The risks of accelerated change

The pace of technological change is accelerating — in artificial intelligence, materials research, engineering, biology: everywhere. Through the Internet, technology now spreads globally as well, and faster than ever. It's also being democratized, as developers find new ways to make technology useful to an ever-wider range of communities, including non-technical people.

It's all great, in many ways, yet rapid change also brings new risks. Artificial intelligence could evolve beyond our control or, less dramatically, replace countless human workers, bringing social and economic instability. Techniques of gene editing might become so easy that anyone may soon be able to hack and modify their own germline, opening huge ethical challenges. Of course, we face looming unknown risks from cybercrime and cyberwarfare.

Humanity has survived until now in no small part because it has found ways to regulate technology to keep the worst negative consequences — things like pollution, weapons and crime — within tolerable bounds. We've used laws, social norms and international agreements. In the case of nuclear weapons, we've also relied on a clear conceptual understanding of the collective suicidal nature of their use. Of course, we're currently struggling to organize ourselves to protect the Earth's environment and to counter CO_2 emissions and global warming.

But our current difficulties could signal something more profound. Growing research centres on the possibility that technological acceleration could propel us into a vastly more unstable era, accompanied by a general reduction in our regulatory effectiveness. A recent analysis of this idea comes from physicists Dimitri Kusnezov and Wendell Jones, who frame the issue using game theory (preprint at http://arxiv.org/abs/1707.06668; 2017), which has played an important role in helping us survive the nuclear age so far.

As traditionally considered, game theory provides knowledge about possible stable outcomes in simple games of strategy involving one or more agents. The classic concept of the Nash equilibrium identifies sets of strategies that, once discovered by those agents, provide a stable fixed point for a game, as no agent has any incentive to depart from their current strategy. The peaceful equilibrium rests on a clear understanding of the mutually assured destruction that would follow from any attempted first strike.

Technological growth could propel us into a more unstable era, with a reduction in our regulatory effectiveness.

As Kusnezov and Jones note, game theory has also helped give form to the bulk of modern international agreements and regulations. As they argue, the negotiation of treaties and regulatory regimes involves the building of a "shared awareness of the characteristics of the choice-space facing the parties", and allows a kind of low-risk, non-violent exploration of the space of available strategies — a search for possible Nash equilibria. A stable outcome and sustainable agreement reflects the mutual discovery of such an equilibrium.

But the usefulness of the Nash equilibrium concept rests on the assumption that agents can understand the logical structure of a situation well enough to discover possible stable strategies. That's only realistic for relatively simple problems, in which agents have ample time to develop a full understanding. In situations of greater complexity — or in games where the rules keep changing — equilibria may be impossible to find, and thus of no practical relevance. As times change and technologies grow more complex, the physicists suggest, we may find this old model no longer helpful.

So what happens in more complex games? This has been explored over the past few years in studies of randomly structured games with P players, each using N possible strategies, where P and N can be arbitrarily large. In such games, players try to learn and adapt, choosing which strategy to use at any moment based on its past performance, and altering their choices along the way in response to the shifting behaviours of other players. Game outcomes in general display a wide range of qualitatively distinct behaviours, sometimes settling into an attracting Nash equilibrium, but in others ending in persisting limit cycles or ongoing chaotic evolution. A sharp boundary in parameter space separates stable (non-chaotic) and unstable (chaotic) regimes of behaviour (J. B. T. Sanders, J. D. Farmer and T. Galla, preprint at http://arxiv.org/abs/1612.08111; 2016), depending on P, N and parameters reflecting the agents' learning process.

As Kusnezov and Jones argue, the stable side of this boundary corresponds to classic game theory, where Nash equilibria reflect likely outcomes. In contrast, our near-future of democratized, rapidly changing technologies look much more likely to live inside the chaotic zone of the space, reflecting game theory beyond Nash equilibria. What does this imply about how we might cope in a more turbulent world?

Kusnezov and Jones don't offer any easy answers, but several immediate lessons follow from their framing. The first is that we're unlikely to find successful regulatory regimes by copying methods from the past. Anything based on shared exploration to find a Nash equilibrium won't work, simply because complexity and the speed of change makes such equilibria increasingly impossible to discover. The second is that we urgently need more study of this regime of complex games, to build up some replacement for traditional game theory, with new concepts attuned to irregularity and chaos.

There's a curious analogy here to fluid turbulence, as the authors of the study note. The general character of a fluid flow follows from the Reynolds number, a non-dimensional ratio of inertial forces to viscous forces. Low Reynolds numbers give regular flows, and high numbers invariably lead to turbulence. By this analogy, the current technological explosion may be driving a kind of transition to turbulence in social and political dynamics. Here a different quantity plays the role of the Reynolds number, one reflecting the complexity of the strategic game, and the increasing difficulty of discovering Nash equilibria.

Just as turbulent flows can't be managed with the methods of laminar fluid dynamics, we're not likely to find the right ways to manage a world of complex technology with ideas developed for a simpler setting. We're going to need something altogether new. ❐

MARK BUCHANAN