

Preface

In a certain sense, this book is part of a trilogy, though it differs from most trilogies in that the books can be read in any order. The books in question came into being out of a growing sense on my part that the predicament of our time could not be understood from within the conventional wisdom that created it, and that the most important element of that conventional wisdom—the heart of a secular belief system that shares most of the characteristics of a religion—was faith in progress.

My first explorations of that theme focused on understanding where the ersatz religion of progress came from and how the mismatch between faith in progress and the insistent reality of our society's failure to progress—or, put more forcefully, of the opening stages of its decline—was likely to play out in the thought, imagination, and beliefs of people in the contemporary world. Those explorations eventually gave rise to a book, *After Progress: Religion and Reason at the End of the Industrial Age*.¹ As that first reconnaissance reached clarity, I began two other related projects, both oriented toward figuring out what sorts of responses might be appropriate to the end of the age of progress.

One of those projects used narrative fiction to try to explore the prospects of a society that abandoned the religion of perpetual progress and, instead, allowed itself and its citizens to pick and choose among the technologies and lifestyles already explored by our species. That narrative became a novel, *Retrotopia*.² The other project approached the same question from the more

conventional angle of nonfiction, and the result is the book you are holding in your hands right now.

As discussed later in this book, the idea of an end to progress is freighted with a great many irrational terrors and strange beliefs. It's far from uncommon for people to insist that any future that isn't defined by the endless elaboration of already overelaborate technologies must somehow involve going back to the caves or sinking into medieval squalor or being gobbled up by any of the other hobgoblins of the past with which the religion of progress threatens unbelievers. These reactions have deep emotional roots; for several centuries now, a vast number of people in the industrial world have allowed their sense of meaning, purpose, and value to depend on their assumed role in the grand onward march of progress from the caves to the stars, and letting go of that self-image is a very challenging thing.

That said, it's not as though we ultimately have a choice. On the one hand, the exhaustion of nonrenewable resources and the buildup of pollutants in the atmosphere, the seas, and the soil are already starting to impose a rising spiral of costs on further attempts to make our technologies even more elaborate than they are today. On the other, it's becoming increasingly clear to people in the industrial world that progress does not necessarily mean improvement, and that older and simpler technologies very often do a better job at their tasks than the latest hypercomplex, high-tech equivalent. A growing number of people are thus beginning to turn aside from the products of progress. That these older, simpler technologies are very often less dependent on nonrenewable resources and less damaging to the biosphere that supports all our lives is just one benefit of that heretical but necessary act.

This book seeks to discuss what the world looks like in the wake of the end of progress: why progress is ending, why it could never have fulfilled the overblown promises made in its name, and what the prospects of our society and species might look like as the age of progress gives way to an age of environmental

blowback and technological unraveling. It's popular to paint those latter prospects in unremittingly bleak colors, but here again that reflects the unthinking assumptions of our age rather than the facts as they actually exist. The burdens that progress have piled upon us, as individuals, as communities, and as a species, are not small, and once the shock has passed off, liberation from those burdens may well be experienced by many of us as a reason for celebration rather than mourning.

That said, there are serious downsides to the end of progress, just as there were equally serious downsides to its beginning and to every step of its historical course. My hope is that this book, as a first survey of the almost entirely unexplored landscape on the far side of progress, will help my readers prepare themselves for the largely unexpected future ahead of us.

My previous books have had a variety of intellectual debts, but this one has depended almost entirely on one source—the readership of my former blog “The Archdruid Report.” For eleven years, from the first tentative posts about peak oil and the future of industrial society all the way to the last posts about the nature of human experience, my readers encouraged me, argued with me, brought me data points that confirmed or challenged the ideas that I’ve offered, and in general created a congenial and thought-provoking environment for the development of my ideas. My thanks go to all.

