

# Efficiency vs. Resilience: Lessons from COVID-19



Moshe Y. Vardi

**Abstract** Why was the world not ready for COVID-19, in spite of many warnings over the past 20 years of the high likelihood of a global pandemic? This chapter argues that the economic goal of efficiency, focused on short-term optimization, has distracted us from resilience, which is focused on long-term optimization. Computing also seems to have generally emphasized efficiency at the expense of resilience. But computing has discovered that resilience is enabled by redundancy and distributivity. These principles should be adopted by society in the “after-COVID” era.

By March 2020, COVID-19 (Coronavirus disease 2019) was spreading around the world. From a local epidemic that broke out in China in late 2019, the disease has turned into a raging pandemic, the likes of which the world has not seen since the 1918 Spanish flu pandemic. By then, thousands have already died, with the ultimate death toll growing into the millions. Attempting to mitigate the pandemic, individuals were curtailing travel, entertainment, and more, as well as exercising “social distancing,” thus causing an economic slowdown. Businesses hoarded cash and cut spending in order to survive a slowdown of uncertain duration. These rational actions by individuals and businesses were pushing the global economy into a deep recession.

Observing the economic consequences of this unexpected crisis, William A. Galston asks in a March 2020 *Wall Street Journal* column<sup>1</sup> “What if the relentless pursuit of efficiency, which has dominated American business thinking for decades, has made the global economic system more vulnerable to shocks?” He went on to argue that there is a tradeoff between efficiency and resilience. “Efficiency comes through optimal adaptation to an existing environment,” he argued, “while resilience requires the capacity to adapt to disruptive changes in the environment.”

---

<sup>1</sup><https://www.wsj.com/articles/efficiency-isnt-the-only-economic-virtue-11583873155>

---

M. Y. Vardi (✉)  
Rice University, Houston, TX, USA  
e-mail: [vardi@cs.rice.edu](mailto:vardi@cs.rice.edu)

A similar point was made by Thomas Friedman in a May 2020 *New York Times* column:<sup>2</sup> “Over the past 20 years, we’ve been steadily removing man-made and natural buffers, redundancies, regulations and norms that provide resilience and protection when big systems—be they ecological, geopolitical or financial—get stressed. . . . We’ve been recklessly removing these buffers out of an obsession with short-term efficiency and growth, or without thinking at all.”

Both Galston and Friedman were pointing out that there is a tradeoff between short-term efficiency and long-term resilience. This tradeoff was also raised, in a different setting, by Adi Livnat and Christos Papadimitriou (2016). Computational experience has shown that simulated annealing, which is a local search—via a sequence of small mutations—for an optimal solution, is, in general, superior computationally to genetic algorithms, which mimic sexual reproduction and natural selection. Why then has nature chosen sexual reproduction as almost the exclusive reproduction mechanism in animals? Livnat and Papadimitriou’s answer is that sex as an algorithm offers advantages other than good performance in terms of approximating the optimum solution. In particular, sexual reproduction favors genes that work well with a greater diversity of other genes, and this makes the species more adaptable to disruptive environmental changes, that is to say, more resilient.

The tradeoff between efficiency and resilience can thus be viewed as a tradeoff between short-term and long-term optimization. Nature seems to prefer long-term to short-term optimization, focusing on the survival of species. Indeed, Darwin supposedly said: “It’s not the strongest of the species that survives, nor the most intelligent. It is the one that is most adaptable to change.”

And yet, we have educated generations of computer scientists on the paradigm that analysis of algorithm only means analyzing their computational efficiency. As Wikipedia states:<sup>3</sup> “In computer science, the analysis of algorithms is the process of finding the computational complexity of algorithms—the amount of time, storage, or other resources needed to execute them.” In other words, efficiency is the sole concern in the design of algorithms. (Of course, the algorithm has to meet its intended functionality.) *The Art of Computer Programming*,<sup>4</sup> a foundational text in computer science by Donald E. Knuth, is focused solely on efficiency. What about resilience? Quoting Galton again: “Creating resilient systems means thinking hard in advance about what could go wrong and incorporating effective countermeasures into designs.” How can we make our algorithms more resilient?

