STRATEGIC STUDIES QUARTERLY

Autonomy and the Future Force

Wg Cdr Andrew Massie, RAF

Abstract

While autonomy is decision making independent of outside control, delegating authority for successfully dispersed and disaggregated operations is antithetical to our current practice. At one end of the spectrum of the human–machine interface is remote control—human input to generate a direct machine response with no authority granted to the machine to decide and act. At the opposite end of the spectrum, recourse to human supervision is absent and the machine intelligence can be exploited to its maximum potential by being freed to react to its environment. This is also the regime where the Department of Defense (DOD) would face the greatest organizational and cultural challenges in exploiting autonomy. The irony is that to harness the full potential of autonomy, we have to trust machines and free decision makers.

✳ ✳ ✳ ✳ ✳

> *As our understanding of the history of technology increases, it becomes clear that a new device merely opens a door; it does not compel one to enter. The acceptance or rejection of an invention, or the extent to which its implications are realized if it is accepted, depends quite as much upon the condition of a society, and upon the imagination of its leaders, as upon the nature of the technological item itself.*
>
> —Lynne White Jr.
> *Medieval Technology and Social Change*

In framing the third offset strategy as being centered upon human–machine collaborative combat networks, Deputy Secretary of Defense Bob Work recognized a social and technology trend that will undoubtedly

---

Wg Cdr Andrew Massie is currently serving as an exchange officer in the strategy cell of the Headquarters USAF plans directorate (HAF/A5SS). He graduated from the School of Advanced Air and Space Studies at Air University. Wing Commander Massie is a Royal Air Force pilot and has varied combat experience, including a deployment to Afghanistan.

have a huge impact upon humanity.[1] The challenge, as historian Lynn White, Jr., proffers, is the extent of our ability to turn this concept into concrete combat capability. If the DOD wants to grasp this new idea and use it to strategic advantage, leaders must seize the opportunity to shape the narrative about machine autonomy and help create a future based on strong US Air Force (USAF) contributions to the multidomain fight. Clearly delegating authority needed for successfully dispersed and disaggregated operations is antithetical to our current practice. Autonomy and its attendant benefits can only be achieved by a change in human–machine relationships to one of mission command. At its core, our ability to harness autonomy is a test of our ability to *trust* machines and, therefore, to delegate authority for decision making and action. Generally this will entail less control and more observation for machines and men.

The deputy secretary has presented five building blocks for this kind of autonomy; however, as they stand, these blocks merely describe a spectrum of activity that ranges from machines that think to machines that think and act. While differentiating between physical and cognitive tasks is important, recognizing environmental complexity and the implications of adversary responses is more important for the DOD. The department must develop a framework to articulate the differing types of tasks and, therefore, highlight those areas where autonomy is a "natural fit" and those where more work is required to inculcate trust or apply safeguards necessary for human–machine collaboration to succeed. This article will therefore propose a framework for understanding autonomy, based upon the nature of the environment in which a task is conducted, to determine the relative propensity for humans to trust machine outputs and therefore employ them effectively. It will then consider the implications of accepting autonomy as a source of strategic advantage in the third offset strategy against great-power adversaries. Ultimately, our ability to recognize and harness the positive opportunity autonomy offers will determine our ability to reap the benefits information technology offers. For this reason, an appreciation of the fundamentals of autonomy is crucial for the DOD to step forward with confidence. To start the process of shaping the future force, we must first clearly articulate what we mean by autonomy.

What is Autonomy?

The Industrial Revolution augmented and substituted manual human labor with machine labor.[2] The implications for the conduct of war were tremendous growth in speed of maneuver, the destructive power of combat forces, and the development of military bureaucracies to manage delivery of military forces on a huge scale. Beyond simple linear growth, the Industrial Revolution—along with later development of the internal combustion engine, the jet engine, and rocket propulsion—enabled powered flight and access to outer space. As we stand at the dawn of an Information Revolution, information technology promises a comparable exponential advantage to that offered by machine over manual labor—but this time in machine cognition and data computation over the human brain. The advantage of the search engine, like the jet engine previously, may dwarf the gains currently conceivable.