1

THE END OF PROGRESS



MOST PEOPLE in the industrial world believe that the future is, by definition, supposed to be better than the past, that growth is normal and contraction is not, that newer technologies are superior to older ones, and that the replacement of simple technologies by complex ones is as unstoppable as it is beneficent. That's the bedrock of the contemporary faith in progress.¹ This faith remains unchallenged by most people today, even though the evidence of our everyday lives contradicts it at every turn.

Most of us know perfectly well that every software “upgrade” these days has more bugs and fewer useful features than what it replaced, and every round of “new and improved” products hawked by the media and shoveled onto store shelves is more shoddily made, more loaded with unwanted side effects, and less satisfactory at meeting human needs than the last round. Somehow, though, a good many of the people who witness this reality, day in and day out, still manage to insist that the future will be, or at least ought to be, a paradise propped up by perfectly functioning machines, in which all the latest clichés about the future will inevitably come true. That the rising tide of technological failure might be something other than an accidental roadbump

on the way to utopia—that it might be trying to tell us something that, by and large, we don't want to hear—has not yet entered our society's darkest dream.

Meanwhile, as problems mount and solutions run short, the contemporary faith in progress drives a common insistence that it's never too late to save the world. No matter how troubling the signs on the horizon, no matter how many predictions of impending trouble have turned into descriptions of troubles we're facing here and now, it's astonishingly rare for anyone to notice that we're past the point where it makes any sense to sit around talking about how somebody ought to fix things one of these days.

The events of our time, though, show no particular interest in waiting until we get around to dealing with them. At least three factors at work in today's world—peak oil, and more generally the peaking of global production of fossil fuels; the ongoing failure of alternative energy technologies to replace fossil fuels; and the accelerating pace of anthropogenic climate change—are already having a major impact on the global economy and, increasingly, on other aspects of human and nonhuman life as well.

Those issues could have been faced and dealt with as soon as it became clear that they were going to be problematic. In every case, there were straightforward fixes available, and if they had been put into place as soon as the facts showed that trouble was on its way, the necessary changes could have been made gradually, without overturning the whole structure of society. But that's not what happened. Instead, obsolete policies stayed frozen in place while the opportunities for constructive change slipped past. Now the bill is coming due.

This doesn't mean that action is useless, much less that we should huddle down, close our eyes, and wait for the end. It means, rather, that business as usual will not last much longer, no matter what we do about it. In the decades ahead, many things that people in the industrial world consider normal will go away

forever. That's going to be profoundly difficult, but it's also profoundly liberating, because the struggle to maintain the status quo has been a massive force blocking the way to constructive change. As the familiar landscapes of the industrial age give way to the unexpected vistas of the near and middle future, the focus of meaningful action will have to shift from preservation to remediation, from "How can we keep our familiar ways of doing things?" to "Now that the familiar things are gone, what can we put in their place?"

The Law of Diminishing Returns

I'm well aware that asking people in the early twenty-first century to doubt the omnipotence and eternal goodness of progress ranks right up there with suggesting to a medieval peasant that God and his saints and angels aren't up there in heaven any more. There are nonetheless two crucial reasons why cumulative technological progress, of the sort that's reshaped the industrial world over the past three centuries, was a temporary, self-limiting process that often imposed costs that outweighed its benefits.

The first is the law of diminishing returns—the principle that the more often you repeat a given action, the fewer benefits you get from each successive repetition and the more the costs mount up. Nearly everything in the world of human experience is subject to this law. The process of extracting petroleum from the earth is a good example: the first oil wells made huge profits for very little expense, while the hunt for the last scrapings of the bottom of the planet's oil barrel that occupies the petroleum industry today has extraordinarily high costs and meager returns.

Is technological progress subject to the same principle? Believers in progress like to insist that this can't be the case, but the evidence suggests otherwise. Consider the way that energy technologies have become more and more expensive to develop over time. The steam engine, the first major energy technology

innovation in modern times, was invented by working engineers in their off hours, using ordinary pipefitting tools. The internal combustion engine and the electrical generator required more systematic effort, but were still well within the reach of a single inventor working in a laboratory. Nuclear fission required an expenditure of money and resources so huge that only a handful of relatively rich nations could afford it. Commercial nuclear fusion power, as we'll see in Chapter Two, is turning out to be so costly that nobody anywhere can afford it at all.

