such as bombing campaigns, military procurement or keeping inventories. Each impacted only a few researchers.

The actors in operations research were trained but not bound to any disciplinary identity, their methods and the problems that they tackled differed. This too adds to the variety in operations research that will be surprising to anyone expecting a coherent story of emergence. The work of operations scientists, or the promotion of the application of a 'scientific method', ended up entering business, systems engineering, procurement military and otherwise, nuclear strategy and policy. Its practitioners were based in universities, businesses, and most famously in think tanks like the RAND corporation in the United States.

The story that Thomas fleshes out here is an important story because the development of operations research (as we know it today) was a significant political development of the last century that still impacts how policy is made by politicians, the military, and even businesses today. The book offers a fascinating account of how mathematics was brought to bear on policy and economics in ways few would have predicted and which makes us question beliefs about the rationality of policy long held and rarely reflected upon. It is also of extra relevance in times in which expert authority to make policy is being vocally rejected by large portions of the population.

Hermione Giffard, Utrecht, The Netherlands

Nick Bostrom, *Superintelligence: Paths, Dangers, Strategies*. Oxford: Oxford University Press, 2014. Pp. 328. £18.99.

An intriguing mix of analytic philosophy, computer science and cutting-edge science fiction, Nick Bostrom's *Superintelligence* is required reading for anyone seeking to make sense of the recent surge of interest in artificial intelligence (AI). This was the book that inspired Stephen Hawking, Elon Musk, Bill Gates, and other prominent technologists to make public statements of concern about the pace and direction of AI R&D. Indeed, it is hard to come away from Bostrom's treatise without sharing at least some of their alarm: *Superintelligence* articulates in fine analytic style a variety of philosophical arguments for the inevitability of global domination by intelligent machines. Whether this leads to humanity's extinction or fulfilment of its 'cosmic endowment' depends, for Bostrom, on the kind of AI humanity builds.

Bostrom states up front that he is not doing the Singularity; but think of this as Singularity 2.0 – the focus on 'accelerating change' and 'exponential growth' has given way to a more formal treatment of technological change

premised on these notions. *Superintelligence* is thus of historical interest, as it introduces a robust conceptual terminology for thinking and talking about AI that is likely to prove influential for years to come, but there is little history here. What there is – Chapter 1 opens with a global overview of 'human progress', measured in GDP – will give historians much to critique.

The topic is greater-than-human or *super*-intelligence. Bostrom discusses a variety of paths to superintelligence, including 'cognitive enhancement' through positive eugenics, but his book is really about *machine* super-intelligence. Given the high speed, small size and large capacity of electronic digital components, machines offer the only viable means of exceeding the information processing capacity of the human brain and thus achieving true *super*intelligence. Put quibbles that computation does not equal intelligence aside – the possibility that an algorithm for 'general intelligence' (g) can and will be discovered is axiomatic to this work of philosophical speculation. If one can agree that there is something called g that is responsible for humans' ability to perform a wide variety of cognitive tasks, then it is possible that computer programmers may one day capture the essence of g in algorithmic form, as 'artificial general intelligence' (AGI). The main premise of the book is this: such an AGI would be capable of recursively reprogramming itself, thereby improving its g, and after many iterations of this process, the AGI might 'take off' into an artificial *super*intelligence (ASI), an entity orders of magnitude more intelligent than anything currently in existence. What do we do then?

Bostrom is at his best when describing the consequences of such an evolution. He predicts the first ASI to emerge from such an 'intelligence explosion' will possess an array of 'cognitive superpowers' that will present its human developers with a classic technological dilemma: How to use the ASI without losing control of it? *Superintelligence* provides a quasi-technical catalogue of ASI control methods and guidance through some of the unintuitive consequences of ASI, such as the 'perverse instantiation', wherein even a 'friendly' ASI may interpret its instructions too literally and plug humanity into *The Matrix* when given instructions to 'promote human happiness', for example. *Superintelligence* argues that once an ASI exists, world domination – either by it, or by the nation in control of it – is logically inevitable. Bostrom calls the resultant monopolar global order a 'singleton' and stresses that the ASI control problem can only be solved by an efficient singleton capable of marshaling sufficient resources to its solution. The book concludes with a warning that failure to solve the ASI control problem will result in existential catastrophe for all of humanity as out-of-control ASI reshapes the surface of the Earth to better suit its needs.

Given the risks, why allow AGI/ASI R&D to proceed? *Superintelligence*

argues that the future of humanity is at stake because the creation of ASI is actually *necessary* to realize our 'cosmic endowment', or full potential. For Bostrom this is the creation of self-replicating space probes that harvest whole planets to build stellar-scale computers powered by sun-enveloping solar-panelled Dyson spheres, in order to run trillions of simulations of happy humans, in perpetuity or until the universe runs down in heat death. In this calculus, *not* creating ASI would thus kill more (virtual) people than have ever lived on Earth, and therefore the potential for so many (simulated) lives to go unrealized counterbalances the risk of existential catastrophe posed by ASI, thus nullifying the argument for relinquishment.

Damned if we do, damned if we don't? Or as Chapter 8 is titled, 'Is the default outcome doom?' Unfortunately, it is in working through the socio-political dimensions of machine intelligence that *Superintelligence* is at its weakest. It was surprising to find that at the end of a book with the subtitle, *Paths, Dangers, Strategies,* the most cogent strategy offered for dealing with what many see as the most likely outcome of AI – deepening socioeconomic inequality as the global economy struggles to adapt to the increasingly rapid automation of production – is merely a 'windfall clause', wherein a corporation voluntarily agrees to distribute any profits in excess of 'say a trillion dollars annually' to the billions of members of humanity evenly. If only corporations were so superintelligent.

Superintelligence opens onto a vast space of possibilities and leaves plenty of room for further work on the problems therein; next-generation technologists will certainly find it useful. However, the shortcomings of apolitical positivism and greedy reductionism are on full display here. If one cannot abide the reification of intelligence into a single quantity, and remain sceptical of claims that it can be amplified along a linear scale of improvement, then the fantastic prognostications of *Superintelligence* may prove entirely irrelevant and even repugnant.

Colin Garvey, Troy, NY, USA

Ronald R. Kline, *The Cybernetics Moment: or Why We Call Our Age the Information Age*. Johns Hopkins University Press, Baltimore. 2015. Pp. 336. $54.95.

Ronald R. Kline's *The Cybernetics Moment* is a valuable addition to the history of cybernetics in two key ways. The first is heralded by the book's title – Kline's working thesis is that to date, most authors dealing with cybernetics have neglected to really account for either its momentary character, or