The Perils of Path Dependence: Lessons from the Nuclear Power Industry

By Kevin A. Bryan

[text] In the 1950s and 1960s, scientists in government and at some of the world's leading industrial firms, including General Electric (GE), Westinghouse, Babcock & Wilcox, Dow, and Union Carbide, worked to develop a new industry: nuclear power. Research was pursued on more than a dozen types of nuclear reactors, each of which implied different paths of future technological development, different complementary products, different regulatory regimes, and different engineer training. Direct U.S. research subsidies via the Atomic Energy Commission (AEC) totaled nearly $40 billion (2013 dollars) during these two decades. Private sector research, indirect subsidies, and military nuclear spending, which inevitably spilled over into civilian programs, far exceeded even that total. This lavish research spending did indeed develop commercially viable nuclear power, with worldwide capacity reaching nearly 50,000 megawatts by 1973, the majority of which was generated by a single style of nuclear plant called "light water". Despite this promising start, however, the nuclear industry would fade to unimportance: not a single plant broke ground in the United States between 1977 and 2013.

This outcome is surely troubling for any firm operating in an evolving technological field. Extensive research spending, on the advice of highly qualified technical minds and with incredible government support, had nonetheless resulted not just in failed firms, but a completely collapsed industry. Worse, this collapse appears to have been at least partially the result of the industry settling on an inferior technological design. As it turns out, the light water reactors that came to dominate the worldwide market were widely considered an inferior technology, both in terms of economic efficiency and safety, compared to alternative reactors that were not fully developed. The belief that light water would prove an inferior technology was held both by researchers and planners in the early 1950s, before the commercial dominance of light water was secured, and by the hindsight of analysts decades later.[1]

Was this simply bad luck? In other words, were there simply unanticipated changes outside the industry that made light water less viable? Did, for example, changes in global military power, government regulation, or scientific advances make alternative technologies more lucrative? In some sense, a story of "bad luck" would not be very interesting. After all, firms always invest under uncertainty about what the future might bring, and ex-ante good investments may look bad after the fact. But this was not the case with the nuclear industry. Rather, the dominance of inferior light water, and the less lucrative industry that followed in its wake, was foreseeable as the outcome of strategic interactions between companies that were harmful for the industry at large, despite being rational for each individual firm. As such, the history of the nuclear industry holds important lessons for companies in any technologically dynamic industry, and it also has implications for marketing, products with network externalities, and investments in networked supply chains.

[subhead] A Brief History

In the early 1950s, U.S. Naval Captain Hyman Rickover selected a pressurized light water reactor as the most promising candidate for a nuclear-powered submarine largely because it was thought, on the basis of very limited technical research, to offer a quick development time. Early experience working on such light water reactors for the Navy unquestionably led GE and Westinghouse to prioritize research along those lines when developed nuclear power reactors, but there was still substantial variety in reactor types being developed by the private sector into the 1960s.

These reactor types involved radically different design choices, with different reactors involving different "coolants" and "moderators" to sustain the nuclear reaction and transform its energy into electricity, different types of fissile material to use as fuel, different methods to control the reaction and the radiation it produces, and so on.[2] Although certain technical problems with light water had been solved while developing the nuclear submarine, even in the early 1950s the question of which reactor type would prove most efficient and profitable in the long run was far from settled. In fact, a previously classified AEC report noted that the light water reactor "seemed most likely to be successful in the short term, by the end of 1957, but it offered a poor long-term prospect of producing economic nuclear power."[3] In other words, the technical difficulties with generating nuclear power were least daunting for light water plants, but inherent limits on the efficiency of those types of plants meant that they would cost more to run when compared with fully developed alternatives.

Participating utilities wanted plants with low operating costs. The problem, though, was that the government wanted plant designs that would generate useful research results. This conflict caused costs to rise rapidly and discouraged utilities from buying early demonstration plants using technologies with greater long-term promise than light water.[4] The result was a market standstill: by 1962, neither government programs nor the salesforces at reactor developers like GE and Westinghouse had persuaded utilities to purchase plants on a commercial basis.