In exactly the same way, and for many of the same reasons, the first advances in health care—basic sanitation, antiseptics, and vaccination—cost very little and brought immense benefits. With every passing year, costs went up and benefits went down, until current health care research is investing billions of dollars in projects that may benefit only a few people, if any. The low-hanging fruit got picked first, leaving more difficult projects for later.

The same is true in every other field. This is why, as a detailed study of patent records has shown, the modern world hit its peak of innovation in the last quarter of the nineteenth century, and the pace of technological progress has actually decreased steadily since that time.² This implies that, at some point, the benefits of continued technological progress will no longer equal the costs and progress will grind to a halt because it no longer pays for itself. Thus there can be such a thing as too much technology, and a very strong case can be made that in the world's industrial nations we've already gotten well past that point.

In a typically cogent article,³ maverick economist Herman Daly sorted out the law of diminishing returns into three interacting processes. The first is *diminishing marginal utility*—that is, the more of anything you have, the less any additional increment of that thing contributes to your well-being. If you're hungry, one sandwich is a very good thing; two is pleasant; three is a luxury; and somewhere beyond that, when you've given sandwiches to all your coworkers, the local street people, and anyone else you can

find, more sandwiches stop being any use to you. When more of anything no longer brings any additional benefit, you've reached the point of futility, at which further increments are a waste of time and resources.

Well before that happens, though, two other factors come into play. First, it costs you almost nothing to cope with one sandwich, and very little more to cope with two or three. After that you start having to invest time, and quite possibly resources, in dealing with all those sandwiches, and each additional sandwich adds to the total burden. Economists call that *increasing marginal disutility*—that is, the more of anything you have, the more any additional increment of that thing is going to cost you, in one way or another. Somewhere in there, too, there's the impact that dealing with those sandwiches has on your ability to deal with other things you need to do; that's the *increasing risk of whole-system disruption*—the more of anything you have, the more likely it is that an additional increment of that thing is going to disrupt the wider system in which you exist.

Next to nobody wants to talk about the way that technological progress has already passed the point of diminishing returns in all three senses: that the marginal utility of each new round of technology is dropping fast; the marginal disutility is rising at least as fast; and whole-system disruptions driven by technology are becoming an inescapable presence in everyday life. Still, an uncomfortable awareness of that fact is becoming increasingly common these days, however subliminal it may be, and is beginning to have a popular culture among many other things.

If you've dug yourself into a hole, as the saying goes, the first thing you need to do is stop digging. If a large and growing fraction of your society's problems are being caused by too much technology applied with too little caution, similarly, it's not exactly helpful to insist that applying even more technology with even less skepticism about its consequences is the only possible answer to those problems.

There's a useful word for something that remains stuck in a culture after the conditions that once made it relevant have passed away, and that word is "superstition." The beliefs that more technology is always better, that every problem must have a technological solution, and that technology always solves more problems than it creates, are among the prevailing superstitions of our time. I'd like to suggest that, comforting and soothing as those superstitions may be, it's high time we outgrow them and deal with the hard realities of a world in which taking such faith-based notions as a guide to the future may not be sensible, or even sane.

Yet there's another reason to ask hard questions about where progress is taking us, and it unfolds from the issue of externalities. Externalities are the costs of an economic activity that aren't paid by the buyer or the seller directly but are pushed off onto some third party. You won't hear much about externalities these days; in many circles, it's considered impolite to mention them, but they're a pervasive presence in contemporary life, and they play a very large role in some of the most intractable problems of our age. Some of those problems were discussed by Garrett Hardin in his famous essay on the tragedy of the commons, and more recently by Elinor Ostrom in her studies of how that tragedy can be avoided.⁴ Still, I'm not sure how often it's recognized that the phenomena they discussed applies not just to commons but to societies as a whole—especially to societies like ours.