The 2015 Defense Science Board (DSB) Summer Study task force went a long way toward describing *how* machine autonomy might offer the DOD a competitive advantage and, therefore, *why* it should be broadly accepted; however, in describing autonomy's use the task force omitted a definition of *what* autonomy is. Without this crucial appreciation, the military professional lacks the insight necessary to generate an informed understanding of autonomy's potential and pitfalls. According to the DSB, autonomy "results from the delegation of a decision to an entity which is authorized to take action within specific boundaries."[3] The crucial takeaway from this definition is that to be autonomous is to be free to make decisions without external intervention. In essence, harnessing autonomy is a test of one's willingness to relinquish control. Under this definition, a broad array of machine tasks can therefore be termed autonomous.

Additionally, we must highlight the critical strengths of the human in the human–machine team and be aware of the *irony of automation*: in a worst case scenario, if we expect a human to step in and override a system, that person requires all of the situational awareness and skill needed to conduct the task absent the machine.[4] So, if the cost of maintaining a large workforce was the driver in accelerating autonomy, the irony of automation might make us reevaluate the expected benefits.

The Machine Autonomy Framework

Since autonomy is the delegation of decision making, a critical facet of USAF understanding of the use of autonomy is related to the question of trust. Like all human interactions, decision making and trust go hand in hand. With a choice, we will give the most responsibility to those whom we believe most capable of conducting a task. Mission command involves communicating intent and an appreciation for *why* a task has been set but does not determine *how* it must be conducted; a competent subordinate will exercise their best judgment dependent upon the circumstances. However, when delegating authority, we set bounds on the activity our subordinates undertake. Approaching one of these boundaries invokes the need to report up the chain for clarification or further guidance. Therefore, supervision is inherent in any command relationship and will vary with circumstance and task complexity. The same logic is true for machine as for man.

As autonomy is decision making independent of outside control, it is critical we recognize there are degrees of autonomy just as there is a spectrum of tasks to be conducted; therefore, the bounds that we place on authorized actions determines the degree of autonomy afforded.[5] At one end of the spectrum of the human–machine interface is remote control—human input to generate a direct machine response. In this instance, no authority is granted to the machine to decide and act; it merely responds directly to a human input. The control philosophy for Reaper or Predator remotely piloted vehicles (RPV) would be illustrative of this interface. In this case, an action is not *autonomous* but *controlled*. It is a direct response to a deliberate stimulus with no need to make an independent decision.

Somewhere in the middle of the spectrum is a machine that can assess its environment, prioritize a list of possible solutions to a problem, rank them, and request an operator's input. The machine can harness the advantages of rapid data manipulation, but a human supervisor is necessary to determine the actual course of action undertaken. Anyone familiar with the health monitoring systems in modern aircraft, such as those tracking fuel or engine performance, will be wholly familiar with the value of this type of activity in reducing operator workload. An extension to this level of collaboration is the recognition that a machine might conduct the task required, such as the routine balancing of fuel between tanks to

maintain aircraft center of gravity, but faced with a nonstandard problem, the decision to act will be commanded by a human.

At the opposite end of the spectrum is a machine afforded the latitude to assess its surroundings, trawl its database for possible responses, rank and weigh those responses, determine the optimal course of action, then enact its derived course of action. As an example, a computer virus detection mechanism or firewall is a capability that should be activated then left to perform its task independently. Recourse to human supervision is absent, and the machine or machine intelligence can be exploited to its maximum potential by being freed to read and react to its environment. It is undoubtedly to this end of the spectrum that most autonomy detractors are drawn and where the specter of the "killer robot" exists. Coincidentally, this is also the regime where the DOD would face the greatest organizational and cultural challenges in exploiting autonomy.