Then, from 1963 to 1966, GE and Westinghouse began to offer "turnkey contracts" for fully built light-water reactors.[5] The goal was to induce lower costs via learning-by-doing, and the plan worked. After 1962, every nuclear plant built in the United States was ordered from one of five companies, four of which focused on light water: GE, Westinghouse, Babcock and Wilcox and Combustion Engineering. The fifth – Gulf Oil – had hoped to usurp the existing industry leaders by developing high-temperature gas-cooled reactors, but a number of technical problems proved insoluble.[6]

Light water reactors thereafter became quickly ensconced in the market due to increasing returns to scale. That is, as the installed base of light water reactors grew, the unit cost of providing electricity from that technology fell, which then led to increasing sales of those reactors, thereby reinforcing the cycle of falling costs. By the mid-1970s, over 70 percent of the world's operating power reactors were light water, with the Canadian heavy water reactor the only alternative still commercially available.[7] Not only did costs plummet relative to other reactor types – just between 1958 and the beginning of the turnkey program, the cost of electricity generated by light water reactors had fallen over 80 percent – but the availability of technical support and safety information also had benefited from increasing returns to scale.[8] Nevertheless, as noted earlier, sales of light water reactors began to dramatically slow toward the mid-1970s (even before the Three Mile Island and Chernobyl disasters). Light water, even in late stages of development, did not achieve the cost efficiency that nuclear power optimists had hoped the industry would reach after 25 years of expensive development. With alternative reactor types receiving only a portion of the research and development given to light water, the economic failure of light water plants ultimately led to the collapse of the nuclear power industry.

[subhead] The Power of Dynamic Increasing Returns

Given this history, three facts need explaining. First, why did firms initially develop light water, a technology that was unlikely to offer the best long-term mix of costs and benefits? Second, as light water began to fail commercially, why didn't firms switch to alternative reactor types? And third, how did government policy affect the behavior of firms?

An understanding of the role of dynamic increasing returns is essential to answering those three questions. According to that powerful market mechanism, the value of using a technology of type A rather than type B can grow over time depending on how many other firms or customers have used technology A or, alternatively, how much experience the firm has with that technology. The existence of dynamic increasing returns implies that relatively minor differences in conditions when technologies are new can "lock in" inferior technologies.[9]

The most famous such story is that of the Stanley Steamer. In the early 20th century, it wasn't obvious at all that gas-powered automobiles were superior to steam-powered ones. In fact, steam-powered cars were quieter, provided a smoother ride, and handled more easily.[10] However, they also required water to be added roughly every 30 miles, and an outbreak of foot-and-mouth disease in New England in 1914 forced the closure of water troughs,[11] leading to customers temporarily shifting away from steam-powered vehicles. The result was a larger user base of gas-powered cars, which then led to more gas stations and mechanics who could work on those vehicles. This growing convenience helped increase the value of gas-powered automobiles, leading to more customers choosing them even after the water troughs had reopened. The eventual dominance of gas power, then, might have occurred even if most drivers preferred steam-powered automobiles when the installed base of each type of vehicle was identical.[12]

Note the importance of increasing returns: if the value of a car were independent of the number of other similar vehicles sold, then a "small event" like the closing of water troughs in New England would not have had such a huge effect. In the nuclear case, the selection of a light-water reactor for the U.S. nuclear submarine gave GE and Westinghouse early experience in building that type of reactor, and learning-by-doing certainly generates increasing returns. That is, the more reactors of a given type that have been built, the easier engineers will find it to improve that technology and build additional reactors of that type. It's important to note here that the choice made by Captain Rickover in the early 1950s was not in any way based on light water being the superior technology for the long-run development of a civilian nuclear energy industry. As such, that decision was the sort of "small event" that can tip subsequent actions down an inferior technology path.[13]