A simplified example may be useful here. Let's imagine a blivet factory, then, that turns out three-pronged blivets in pallet loads for customers. The blivet-making process, like all other manufacturing, produces waste as well as blivets, and we'll assume for the sake of the example that blivet waste is moderately toxic and causes health problems in people who ingest it. The blivet factory produces one barrel of blivet waste for every pallet of blivets it ships. The cheapest option for dealing with the waste, and thus the option that economists favor, is to dump it into the river that flows past the factory.

Notice what happens if the blivet manufacturer follows this approach. The manufacturer has maximized his own benefit from the manufacturing process by avoiding the expense of finding some other way to deal with all those barrels of blivet waste. His customers also benefit, because blivets cost less than they would if the cost of waste disposal was factored into the price. On the other hand, the costs of dealing with the blivet waste don't disappear; they are imposed on the people downstream who get their drinking water directly or indirectly from the river and who suffer from health problems because there's blivet waste in their water. The blivet manufacturer is thus externalizing the cost of waste disposal; his increased profits are being paid for at a remove by the increased health care costs of everyone downstream.

That's how externalities function. Back in the days when people actually talked about the downsides of economic growth, there was a lot of discussion of how to handle externalities, and not just on the leftward end of the spectrum. I recall a thoughtful book titled *TANSTAAFL*—that's an acronym, for those who don't know their Heinlein, for "There Ain't No Such Thing As A Free Lunch"⁵—which argued, on solid libertarian-conservative grounds, that the environment could best be preserved by making sure that everyone paid full sticker price for the externalities they generated. Today's crop of American pseudo-conservatives, of course, turned their back on all this a long time ago and insist at the top of their lungs on their allegedly God-given right to externalize as many costs as they possibly can. This is all the more ironic in that most pseudo-conservatives claim to worship a God who said some very specific things about "what ye do unto the least of these," but that's a subject for a different time.

The Externality Trap

Economic life in the industrial world these days can be described, without too much inaccuracy, as an arrangement set up to allow a privileged minority to externalize nearly all their costs onto the rest of society while pocketing as many as possible of the benefits

themselves. That's come in for a certain amount of discussion in recent years,⁶ but I'm not sure how many of the people who've participated in those discussions have given any thought to the role that technological progress plays in facilitating the internalization of benefits and the externalization of costs that drive today's increasingly inegalitarian societies. Here again, our example will be helpful.

Before the invention of blivet-making machinery, let's say, blivets were made by old-fashioned blivet makers, who hammered them out on iron blivet anvils in shops that were to be found in every town and village. Like other handicrafts, blivet-making was a living rather than a ticket to wealth; blivet makers invested their own time and muscular effort in their craft and turned out enough in the way of blivets to meet the demand. Notice also the effect on the production of blivet waste. Since blivets were being made one at a time rather than in pallet loads, the total amount of waste was smaller; the conditions of handicraft production also meant that blivet makers and their families were more likely to be exposed to the blivet waste than anyone else, and so they had an incentive to invest the extra effort and expense to dispose of it properly. Since blivet makers were ordinary craftspeople rather than millionaires, furthermore, they weren't as able to buy exemption from local health laws.

The invention of the mechanical blivet press changed that picture completely. Since one blivet press could do as much work as fifty blivet makers, the income that would have gone to those fifty blivet makers and their families went instead to one factory owner and his stockholders, with as small a share as possible set aside for the wage laborers who operated the blivet press. The factory owner and stockholders had no incentive to pay for the proper disposal of the blivet waste, either—quite the contrary, since having to meet the disposal costs cut into their profit, buying off local governments was much cheaper, and, if the harmful effects of blivet waste were known, you can bet that the owner and shareholders all lived well upstream from the factory.

Notice also that a blivet manufacturer who paid a living wage to his workers and covered the costs of proper waste disposal would have to charge a higher price for blivets than one who did neither and thus would be driven out of business by his more ruthless competitor. Externalities aren't simply made possible by technological progress, in other words; they're the inevitable result of technological progress in a market economy, because the more a firm externalizes the costs of production, the more readily it can outcompete rival firms, and the firm that succeeds in externalizing the largest share of its costs is the most likely to survive and prosper.

