during the Iraq War’, *Wikipedia*, online.

8 Department of Defence (DoD), *2016 Integrated Investment Program*, Australian Government, 2016, 119, online.

9 Based on a unit cost of around US$32 million each (from the US Army’s FY 2020 budget justification book—online), it’s likely that the total acquisition cost for a fleet of 29 Apaches and all necessary fundamental inputs to capability would approach A$3 billion once we take exchange rate, escalation and contingency (i.e. risk provisions) into account. The cost of 24 MH-60R maritime combat helicopters was around A$2.5 billion (according to the 2019–20 Defence Portfolio Budget Statements).

10 The following Senate Estimates committee illustrates the difficulty in modifying aircraft such as the Tiger: ‘Mr Fairweather: You’re quite correct. The technology was developed, and then to look at the integration onto the Tiger’s very complex airframe and to install that equipment on the airframe is a very expensive and lengthy proposition. We need to understand how it fits and operates on the aircraft aerodynamically and also how it affects it when the aircraft deploys into the water. That requires a lot of analysis around the carriage of the aircraft and the physical and fatigue elements of it by fitting that piece of equipment to the aircraft—the way it’s structurally mounted on. The Tiger structure is very complex and to understand the fatigue and structural impacts of those loads on the airframe need a lot of detail. So, it is going through the certification process to make sure that by putting that piece of equipment on the aircraft, which may address one safety risk, is not going to introduce other safety risks that actually have a worse case.’ Senate Foreign Affairs, Defence and Trade Legislation Committee Estimates, *Hansard*, 5 April 2019, 33, online.

11 Clayton M Christensen, *The innovator’s dilemma: when new technologies cause great firms to fail*, Harvard Business School Press, Boston, Massachusetts, 1997.

12 David H Berger, *Commandant’s planning guidance*: *38th Commandant of the Marine Corps*, US Marine Corps, July 2019, online.

13 Congressional Budget Office, *An analysis of the Navy’s fiscal year 2020 Shipbuilding Plan*, October 2019, online.

14 Megan Eckstein, ‘Navy betting big on unmanned warships defining future of the fleet’, *USNI News*, 8 April 2019, online.

15 Congressional Research Service, *Navy large unmanned surface and undersea vehicles: background and issues for Congress*, updated 15 November 2019, online.

16 Sharon Weinberger, ‘China has already won the drone wars’, *Foreign Policy*, 10 May 2018, online.

17 Australian Army, *Robotic and Autonomous Systems Strategy*, October 2018, online.

18 On 28 November 2019, the Minister for Defence and Minister for Defence Industry announced that the government had down-selected the General Atomics MQ-9B Sky Guardian as the Reaper variant to be acquired by Defence. The media release also stated that ‘government consideration’—presumably second-pass approval to actually acquire the system—wouldn’t occur until 2021–22; Linda Reynolds, Melissa Price, ‘Cutting edge remotely piloted platform chosen in billion dollar project’, joint media release, 28 November 2019, online. Defence confirmed at Senate Estimates committee hearings on 29 November 2019 that initial operating capability was scheduled for 2025–26; Senate Foreign Affairs, Defence and Trade Legislation Committee Estimates, 29 November 2019, 23, online. The US Air Force has had the Reaper in service since 2007 and an armed version of its predecessor, the Predator, since around 2001, so Australia will be around 25 years behind the US in introducing an armed remotely piloted aircraft into service.

19 The Reaper and the Predator have also been shot down.

20 Liberal Party of Australia, 'Boost for WA economy and jobs with new naval commitment', media release, 29 April 2019, online; Marcus Hellyer, 'Morrison's shipbuilding announcements are about more than jobs', *The Strategist*, 30 April 2019, online. The US Navy is also retiring its manned minehunters in favour of autonomous systems operated from littoral combat ships or vessels of opportunity; David Axe, 'America is getting rid of their minesweepers', *The National Interest*, 3 April 2019, online.

21 Shervin Shahrdar, Luiza Menezes, Mehrdad Nojoumian, 'A survey on trust in autonomous systems', Computing Conference 2018, 10–12 July 2018, London.

22 The programs' priorities and investments awarded are described in detail in Department of Defence, *Defence industry and innovation: 2017–18 defence industry and innovation programs annual report*, Australian Government, Canberra, 2019, online. According to the report, 51% of the Next Generation Technologies Fund's investment for that year went into trusted autonomous systems, which sounds like a lot, but it was less than $7 million for the year.

23 Trusted Autonomous Systems Defence Cooperative Research Centre, online.

24 As context, see Martin Callinan, *Defence and security R&D: a sovereign strategic advantage*, ASPI, Canberra, 2019, online.

25 The *2016 Defence White Paper* has 13 references to a technology 'edge', leading 'edge' technologies, and leading edge innovation, as well as references to technological advantages in the region.

26 See, for example, Senate Foreign Affairs, Defence and Trade Legislation Committee Estimates, *Hansard*, 24 October 2018, 62.

27 Department of Defence, *Budget 2019–20. Budget Related Paper No. 1.4A. Defence portfolio*, Canberra 2019, Table 55, online.

28 Shardar et al., 'A survey on trust in autonomous systems', 7.

29 US Government Accountability Office, *Weapon systems cybersecurity: DOD just beginning to grapple with scale of vulnerabilities*, October 2018, online.

30 The example is used by Connor McLemore, Charles Clark, 'The devil you know: trust in military applications of artificial intelligence', *War on the Rocks*, 23 September 2019, online.

31 On the difficulty of trusting truly autonomous systems, see Heather M Roff, David Danks, '"Trust but verify": the difficulty of trusting autonomous weapons systems', *Journal of Military Ethics*, 2018, 17(1):2–20.