Tasks and Trust

The development of a useful understanding of the spectrum of tasks and their associated levels of trust requires a framework to distinguish between the nature of differing military tasks and the intendant effects upon the need for human supervision.[6]

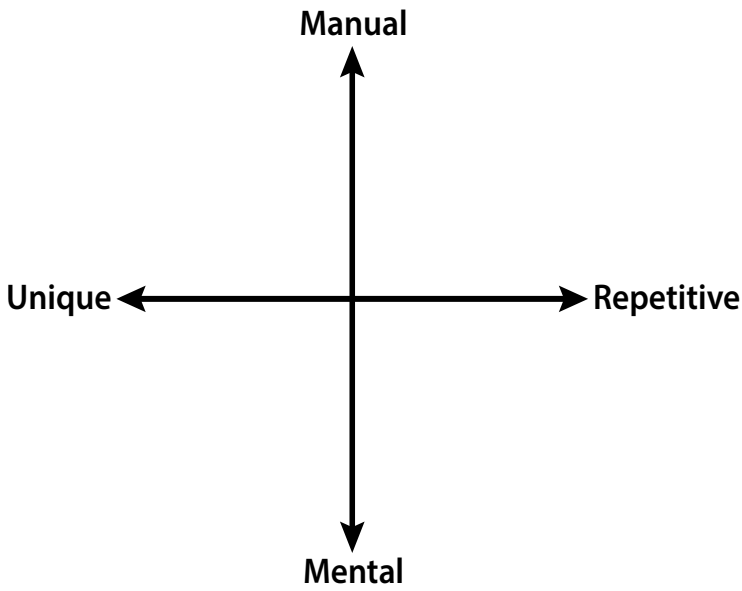

Figure 1. A framework for task classification

Along the horizontal axis in figure 1, tasks can be determined by environmental novelty. Those tasks on the left-hand side of the chart are described as unique and those on the right as repetitive. A variation in task or variation in the environment largely determines the changing factor along the horizontal axis. Due to these two factors, repetitive tasks are those where the environment in which the task is conducted and the task's output is unchanging. On the other hand, unique tasks are conducted in a changing and unpredictable environment or reflect a demand for variable outputs depending upon a specific requirement.

Crucial to the application of autonomy in military affairs is the recognition of the roles of unpredictability and adversarial action in conflict. While the advantages afforded to industrial production facilities are obvious examples of a manual repetitive task (the right-hand side of the chart), the battlefield and shop floor are dichotomous due to the presence of a reacting adversary (the left-hand side of the chart). In integrating autonomous machines into our inventory we must recognize the presence of a thinking and noncooperative actor as the baseline standard for interface in many military tasks—a concept Clausewitz articulated on the first page of book one of *On War*.[7]

Figure 2 deductively shows the implications of environment novelty upon the level of human–machine collaboration. Where outcome certainty is low, trust will be low, and the need for human supervision will be high to ensure the expected task is conducted appropriately. While this will undoubtedly change with time, in the near term, it is intuitive to say one will have low trust of machine decision-making success in complex changing environments and, therefore, will need to ensure a high degree of human supervision. A current example of this is the level of human supervision applied in the operation of the MQ-1 and MQ-9 RPVs. High environmental uncertainty, low trust, and high human supervision lead us naturally to a default human–machine relationship of strict control and, at its most extreme, remote control—or nil autonomy. While it may sound trite, the experience of any new instructor pilot with a novice student will attest to the desire to be prescriptive and offer direct commands over a more laissez-faire approach: the instructor's "skin is in the game," and mission success dictates this default human response. With experience and exposure comes greater subtlety in response. The same will be true of our interaction with machines over time.

The vertical axis in figure 1 differentiates between machine output. Mental tasks are referred to in the DSB study as "autonomy at rest," while manual tasks are referred to as "autonomy in motion." On the same vertical axis, we see these dual possibilities information technology offers: one is machines used to do lower-level mental tasks; the other is one of empowering machines themselves to enter the human realm of decision making in some limited capacity. With mental tasks, artificial intelligence (AI) offers the opportunity to harness the power of data computation to perform tasks that free the human to exercise unique attributes of creativity and intuition. On the other hand, by pairing AI with robotics, we gain the ability to advance the power of machine labor with machine cognition.