Each further step in the progress of blivet manufacturing, in turn, tightened the same screw another turn. Today, to finish up the metaphor, the entire global supply of blivets is made in a dozen factories in distant Slobbovia, where sweatshop labor under ghastly working conditions and the utter absence of environmental regulations make the business of blivet fabrication more profitable than anywhere else. The blivets are as shoddily made as possible; the entire blivet supply chain, from the open-pit mines worked by slave labor that provide the raw materials to the big box stores with part-time, poorly paid staff selling blivetronic technology to the masses, is a human and environmental disaster. Every possible cost has been externalized, so that the two multinational corporations that dominate the global blivet industry can maintain their profit margins and pay absurdly high salaries to their CEOs.

That in itself is bad enough, but let's broaden the focus to include the whole systems in which blivet fabrication takes place: the economy as a whole, society as a whole, and the biosphere as a whole. The impact of technology on blivet fabrication in a market economy has predictable and well-understood consequences for each of these whole systems, which can be summed up precisely in the language we've already used. In order to maximize its own profitability and return on shareholder investment, the blivet industry externalizes costs in every available direction.

Since nobody else wants to bear those costs, either, most of them end up being passed onto the whole systems just named, because the economy, society, and the biosphere have no voice in today's economic decisions.

Like the costs of dealing with blivet waste, though, the other externalized costs of blivet manufacture don't go away just because they're externalized. As externalities increase, they tend to degrade the whole systems onto which they're dumped—the economy, society, and the biosphere. This is where the trap closes tight, because blivet manufacturing exists within those whole systems and can't be carried out unless all three systems are sufficiently intact to function in their usual way. As those systems degrade, their ability to function degrades also, and eventually one or more of them breaks down—the economy plunges into a depression; the society disintegrates into anarchy or totalitarianism; the biosphere shifts abruptly into a new mode that lacks adequate rainfall for crops—and the manufacture of blivets stops because the whole system that once supported it has stopped doing so.

Notice how this works out from the perspective of someone who's benefiting from the externalization of costs by the blivet industry—the executives and stockholders in a blivet corporation, let's say. As far as they're concerned, until very late in the process, everything is fine and dandy: each new round of technological improvements in blivet fabrication increases their profits, and if each such step in the onward march of progress also means that economies go haywire, democratic institutions implode, toxic waste builds up in the food chain, or what have you, hey, that's not their problem—and after all, that's just the normal, praiseworthy creative destruction of capitalism, right?

That sort of insouciance is easy for at least three reasons. First, the impacts of externalities on whole systems can pop up a very long way from the blivet factories. Second, in a market economy, everyone else is externalizing their costs as enthusiastically as the blivet industry, and so it's easy for blivet manufacturers

(and everyone else) to insist that whatever's going wrong is not their fault. Third, and most crucially, whole systems as stable and enduring as economies, societies, and biospheres can absorb a lot of damage before they tip over into instability. The process of externalization of costs can thus run for a very long time, and become entrenched as a basic economic habit, long before it becomes clear to anyone that continuing along the same route is a recipe for disaster.

Even when externalized costs have begun to take a visible toll on the economy, society, and the biosphere, furthermore, any attempt to reverse course faces nearly insurmountable obstacles. Those who profit from the existing order of things can be counted on to fight tooth and nail for the right to keep externalizing their costs: after all, they have to pay the full price for any reduction in their ability to externalize costs, while the benefits created by not imposing those costs on whole systems are shared among all participants in the economy, society, and the biosphere respectively. Nor is it necessarily easy to trace back the causes of any given whole-system disruption to specific externalities benefiting specific people or industries. It's rather like loading hanging weights onto a chain; sooner or later, as the amount of weight hung on the chain goes up, the chain is going to break, but the link that breaks may be far from the last weight that pushed things over the edge, and every other weight on the chain made its own contribution to the end result.