Of course, fault tolerance has been part of the canon of computing-system building for decades. Jim Gray’s 1998 Turing Award citation<sup>5</sup> refers to his invention of transactions as a mechanism to provide crash resilience to databases. Leslie Lamport’s 2013 Turing Award citation<sup>6</sup> refers to his work on fault tolerance in

<sup>2</sup><https://www.nytimes.com/2020/05/30/opinion/sunday/coronavirus-globalization.html>

<sup>3</sup>[https://en.wikipedia.org/wiki/Analysis\\_of\\_algorithms](https://en.wikipedia.org/wiki/Analysis_of_algorithms)

<sup>4</sup>[https://en.wikipedia.org/wiki/The\\_Art\\_of\\_Computer\\_Programming](https://en.wikipedia.org/wiki/The_Art_of_Computer_Programming)

<sup>5</sup>[https://amturing.acm.org/award\\_winners/gray\\_3649936.cfm](https://amturing.acm.org/award_winners/gray_3649936.cfm)

<sup>6</sup>[https://amturing.acm.org/award\\_winners/lamport\\_1205376.cfm](https://amturing.acm.org/award_winners/lamport_1205376.cfm)

distributed systems. Nevertheless, computer science has yet to fully internalize the idea that resilience, which to include reliability, robustness, and more, must be pushed down to the algorithmic level. Case in point is search-result ranking. Google's original ranking algorithm was PageRank,<sup>7</sup> which works by counting the number and quality of links to a page to determine how important the website is. But PageRank is not resilient to link manipulation, hence "search-engine optimization."

As pointed up by Friedman and Galston, the relentless pursuit of economic efficiency prevented us from investing in getting ready for a pandemic, in spite of many warnings over the past several years, and pushed us to develop a global supply chain that is quite far from being resilient. Does computer science have anything to say about the relentless pursuit of economic efficiency? Quite a lot, actually.

Economic efficiency means<sup>8</sup> that goods and factors of production are distributed or allocated to their most valuable uses and waste is eliminated or minimized. Free-market advocates argue<sup>9</sup> that through individual self-interest and freedom of production and consumption, economic efficiency is achieved and the best interests of society, as a whole, are fulfilled. But efficiency and optimality should not be conflated. The First Welfare Theorem,<sup>10</sup> a fundamental theorem in economics, states that under certain assumptions a market will tend toward a competitive, Pareto-optimal equilibrium; that is, economic efficiency is achieved. But how well does such an equilibrium serve the best interest of society?

In 1999, Elias Koutsoupias and Papadimitriou undertook (Koutsoupias and Papadimitriou 1999) to study the optimality of equilibria from a computational perspective. In the analysis of algorithms, we often compare the performance of two algorithms (e.g., optimal vs. approximate or offline vs. online) by studying the ratio of their outcomes. Koutsoupias and Papadimitriou applied this perspective to the study of equilibria. They studied systems in which non-cooperative agents share a common resource and proposed the ratio between the worst possible Nash equilibrium and the social optimum as a measure of the effectiveness of the system. This ratio has become known as the "Price of Anarchy,"<sup>11</sup> as it measures how far from optimal such non-cooperative systems can be. They showed that the price of anarchy can be arbitrarily high, depending on the complexity of the system. In other words, economic efficiency does not guarantee the best interests of society, as a whole, are fulfilled.