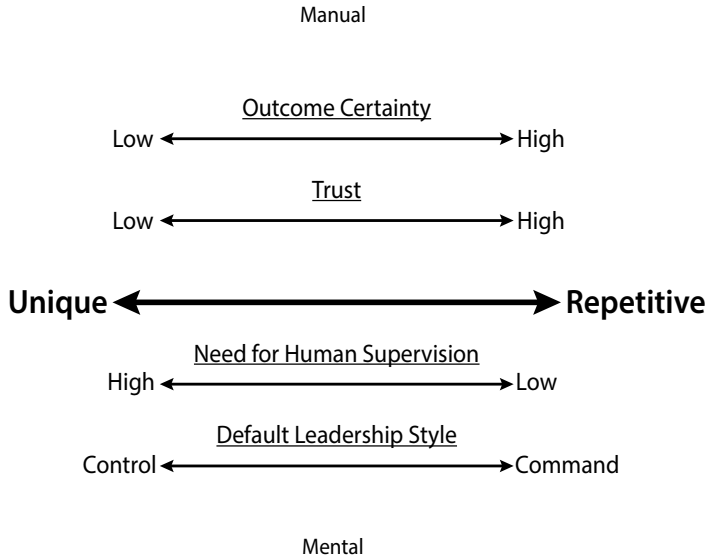

Figure 2. Insights on the horizontal axis

Where outcome certainty is high, trust is high, as the machine can comfortably and reliably meet the task. The autopilot used in climb, cruise, and descent of modern airliners is a clear example of our confidence in machine decision making and action. The need for human supervision exists but is low, and the default human–machine interaction can be "hands-off"—to command a range of activities and then sit back and monitor. However, this need not be a benign environment. In a high–intensity peer conflict, autonomy may be delegated to defensive systems, such as a Patriot battery, to scan a cleared free-fire area, detect

movement, determine a threat through prescribed algorithms, and engage on machine command. The important point to note is the inverse relationship between our confidence in an outcome and the need for direct input: low confidence equals control, high confidence equals command but with more autonomy.

Commanding or Controlling

The application of this two-axis approach to different types of tasks illuminates significant insight for the DOD and the USAF. Regardless of whether the task is mental or manual, it is clear that the novelty of the environment in which a machine (or human) is operating is a significant discriminator in terms of autonomy. This should be no surprise, as at its core, autonomy is a question of delegated decision making: the more novel an environment, the more challenging to delegate authority. In complex, wide–area security operations over the last 15 years, we have learned the hard lesson that a significant amount of trust must be afforded lower echelon decision makers to achieve operational and tactical goals. Higher echelons must take greater risk in freeing units to exploit the increase in situational awareness and fleeting advantage. The same lesson will apply to mission command for machines and will necessitate a gradual lessening of restrictions, through training for human supervisors and better and more-rapidly programmable machine decision-making code. As the British strategist J. F. C. Fuller noted, "The more mechanical become the weapons with which we fight, the less mechanical must be the spirit which controls them."[8]

The "teams" or relationships we form with machines will therefore be largely determined by environmental novelty—or in military terms, proximity to an adversary. The more our environment favors repetitive, manual tasks—such as base logistics—the greater opportunity for machine automation. Similarly, where analysis warrants the assessment of longer-term trends and activity, the better suited our analysis will be to machine intelligence. As we approach contact with an adversary and environmental novelty increases, we are in the realm of tacit knowledge and rapid environmental assessment. As a recent study by Oxford University and Citigroup noted on the implications of autonomy in the workplace, this is specifically the area where human interaction will hold preeminence.[9]

Human preeminence need not mean machine absence; indeed, this may be the greatest value of Deputy Secretary Work's emphasis on autonomy.

As in all technology endeavors, robotics and AI may provide significant advantage by augmenting or amplifying human activity. Rather than seeing the human–machine interaction as a zero-sum or an either-or relationship, we must find the synergy between the man and machine. Wearable technology and robot assistants, or "co-bots" (collaboration robots), offer the synthesis of the best of both worlds—the interaction of human intuition and tacit or social knowledge with machine intelligence and manual strength.[10] In this regime, the area of interest will be the nature of the interaction or relationship—just as in our use of animals to perform military tasks. It may well be, similar to an attack dog, the human commands the machine to act and employ its strength to the team's advantage. Alternatively, and more conceptually challenging, like the explosive-sniffing dog, the machine may lead the human to action. It is undoubtedly in the development of teams and co-bots that the benefits of autonomy will be decisive militarily. In doing so, we must be prepared to lead, to trust, and to follow.