A society that's approaching collapse because too many externalized costs have been loaded onto the whole systems that support it thus shows certain highly distinctive symptoms. Things are going wrong with the economy, society, and the biosphere, but nobody seems to be able to figure out why; the measurements that economists use to determine prosperity show contradictory results, with those that measure the profitability of individual corporations and industries giving much better readings than those that measure the performance of whole systems; the rich

are convinced that everything is fine, while outside the narrowing circles of wealth and privilege, people talk in low voices about the rising spiral of problems that beset them from every side. If this doesn't sound familiar to you, dear reader, you probably need to get out more.

At this point it may be helpful to sum up the argument I've developed here:

- a) Every increase in technological complexity tends also to increase the opportunities for externalizing the costs of economic activity.
- b) Market forces make the externalization of costs mandatory rather than optional, since economic actors that fail to externalize costs will tend to be outcompeted by those that do.
- c) In a market economy, as all economic actors attempt to externalize as many costs as possible, externalized costs will tend to be passed on preferentially and progressively to whole systems such as the economy, society, and the biosphere, which provide necessary support for economic activity but have no voice in economic decisions.
- d) Given unlimited increases in technological complexity, there is no necessary limit to the loading of externalized costs onto whole systems, short of systemic collapse.
- e) Unlimited increases in technological complexity in a market economy thus necessarily lead to the progressive degradation of the whole systems that support economic activity.
- f) Technological progress in a market economy is therefore self-terminating, and ends in collapse.

Secular Stagnation

There are, of course, arguments that could be deployed against this thesis. For example, it could be argued that progress doesn't have to generate a rising tide of externalities. The difficulty with this argument is that externalization of costs isn't an accidental side effect of technology but an essential aspect—it's not a bug,

it's a feature. Every technology is a means of externalizing some cost that would otherwise be borne by a human body. Even something as simple as a hammer takes the wear and tear that would otherwise affect the heel of your hand, let's say, and transfers it to something else: directly, to the hammer; indirectly, to the biosphere by way of the trees that had to be cut down to make the charcoal to smelt the iron, the plants that were shoved aside to get the ore, and so on.

The more complex a technology becomes, the more costs it generates, since every bit of added complexity has to be paid for. Each more-complex technology thus has to externalize its additional costs in order to compete against the simpler technology it replaces. In the case of a hypercomplex technosystem such as the internet, the process of externalizing costs has gone so far, through so many tangled interrelationships, that it's next to impossible to figure out exactly who's paying for how much of the gargantuan inputs needed to keep the thing running. This lack of transparency feeds the illusion that large systems are cheaper than small ones, by making externalities of scale look like economies of scale.

It might be argued instead that a sufficiently stringent regulatory environment, forcing economic actors to absorb all the costs of their activities instead of externalizing them onto others, would be able to stop the degradation of whole systems while still allowing technological progress to continue. The difficulty here is that increased externalization of costs is what makes progress profitable. All other things being equal, a complex technology will be more expensive in real terms than a simpler technology, for the simple fact that each additional increment of complexity has to be paid for by an investment of energy and other forms of real capital.

Strip complex technologies of the subsidies that transfer some of their costs to the government, the perverse regulations that transfer some of their costs to the rest of the economy, the

bad habits of environmental abuse and neglect that transfer some of their costs to the biosphere, and so on, and pretty soon you're looking at hard economic limits to technological complexity, as people forced to pay the full sticker price for complex technologies maximize the benefits they receive by choosing simpler, more affordable options instead. A regulatory environment sufficiently strict to keep technology from accelerating to collapse would thus bring technological progress to a halt by making it unprofitable.