A few years later, Constantinos Daskalakis, Paul Goldberg, and Papadimitriou asked (Daskalakis et al. 2006) how long it takes until economic agents converge to an equilibrium. By studying the complexity of computing mixed Nash equilibria, they provide evidence that there are systems in which convergence to such equilibria can take an exceedingly long time. The implication of this result is that economic

---

<sup>7</sup><https://en.wikipedia.org/wiki/PageRank>

<sup>8</sup>[https://www.investopedia.com/terms/e/economic\\_efficiency.asp](https://www.investopedia.com/terms/e/economic_efficiency.asp)

<sup>9</sup><https://www.investopedia.com/terms/i/invisiblehand.asp>

<sup>10</sup>[https://en.wikipedia.org/wiki/Fundamental\\_theorems\\_of\\_welfare\\_economics](https://en.wikipedia.org/wiki/Fundamental_theorems_of_welfare_economics)

<sup>11</sup>[https://en.wikipedia.org/wiki/Price\\_of\\_anarchy](https://en.wikipedia.org/wiki/Price_of_anarchy)

systems are very unlikely ever to be in an equilibrium, because the underlying variables, such as prices, supply, and demand, are very likely to change while the systems are making their slow way toward convergence. In other words, economic equilibria, a central concept in economic theory, are mythical rather than real phenomena. This is not an argument against free markets, but it does oblige us to view them through a pragmatic, rather than ideological, lens.

Our digital infrastructure, which has become a key component of the economic system in developed countries, is one of the few components that did not buckle under the stress of COVID-19. Indeed, in March 2020, many sectors of our economy switched in haste to the WFH mode, “working from home.” This work from home, teach from home, and learn from home were enabled (to an imperfect degree, in many cases) by the Internet. From its very roots of the ARPANET in the 1960s, resilience, enabled by seemingly inefficient distributivity and redundancy, was a prime design goal for the Internet (Yoo 2018). Resilience via distributivity and redundancy is one of the great principles of computer science that deserves more attention by the business community.

In summary, resilience is a fundamental, but under-appreciated, societal need. Both computing and economics need to increase their focus on resilience. It is important to note that markets and people tend to underprepare for low-probability or very long-term events. For example, car insurance is inefficient for insurance holder, though it offers resilience. Yet people are required to purchase car insurance, because many will not buy it otherwise. In other words, societal action is required to ensure resilience. It is important to remember this point, as many now argue that COVID-19 is just the “warm-up act” for the climate crisis.<sup>12</sup>

The big question is how the AC (“after-COVID”) world will differ from the BC (“before-COVID”) world. Fareed Zakaria wrote<sup>13</sup> in *The Washington Post* in October 2020: “The pandemic upended the present. But it’s given us a chance to remake the future.” Matt Simon wrote<sup>14</sup> in *Wired* in December 2020: “The COVID-19 pandemic has brought incalculable suffering and trauma. But it also offers ways for people—and even societies—to change for the better.” I believe that resilience must be a key societal focus in the AC world.

---

<sup>12</sup><https://www.chronicle.com/article/covid-19-is-just-the-warm-up-act-for-climate-disaster>

<sup>13</sup><https://www.washingtonpost.com/opinions/2020/10/06/fareed-zakaria-lessons-post-pandemic-world/>

<sup>14</sup><https://www.wired.com/story/who-will-we-be-when-the-pandemic-is-over/>

## References

- Daskalakis, C., Goldberg, P.W., and Papadimitriou, C.H. (2006) The complexity of computing a Nash equilibrium. *STOC 2006*: 71-78
- Livnat, A. Papadimitriou, C.H. (2016) Sex as an algorithm: the theory of evolution under the lens of computation. *Commun. ACM* 59(11), pp. 84-93
- Koutsoupias, E. and Papadimitriou, C.H. (1999) Worst-case Equilibria. *Proc. 16th Annual Symposium on Theoretical Aspects of Computer Science, Lecture Notes in Computer Science 1563*, Springer, pp. 404-413
- Yoo, C.S. (2018) Paul Baran, Network Theory, and the Past, Present, and Future of the Internet. *Colo. Tech. LJ* 17, p.161.

**Open Access** This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