Increasing returns – and hence the potential for lock-in – appear in many forms beyond simple learning-by-doing. For example, existing technologies generate a series of contracts and investments by firms upstream and downstream that remain as constraints even when an alternative technology appears to be a strict improvement. This explains the substantial difficulty that electric cars will have in usurping the dominant position of gas-powered vehicles.[14] Network effects, or the benefits of coordination on a single standard by producers or buyers, can also lock in inferior early movers.[15] In addition, "bad luck" during the early stages of research that attempts to learn about the viability of alternative research agendas in the face of uncertainty can lead to firms focusing on technologies that have a few "lucky draws" in their early development.[16] Lastly, increasing returns can result from small events that affect the development of complementary technologies for suppliers or buyers, minor differences in preferences between early and late adopters, and small gifts of luck along different research lines.

Are "small events" the only source of dynamic increasing returns that can lead an industry down a potentially inefficient path? If they were, we would again be in a situation of "bad luck": there is nothing General Electric or Union Carbide could have done if Rickover's idiosyncratic choice as a military bureaucrat had locked in light water technology. However, the widespread adoption of an inferior technology often occurs from dynamic increasing returns that are entirely within the control of the companies involved. Specifically, lock-in can inadvertently result from the *deliberate* actions of *rational* firms.

[subhead] When Rational Actions Lead to Disastrous Results

Why would firms take actions that deliberately lock in an inferior technology? The answer comes from the disconnect between a firm maximizing its own profits, and an industry trying to maximize the profits of all its firms. Consider the following two examples.

Two firms, A and B, are competing to develop a commercial nuclear power plant. They can either put their researchers to work on a gas-graphite or a light-water demonstration plant. If firm A works on gas-graphite reactors, the research will take two years. With the demonstration plant as a model, a full-size gas-graphite plant could then be developed by both firms, with firm A's expertise in the demonstration plant enabling it to earn $6 billion and firm B to earn $3 billion. If, however, either firm works on a light-water demonstration plant, it can be built in only one year, after which both firms would find it in their interest to build full-scale light-water plants. In this case, the inventor of the first demonstration plant will earn $4 billion and the other firm will earn $2.5 billion. In both cases, the differential earnings reflect the relative ease of scaling up the demonstration plant by its inventor. Assume that, in either case, once a demonstration plant has been built, dynamic increasing returns mean that both firms will find it most profitable to shift all of their research onto scaling up that demonstration plant.

What should firm A do? It will reason as follows. "If I work on the gas-graphite reactor, then firm B will have its scientists develop the light-water plant as they will finish their demonstration plant before I finish mine. And once their demonstration plant is complete, I will direct my scientists to scale up the light-water plant, earning me $2.5 billion. Alternatively, I can simply start by working on a light-water demonstration plant, giving me roughly a 50/50 shot of being first to develop it, and hence earning $4 billion half the time, and $2.5 billion the other half, for an average expected profit of $3.25 billion. Therefore, I am going to direct my scientists to work on the light-water plant." Firm B will, of course, reason in precisely the same way, and hence both firms will work on light water even though they both accurately perceive that gas graphite would maximize industry profits.

The problem is a *strategic externality*: once the demonstration plant for light water is built by any firm, it's no longer worth continuing development of the alternative technology. In essence, rational firms do not care that their research today affects the value of research projects initiated by rivals. That is, the effort of other firms generates increasing returns to scale that will affect the future profits of rival companies, but those other firms rationally will not take that into account.