Notice, however, the flipside of the same argument: a society that chose to stop progressing technologically could maintain itself indefinitely, so long as its technologies weren't dependent on nonrenewable resources or the like. The costs imposed by a stable technology on the economy, society, and the biosphere would be more or less stable, rather than increasing over time, and it would therefore be much easier to figure out how to balance out the negative effects of those externalities and maintain the whole system in a steady state. Societies that treated technological progress as an option rather than a requirement, and recognized the downsides to increasing complexity, could also choose to reduce complexity in one area in order to increase it in another, and so on—or they could just raise a monument to the age of progress and go do something else instead.

The costs of progress are already starting to take an increasing bite out of the global economy. That bite shows up in the fact that none of the twenty biggest industries in today's world could break even, much less make a profit, if they had to pay for the damage they do to the environment.⁷ The conventional wisdom these days is that it's unfair to make those industries pay for the costs they impose on the rest of us. That attitude is exemplified by fracking firms in North Dakota, among many others, who proposed at height of the fracking bubble that they should be exempted from rules for handling radioactive waste from their wells, because following the rules would prevent them from making a profit.⁸ That the costs externalized by the fracking industry

will sooner or later be paid by others, as radionuclides in fracking waste begin to generate cancer clusters, is not something that our current economic thought is able to grasp.

The crucial point to keep in mind is that externalized costs don't just go away. They will be paid by someone; the only question is who pays them. That's the central argument of *The Limits to Growth*, still the most accurate (and thus inevitably the most reviled) of the studies that tried unavailingly to guide industrial society away from self-inflicted ruin: on a finite planet, once an inflection point is passed, the costs of economic growth rise faster than growth does, and sooner or later force the global economy to its knees.⁹ The mere fact that those costs aren't carried on the balance sheets of the companies that generate them doesn't make those costs go away; it just keeps them from being taken into account by policymakers.

One way in which those costs may already be having an impact is in the phenomenon of secular stagnation.¹⁰ It so happens that when you subtract the paper wealth manufactured by derivatives and other forms of financial make-believe, the global economy has been stuck in a period of slow, no, or negative growth since 2009. There are plenty of economists, mind you, who insist that this can't happen, and even among those who admit that what's happening can indeed happen, there's no consensus as to how or why such a thing could occur. I'd like to suggest that the most important cause of secular stagnation is the increasing impact of externalities on the economy. The dubious book-keeping that leads economists to think that externalized costs go away because they're not entered into anyone's books doesn't actually make them disappear, after all. Instead, they become an unrecognized burden on the economy as a whole, an unfelt headwind blowing with hurricane force in the face of economic growth.

Thus the insistence by fracking firms that they ought to be allowed to externalize even more of their costs in order to maintain

their profit margin is self-defeating, even if the firms themselves don't realize that. If in fact the buildup of externalized costs is what's causing the ongoing economic slowdown that's driving down commodity prices, forcing down interest rates, and resurrecting the specter of deflationary depression, the fracking firms in question are trying to respond to secular stagnation by doing more of what causes secular stagnation.

In theory, this sort of self-defeating behavior would be recognized for what it is and set aside as a bad idea. In the real world, by contrast, fracking firms, like every other business concern these days, can be expected to put their short-term cash flow ahead of the survival of their industry, or for that matter of industrial civilization as a whole. That's business as usual—and it's made even easier than it otherwise would be by certain habits of thought that make it hard to think clearly about technology and progress.

Technologies in the Plural

Here's an example. When talking heads these days babble about technology in the singular, as a uniform, monolithic thing that progresses according to some relentless internal logic of its own, they're spouting balderdash.¹¹ In the real world, there's no such monolith, no technology in the singular. Instead, there are technologies in the plural, clustered more or less loosely in technological suites that may or may not have any direct relation to one another.

An example, again, might be useful here. Consider the technologies necessary to build a steel-framed bicycle. The metal parts require the particular suite of technologies we use to smelt ores, combine the resulting metals into useful alloys, and machine and weld those into shapes that fit together to make a bicycle. The tires, inner tubes, brake pads, seat cushion, handlebar grips, and paint require a different suite of technologies drawing on various branches of applied organic chemistry, and a few other suites also

have a place: for example, the one that's needed to make and apply lubricants.