32 Steven E Davis, *Individual differences in operator's trust in autonomous systems: a review of the literature*, DST-Group-TR-3587, Edinburgh, 2019, 4.

33 Including the AH-64 Apache: 'Army's helicopter-drone teams to get capability boost in 2019', *DefenseNews*, 8 October 2017, online.

34 And this brings trust issues of its own; if the AI doesn't trust a human to achieve a certain level of performance, it may assign excessive amounts of human resources to ensure the job gets done, or not send manned platforms at all. Trust cuts both ways.

35 Australian National Audit Office (ANAO), *Management of the Tiger Armed Reconnaissance Helicopter Project—Air 87*, Performance Audit report no. 36 of 2005–06, Australian Government, Canberra, 2 May 2006, online; ANAO, *Tiger—Army's armed reconnaissance helicopter*, Performance Audit report no. 11 of 2016–17, Australian Government, Canberra, 1 September 2016, online.

36 According to media reporting, Defence received responses to the RFI from Airbus Australia, Boeing (proposing its AH-64E Apache), and Bell and BAE Systems (proposing the Bell AH-1Z Viper).

37 Senate Foreign Affairs, Defence and Trade Legislation Committee Estimates, 30 May 2017, Canberra, 26.

38 Stephen Jobson, 'Tiger—Armed reconnaissance. Attack. Escort. Land. Amphibious. Maritime. Air. Special Forces. Protecting and saving lives', *LinkedIn*, 13 July 2018, online.

39 'Aussie Tiger first to achieve 2K hours', *Australian Defence Magazine*, 15 May 2017, online.

40 Roger Brennan, 'Tigers take flight on IPE19', news release, Department of Defence, Canberra, 16 May 2019, online.

41 During which it was probably lasered by Chinese vessels in the South China Sea; Euan Graham, 'Australian pilots hit with lasers during Indo-Pacific exercise', *The Strategist*, 28 May 2019, online.

42 'Brig. King: The intention was to piggyback off other nations who were going to be the lead. Only this month our first Tiger became the first Tiger that reached 2,000 aircraft flying hours. So no other Tiger in the world has reached 2,000 hours. Our aircraft was the first to reach that point. That was never our intended strategy.' Senate Foreign Affairs, Defence and Trade Legislation Committee Estimates, *Hansard*, 30 May 2017, Canberra, 23.

43 The request for information documentation is no longer on AusTender.

44 Jared Keller, 'The Air Force's MQ-9 Reaper drone just got a whole lot deadlier', *Task & Purpose*, 9 May 2017, online.

45 'Closed ATM view2—HSD/RFI/9467/1, LAND 129 Phase 3 Tactical unmanned aerial system replacement or enhancement', *AusTender*, online.

46 DoD, *2016 Integrated Investment Program*, 112.

47 Speech by Senator Linda Reynolds, Royal Australian Navy Sea Power Conference, Sydney, 8 October 2019, online.

48 Australian Army, *Robotic and Autonomous Systems Strategy*, 20.

49 David Epstein, *Range: how generalists triumph in a specialized world*, Macmillan, 2019.

50 Two classic examples being the French Army's failure to capitalise on the mitrailleuse machine-gun and its failure to recognise the true potential of the tank until that was demonstrated in particularly abrupt fashion by the Wehrmacht.

51 Michael Shoebridge, 'Special forces' approach to technological change a model for others', *The Strategist*, 13 August 2019, online.

ACRONYMS AND ABBREVIATIONS

ADF	Australian Defence Force
AI	artificial intelligence
ANAO	Australian National Audit Office
ARH	armed reconnaissance helicopter
ISR	intelligence, surveillance and reconnaissance
LHD	landing helicopter dock
R&D	research and development
RAAF	Royal Australian Air Force
RAS	robotic and autonomous system
RFI	request for information
UAS	unmanned aerial system
UAV	unmanned aerial vehicle
UUV	unmanned underwater vehicle

Some recent ASPI publications

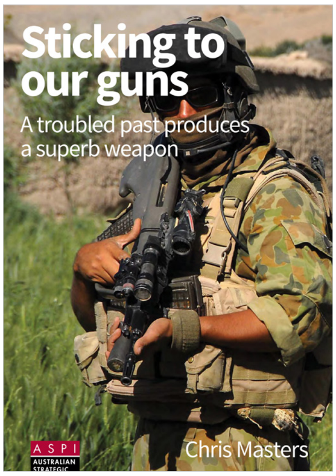

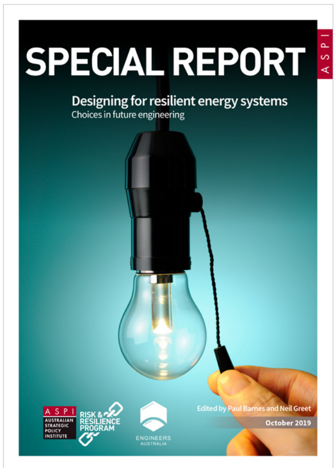

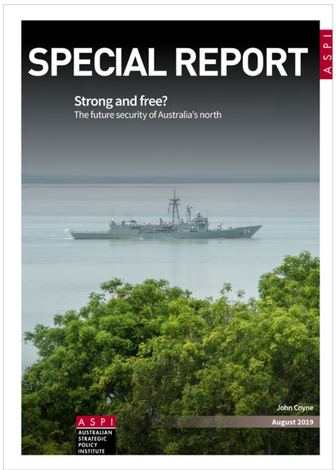

Accelerating autonomy
Autonomous systems and the Tiger helicopter replacement

ISSN 2200-6648