Implications for the Third Offset

Clearly there are cultural, practical, and political challenges facing autonomy in enhancing military advantage. Conversely, the enormous benefits that come with pairing machine cognition with machine labor are apparent to the military practitioner. Indeed, it has been articulated by the deputy secretary as the single greatest advantage, in concert with an educated workforce, the United States can leverage against its likely adversaries. The current description, interestingly, seeks to differentiate between tasks by the manual-mental "output" that are **absent** unique-repetitive environmental complexity. Those differentiated tasks are depicted in figure 3 and explained below.[11]

- **Learning Machines or Systems** represent machine decision making on a network that allows machines to learn from and communicate with each other in order to counter machine attacks such as a cyber virus. Learning machines maximize machine task autonomy with minimal human supervision but perform a wholly cognitive and virtual function, such as Google's "Deep Mind" system. This concept also recognizes that cyber weapons may be employed at a speed too great for human response; machine defense may be essential to counter machine offense.

- **Human–Machine Collaboration** represents a situation in which machines benefit from huge databases to highlight patterns and trends to facilitate human decision making. This is a largely cognitive task that requires human action to translate data to an action. An example may be the development of a digital "air-operations planner" that monitors all air mission activity and battle damage assessment on operations and presents alternate courses of action to the combined force air component commander (CFACC) for the next day's air tasking order or dynamic solutions to an unfolding significant event.

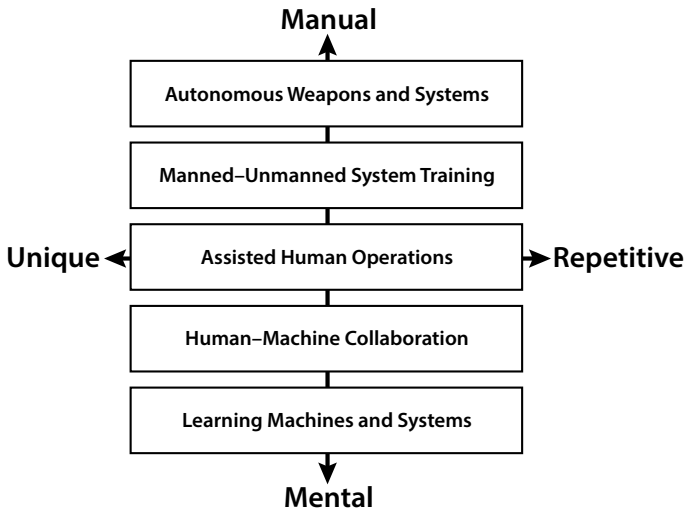

Figure 3. The "Big Five" using the framework for task classification

- **Assisted Human Operations** are tasks with similar output to human–machine collaboration but with a greater emphasis on deliverables or wearable hardware at the tactical level of war. As an excellent example, the Air Force Future Operating Concept (AFFOC) offers an aerial resupply port of the future, where networked supply chains in real time across an area of responsibility prioritize and palletize aircraft loads based upon evolving theater priorities.[12] The only science-fiction element to this vignette is its military application: this is a business practice widely employed by commerce giants such as Amazon and Walmart today.
- **Manned–Unmanned System Teaming or Human–Machine Combat Teaming** deals with tasks consisting of physically cooperating human and autonomous systems on the battlefield. Human interaction

and supervision is still necessary for mission success, albeit in a limited capacity. The clearest example for the USAF is the integration of autonomous wingmen into the combat air forces to enhance lethality or situational awareness. The AFFOC vignettes on a future close air support and air superiority mission invoke autonomous wingmen in concert with a manned combat platform, allowing man and machine to contribute their requisite strengths to the benefit of the overall mission—increasing payload, survivability, and the merits of disaggregated command and control (C2) to grasp fleeting changes in local conditions.