It gets worse. Consider a second example in which the two firms can either put their scientists to work on a full-scale nuclear plant, or on a demonstration plant that can be sold only in small numbers. If the full-scale plant is developed directly, the inventor earns $8 billion in the eventual market and the rival earns $1 billion, with the difference reflecting the large gap in knowledge around building full-scale plants in the early stages of a new industry. However, if the demonstration plant is developed first, the inventor earns $1 billion and the eventual market in full-scale plants will be split between both firms. In this case, the openness in development will lead to the full-scale market itself to be worth $10 billion, with $5 billion going to each firm. Note that industry profits are maximized when the demonstration plant is invented first and then developed into a full-scale plant by the industry at large; the assumption here is that it is most efficient for the industry as a whole to test things on a demonstration plant before committing to a costly full scale plant. Assume further that, because the full-scale plant is more difficult to develop, if one firm works on a full-scale plant and the other on a demonstration plant, the full-scale plant will be invented first 1/3 of the time, and the demonstration plant will be invented first 2/3 of the time.

In this scenario, firm A will reason in the following way. "If firm B works on a demonstration plant and I work on the full-scale plant, then I will invent my plant first 1/3 of the time, earning $8 billion; the other 2/3 of the time, my rival will invent the demonstration plant, and I can then switch my scientists to scaling up that project, earning $5 billion. The average expected profit is therefore $6 billion. If, however, I work on the demonstration plant, I will invent it before my rival half the time, earning an average expected profit of $500 million from selling demonstration plants to technologically savvy utilities, and will earn another $5 billion from developing that demonstration plant into a full-scale industry, for an overall average expected profit of $5.5 billion. So, if firm B works on a demonstration plant, I ought to work on the full-scale plant."

What if firm B works on a full-scale plant? “If that happens and I work on a demonstration plant, I will finish first 2/3 of the time, earning the $1 billion profit from selling that early-stage plant, and another $5 billion from helping develop that demonstration plant to scale. But 1/3 of the time, my rival will invent the full-scale plant before me, and I will earn $1 billion. My total average expected profit is therefore $4.33 billion. If, however, I work on a full-scale plant, then I finish the full-scale plant first half the time, earning $8 billion dollars, and my rival finishes that plant first half the time, leaving me to earn $1 billion dollars, hence giving me an expected profit of $4.5 billion. Therefore, no matter what firm B does, I should work on the full-scale plant.” Again, firm B will reason the same way, and hence the industry will inefficiently develop the more difficult but less lucrative full-scale plant. The problem is a different type of strategic externality: rational firms do not account for the fact that their research makes possible follow-on research by rivals, and hence they care only about the profitable value of follow-on research that they can capture directly.

The lesson of these examples is that competitive firms can have incentives to endogenously introduce – that is, to create as a result of their own rational incentives – the increasing returns to scale that will lock in subpar technologies. Earlier workhorse models of path dependence[17] have described the lock-in of inferior technologies as a result of unforeseeable “small events” like the personal preferences of certain customers. But the market dynamics of increasing returns to scale can result either from the initial advantage of an inferior technology generated by small events, or from the active and rational choices of competing firms! It is the latter that is worrying: unlike the case where industries are harmed by the idiosyncracies of Hyman Rickover, or the randomness of a foot-and-mouth outbreak, there are things firms in a growing industry can do to avoid the peril of a path dependence creating by their own choices.

[subhead] Six Ways to Avoid Path Dependence

What, then, should firms in technologically dynamic industries do? The concern, of course, is that competitors may have the incentive to pursue research strategies that decrease, via path dependence, the value of the industry. Or, worse, what to do when those strategies would lead the industry toward complete collapse? The following six lessons are worth noting.

1) Coordinate to reduce strategic incentives that lower industry profits.

The fundamental cause of endogenous lock-in is increasing returns to scale at the level of the industry, combined with individual firms maximizing only their own (rather than the industry's) payoff. A gap between individual and industry-wide incentives can occur in many strategic interactions. When firms are deciding whether to undercut a rival's price, for instance, they typically don't consider how that decision might affect the rival's profits – and they might not fully realize how it could also affect the market as a whole. In the case of pricing, of course, companies are generally prohibited via antitrust law from colluding. Antitrust law does not, however, generally prohibit the coordination of research activity at the industry level, because more efficient R&D can benefit both consumers and firms.

Consider, for instance, a patent pool, in which firms in an industry agree to freely license patents to one another on the condition that each participant pursues a certain level of R&D, or a research joint venture where members jointly decide how much to spend on R&D and where to allocate that spending. Because firms in this patent pool or research joint venture share technology, the incentive for everyone is to choose a research portfolio that maximizes industry profit. In the late 1980s, for example, the semiconductor research consortium SEMATECH[18] enabled 14 U.S. semiconductor companies to eliminate duplicate research spending and reduced the incentive for any individual firm to pursue research along a technological path that did not maximize industry profits.

2) Ensure outside firms have incentives to build complements to the optimal research line.

Firms should also consider inducing the development of products that are complementary to the research line that maximizes industry profits. These complements, if invented, increase the profitability for every firm in the industry to pursue the research line that maximizes industry profits as a whole. The strategy of subsidizing complements, particularly in the early stages of dynamic increasing returns, is well-known from the literature on platforms.[19] An important insight from this literature is that, if complements are costly to produce, then firms might consider making a credible commitment not to "step on the toes" of a complement's producer by later entering its market. Consider, for example, an online retailer like Amazon with a platform that connects buyers and sellers via a reseller marketplace. That platform benefits from the existence of third-party sellers of unusual products that Amazon, itself, does not carry. These sellers, however, would be reluctant to pay the cost of sourcing unusual products and setting up an Amazon sales page unless they were sure that Amazon wouldn't undercut their business should it prove successful.

Similarly, with respect to R&D, the producers of inventions that are complementary to the research line that maximizes industry profits would be hesitant to develop those complements if they expected firms in the industry to later squeeze their margins or attempt to "invent around" their complement. In many cases, making a credible commitment to avoid such squeezes can be a huge challenge, particularly when the precise form of the complement cannot be known before it is invented. In general, industries with very few large players on the research side and with enough stability for informal agreements ("relational" contracts) to be sustained are generally better positioned to attract the complements that will help ensure that inferior technological lines are not worth pursuing.

3) Link internal incentives to the direct goal.

Within organizations, the actual research is performed by individuals or teams, not by the firms themselves. This means that distortions that induce harmful path dependence can exist *within* a firm for precisely the same reason that they can exist within an industry, even when the firm is a monopolist. Consider, for example, a firm that pays its scientists a bonus that's proportional to the profits generated by any inventions. Moreover, assume that the firm allows its individual scientists to choose what projects to work on. Those conditions are precisely analogous to the incentive system at the industry level that was discussed earlier. Namely, individual inventors will not sufficiently consider how their research might affect the incentives of other inventors.

As it turns out, *every* innovation policy either conditions payments on the value of potential future inventions, or it leaves the firm vulnerable to harmful path dependence.[20] That is to say, if senior executives truly do not know which inventions are the best ones to pursue, then there is no "neutral" salary and bonus policy that will avoid lock-in. For instance, giving bonuses to scientists for successful new products may induce them to work on too many low-value incremental inventions. On the other hand, giving bonuses to scientists only for difficult advances might encourage them to avoid incremental steps that could be necessary to develop the most profitable new technology. To walk the fine line between these two cases, a firm might consider rewarding only those breakthroughs that actually lead to the most profitable new technology, where the measure of "profitable" must adjust for the difficulty or cost of developing that technology. Ideally, a firm would initially query outside experts or perform internal data collection among the research staff to identify the optimal technological path. Then, effort along that path can be encouraged using simple bonuses for the "right" inventions.

4) Beware the difficulty of "leapfrogging" bad investments.