- **Autonomous Weapons and Systems** represent tasks that benefit from all four layers previously described to apply learning machines to advanced robotics and deliver a machine that is able to conduct its task against a reacting adversary without human input. While this may seem far-fetched for an air force that has engaged in 15 years of wide-area security operations, conducting high-tempo operations in a highly contested environment offers a very different operating concept. If the United States were able to embrace autonomous weapons in defense of currently vulnerable and distant operating bases, with much greater emphasis on early detection and engagement, the tyranny of distance might be repainted as an opportunity. With clear delineation between friend and foe, clear fire corridors for autonomous kinetic, cyber, and electronic-warfare weapons might offer a decimating form of defense against any potential aggressor.

The obvious takeaway from placing these five capabilities on a quad chart that shows the vulnerability to adversarial action is that there immediately are undoubtedly huge benefits to the military application of learning systems, human–machine collaboration, and assisted human operations. Indeed, during the last 15 years, many of these benefits are already being exploited in understanding enemy networks and their subsequent targeting. Furthermore, cyber defense already rests largely in learning systems and human–machine collaboration. Similarly, those who have worked on exchange tours with industry would recognize these five capabilities are widely used and see that the DOD could undoubtedly do more to employ such abilities. The advantage Deputy Secretary Work

seeks will be realized when this cognitive computational power can be reliably delivered into a machine that also performs the task at hand.

It is in autonomous weapons systems and manned–unmanned system teaming that the most benefit can be derived but the greatest military risk exists. Machines promise significant opportunities in delivering lethality and performance beyond that of the limited human physiology. However, their application is fraught with risk due to the question of outcome certainty and the necessity to monitor them. The niche for autonomous weapons systems does exist, but its fragility to adversary action or, conversely, the time and cost of development is significant. Thus, human creativity will continue to be essential in delivering battlefield success against reacting and intelligent adversaries. As the recent evaluation of Google's AlphaGo machine algorithm against a human expert demonstrated, learning machines come with significant advantages. Such machines are guaranteed to perform to expert levels when fielded and will continue to learn thereafter. However, in a crucial one-off engagement, like combat, they can be undermined by genius or confused by human error.[13]

The answer lies, as in most polemics, somewhere in the middle—in advancing the concept of manned–unmanned system teaming to determine where full mission autonomy might be granted, under specific rules of engagement (ROE) or circumstance, and where the final determination of action must rest with a "man in the loop" or on the spot. The emphasis must be upon teaming or the appropriate mix of interaction that generates the greatest military advantage.

The final critical deduction from a study of autonomy is the promise and challenge of disaggregated and dispersed operations. As a facet of the third offset, the necessity to operate in a highly-contested environment, using networks of platforms to defeat massed firepower, is a robust deduction. However, there are grave limitations between that mode of operating and our current C2 structure. A generation of leaders has lived in an operational environment where risk has been held at a fairly high level and decision making for the use of lethal force has been largely held with higher echelons. ROE do exist for tactical action, but they have been extremely constrained. Operating with greater emphasis on command, rather than control, will be challenging but not insurmountable. Significant capital must be expended in training and simulation to prepare commanders to grant their machinery more autonomy, and more importantly, this way of thinking must be inculcated into USAF

leaders such as CFACCs. If the adversaries we expect to face take the battlefield, the long screw driver will be consigned to history, and the strategic corporal and captain will own the day. This may well be as great a cultural challenge, in releasing the reins, as the simple introduction of the technology itself. The challenge we face is that in an Information Age war, the initial moves may be so debilitating that little time is available to adapt or react. Our drive to field centrally controlled, exquisite capabilities over networked, disaggregated, human–machine mission capacities may deny us a second chance and may be so cost prohibitive as to deter action. Being "not too wrong" necessitates a balanced capability mix to allow an opportunity to adapt rapidly to a threat environment.