Minor mistakes in the early development of a new technology can ultimately be very costly under conditions of dynamic increasing returns. Every invention can affect not only current profits but future profit potential as well, depending on customer adoption, the steepness of the learning curve, and other factors. In the nuclear power industry, once dozens of light-water plants had been ordered, the technical staff had been trained, and safety regulations had been promulgated, the benefit of switching research efforts back to even a superior reactor type was very limited and hence did not occur.

In the literature on industrial organization,[21] the concept of "leapfrogging" helps explain some of the dynamics involved. Firms or nations with an existing product of technology (say, a developed landline network) may be slower to introduce a new product or adopt a new technology (say, cellular phones) than a firm or nation that doesn't possess the existing technology. The problem for R&D is that even experimental research may generate breakthroughs that can, because of increasing returns to scale, unintentionally place an industry on a harmful technological path. Thus, there's a danger in pursuing a research agenda that attempts to develop many technologies simultaneously. Breakthroughs in an area that's not along the profit-maximizing research path may, via returns to scale in early consumer adoption or rival follow-up inventions, lead the inventing firm to earn less profit that it would have earned had it not developed that invention at all. The desire to pursue an open research agenda – to "not put all of your eggs in one basket" – can actually be harmful!

5) Consider the distortionary effects of "free" money.

In general, government subsidies – "free" money – may seem unambiguously good. If those subsidies are aimed at a technology that leads to harmful lock-in, however, the industry would clearly be better off taking a pass. More subtly, even "technologically neutral" subsidies can induce harmful lock-in. Consider the impact of a hypothetical $10 billion prize offered by the U.S. government in 1950 for the first private-sector firm to develop a nuclear reactor attached to a power grid. Although this prize is "neutral" with respect to technology, it will shift incentives in the private-sector equilibrium toward the development of reactors that can be quickly constructed at a minimum viable scale, rather than toward reactors that are easiest to develop into commercially viable products. Once that minimally viable reactor has been developed, the difficulty of "leapfrogging" might then prevent the industry from ever returning to pursue what initially might have been the most promising technology. Simply put, neutral incentives do not necessarily induce neutral outcomes.

The same is true even when it comes to credit for scientific breakthroughs. Take, for instance, the development of the early airplane.[22] Worldwide fame was on the line for the first person to successfully fly a controllable, heavier-than-air, powered plane, and teams from as far as Australia and Brazil worked on developing a solution. The Wright brothers "invented" the airplane in 1903, garnering worldwide adulation, but the particular nature of their design – wing-warping rather than ailerons for lateral control, and fabric rather than metal for the fuselage – meant that the Wright plane was substantially more difficult to develop into a viable commercial product than the planes that were being developed contemporaneously in Europe. In other words, the desire for scientific credit pushed the Wrights to develop an initially simpler technology that was challenging to commercialize, and hence their victory may have set back the development of the commercial airplane rather than encouraged it. Consequently, firms must be aware of how "free" and "neutral" incentives like government subsidies and scientific credit can, in fact, push an industry onto an inefficient research path.

6) R&D is not the only investment with strategic externalities.

Although this chapter has been concerned largely with endogenous path dependence in R&D, precisely the same type of analysis will potentially apply in any situation in which 1) competitive firms 2) take actions with increasing returns to scale, and 3) those returns to scale are at the industry level and hence each firm's actions generate strategic externalities on other companies. Situations that naturally fit those three criteria include category demand inducement via marketing spending, the sequential release of products with network externalities, and investments in networked supply chains. In each of these cases, among many others, non-coordination at the industry level can lead to lower industry profits in the long run, and these lower profits can persist even when it's known by all firms that the path pursued was suboptimal.

[double line break]

Would the nuclear industry have prospered had the majority of R&D and early learning-by-doing been directed toward reactors other the light water? That's a question best left to historians of technology. But the perilous path dependence that waylaid the industry – a path dependence that was foreseeable in the rational actions of competing firms – could have been avoided with savvy strategic thinking. Moreover, the lessons learned from that market collapse do not pertain only to the nuclear energy industry; in fact, they are applicable to all companies in markets that are governed by similar dynamics.

ENDNOTES

[1] R. Cowan, "Nuclear Power Reactors: A Study in Technological Lock-in," The Journal of Economic History 50.3 (1990): 541-567.

[2] M. Hertsgaard, "Nuclear Inc.: The Men and Money Behind Nuclear Power" (New York: Pantheon Books, 1983).

[3] R.G. Hewlett and J.M. Holl, "Atoms for Peace and War: 1953-1961" (Oakland, CA: University of California Press, 1989).

[4] W. Allen, "Nuclear Reactors for Generating Electricity: U.S. Development from 1946 to 1963," technology report, RAND (1977).

[5] W. Allen, "Nuclear Reactors for Generating Electricity: U.S. Development from 1946 to 1963," technology report, RAND (1977).

[6] W.J. Nuttall, "Nuclear Renaissance: Technologies and Policies for the Future of Nuclear Power" (Boca Raton, FL: CRC Press, 2004).

[7] R. Cowan, "Nuclear Power Reactors: A Study in Technological Lock-in," The Journal of Economic History 50.3 (1990): 541-567.

[8] R. Cowan, "Nuclear Power Reactors: A Study in Technological Lock-in," The Journal of Economic History 50.3 (1990): 541-567.

[9] W.B. Arthur, "Competing Technologies, Increasing Returns, and Lock-in by Historical Events," The Economic Journal 99 (1989): 116-131.

[10] J. Mokyr, "The Lever of Riches" (Oxford, UK: Oxford University Press, 1992).

[11] C.C. McLaughlin, "The Stanley Steamer: A Study in Unsuccessful Invention," Explorations in Entrepreneurial History 7, no. 1 (1954).

[12] W.B. Arthur, "Competing Technologies, Increasing Returns, and Lock-in by Historical Events," The Economic Journal 99 (1989): 116-131.

[13] R. Cowan, "Nuclear Power Reactors: A Study in Technological Lock-in," The Journal of Economic History 50.3 (1990): 541-567.

[14] R. Cowan and S. Hulten, "Escaping Lock-in: The Case of the Electric Vehicle," Technology and the Environment 53, no. 1 (1996).

[15] J. Farrell and G. Saloner, "Installed Base and Compatibility: Innovation, Product Preannouncements, and Predation," American Economic Review 76, no. 5 (1986).

[16] The experience of chemical pesticides *contra* alternative forms of pest management is covered in R. Cowan and P. Gunby, "Sprayed to Death: Path Dependence, Lock-in, and Pest Control Strategies," The Economic Journal 106 (1996): 521-542. A theoretical discussion is contained in M. Rothschild, "A Two-Armed Bandit Theory of Market Pricing," Journal of Economic Theory 9, no. 2 (1974).

[17] See, for example, W.B. Arthur, "Competing Technologies, Increasing Returns, and Lock-in by Historical Events," The Economic Journal 99 (1989): 116-131.

[18] D.A. Irwin and P.J. Klenow, “High-tech R&D Subsidies: Estimating the Effects of Sematech,” Journal of Industrial Economics 40.3-4 (1996): 323-334.

[19] See, for example, A.M. Brandenburger and B. Nalebuff, “Co-opetition” (New York: Crown, 1996); and A. Gawer and R. Henderson, “Platform Owner Entry and Innovation in Complementary Markets: Evidence from Intel,” Journal of Economics and Management Strategy 16, no. 1 (2007).

[20] K. Bryan and J. Lemus, “The Direction of Innovation,” working paper (2016).

[21] D. Fudenberg, R. Gilbert, and J. Stiglitz, “Preemption, Leapfrogging, and Competition in Patent Races,” European Economic Review 22, no. 1 (1983).

[22] K. Bryan, “Industrial Reversals of Fortune: The Meaning of Invention in the Early Airplane Industry,” working paper (2016).