Conclusion

As the venerated British general Graeme Lamb noted about leadership in complex environments, the solution to a future characterized by autonomy may be to operate "in command, but out of control."[14] When it comes to autonomy, the third offset is as much about software, or organizational culture and concepts, as it is equipment. Any discussion of autonomy must capture and leverage this insight. An important inference is that leaders, decision makers, and planners will lead *and* follow; they must become comfortable in both roles as humans guiding and following autonomous systems.

Autonomous machines, like people, offer greater potential with increased latitude in determining their own course of action. The challenge with men or machines is trusting their judgment in a complex and contested environment. In this final regard, we hold a significant advantage. Western militaries have a long history of devolved command responsibility. This autonomy for man and machine is an opportunity to adapt in contact and may well be our unique advantage against the most likely peer in an era of information age war. While a conversation on autonomy may drive the audience to the subject of hardware and equipment, it is clear that building trusting organizational constructs is as, or more, important. The ultimate irony may be that to get the most from our machines, we have to free our men and women. SSQ

Notes

1. Sydney J. Freedberg Jr., "People, Not Tech: DepSecDef Work on 3rd Offset, [Joint Interagency Combined Space Operations Center] JICSPOC," *Breaking Defense*, http://breakingdefense.com/2016/02/its-not-about-technology-bob-work-on-the-3rd-offset-strategy.

2. Yuval Noah Harari, *Sapiens: A Brief History of Humankind* (New York: Harper, 2015), 334.

3. Defense Science Board (DSB), Department of Defense (DOD), *DSB Summer Study on Autonomy* (Washington, DC: DOD, July 2015 [publication forthcoming]), 5.

4. Lisanne Bainbridge, "Ironies of Automation," *Automatica* 19, no. 6 (1983): 775–79, http://www.ise.ncsu.edu/nsf_itr/794B/papers/Bainbridge_1983_Automatica.

5. For a more complete and useful descriptor see Sheridan and Verplank's 10 point levels of automation contained in Liang Sim, M. L. Cummings, and Cristin A. Smith, "Past, Present and Future Implications of Human Supervisory Control in Space Missions," *Acta Astronautica* 62, no. 10–11 (2008): 648–55, http://dx.doi.org/10.1016/j.actaastro.2008.01.029.

6. "Automation Angst," *Economist* (UK), 16 August 2015, http://www.economist.com/node/21661017.

7. Carl von Clausewitz, *On War*, edited and translated by Michael Howard and Peter Paret (Princeton, NJ: Princeton University Press, rev. 1984), 75.

8. J. F. C. Fuller, *Generalship: Its Diseases and Their Cure; a Study of the Personal Factor in Command* (Harrisburg, PA: Military Service Publishing, March 1936), 13, https://archive.org/details/GeneralshipItsDiseasesAndTheirCure.

9. Carl Benedikt Frey et al., *Technology at Work v2.0: The Future Is Not What It Used to Be* (New York: CitiGroup and Oxford Martin School, January 2016), http://www.oxfordmartin.ox.ac.uk/downloads/reports/Citi_GPS_Technology_Work_2.pdf.

10. Ibid., 92.

11. Robert Work, "Deputy Secretary of Defense Speech: Reagan Defense Forum: The Third Offset Strategy" (speech, Reagan Defense Forum, Reagan Presidential Library, Simi Valley, CA, 7 November 2015), http://www.defense.gov/News/Speeches/Speech-View/Article/628246/reagan-defense-forum-the-third-offset-strategy.

12. *Air Force Future Operating Concept: A View of the Air Force in 2035*, September 2015, 27. http://www.af.mil/Portals/1/images/airpower/AFFOC.pdf.

13. Choe Sang-Hun, "South Korean Gets 'Priceless' Victory over Computer in Go Match," *New York Times*, 13 March 2016, http://www.nytimes.com/2016/03/14/world/asia/south-korean-gets-priceless-victory-over-computer-in-go-match.html.

14. Graeme Lamb, "In Command and Out of Control," *Medium.com* (web site), 21 October 2015, https://medium.com/the-bridge/in-command-and-out-of-control-aab523b92fe1.

Disclaimer