The suites that make a bicycle have other uses. If you can build a bicycle, as Orville and Wilbur Wright demonstrated, you can also build an aircraft and a variety of other interesting machines as well. That said, there are other technologies—say, the ones needed to manufacture medicines or precision optics or electronics—that require very different technological suites. You can have everything you need to build a bicycle and still be unable to make a telescope or a radio receiver, and vice versa.

Strictly speaking, therefore, nothing requires all the different technological suites to move in lockstep. It would have been quite possible for different technological suites to have appeared in a different order than they did; it would have been just as possible for some of the suites central to our technologies today to have never gotten off the ground, while other technologies we never tried emerged instead. Imagine, for example, an alternative reality in which solar water heaters (in our world, worked out by 1920) and passive solar architecture (mostly developed in the 1960s and 1970s) were standard household features, canal boats (dating from before 1800) and tall ships (ditto) were the primary means of bulk transport, shortwave radio (developed in the early twentieth century) was the standard long-range communications medium, ultralight aircraft (largely developed in the 1980s) were still in use, and engineers crunched numbers using slide rules (perfected around 1880).

There's no reason why such a pastiche of technologies from different eras couldn't work. We know this because what passes for modern technology is a pastiche of the same kind, in which (for example) cars whose basic design dates from the 1890s are gussied up with computers invented a century later. Much of modern technology, in fact, is old technology with a new coat of paint and a few electronic gimmicks tacked on, and it's old technology that originated in many different eras, too. Part of what

differentiates modern technology from older equivalents, in other words, is mere fashion. Another part, though, moves into more explosive territory.

One reader of my blog “The Archdruid Report” once enlivened the discussion on the comments page with the story of the one and only class on advertising she took at college. The teacher invited a well-known advertising executive to come and talk about the business, and one of the points he brought up was the marketing of disposable razors. The old-fashioned steel safety razor, the guy admitted cheerfully, was a much better product: it was more durable, less expensive, and gave a better shave than disposable razors. Unfortunately, it didn’t make the kind of profits for the razor industry that the latter wanted, and so the job of the advertising company was to convince shavers that they really wanted to spend more money on a worse product instead.

I know it may startle some people to hear a luxuriantly bearded man talk about shaving, but I do have a certain amount of experience with the process—though admittedly it’s been many years. The executive was quite correct: an old-fashioned safety razor with interchangeable blades gives better shaves than a disposable. What’s more, an old-fashioned safety razor combined with a shaving brush, a cake of shaving soap, a mug, and a bit of hot water from the teakettle produces a shaving experience that’s vastly better, in every sense, than what you’ll get from squirting chemical-laced foam out of a disposable can and then scraping your face with a disposable razor; it takes the same amount of time, costs much less on a per-shave basis, and has a drastically smaller ecological footprint to boot.

Notice also the difference in the scale and complexity of the technological suites needed to maintain these two ways of shaving. To shave with a safety razor and shaving soap, you need the metallurgical suite that produces razors, the very simple household-chemistry suite that produces soap, the ability to make pottery and brushes, and some way to heat water. To shave

with a disposable razor and a can of squirt-on shaving foam, you need fossil fuels for plastic feedstocks, chemical plants to manufacture the plastic and the foam, the whole range of technologies needed to manufacture and fill the pressurized can, and so on—all so that you can count on getting an inferior shave at a higher price, and the razor industry can boost its quarterly profits.

That's a small, and arguably silly, example of a vast and far from silly issue. These days, when you see the words "new and improved" on a product, rather more often than not, the only thing that's been improved is the bottom line of the company that's trying to sell it to you. When you hear equivalent claims about some technology that's being marketed to society as a whole, rather than sold to you personally, the same rule applies at least as often.

What, after all, defines a change as "progress"? There's a wilderness of ambiguities, some of them quite deliberate, hidden in that apparently simple word. The contemporary faith in progress presupposes that there's an inherent dynamic to history and that things therefore change, or tend to change, or at the very least ought to change, from worse to better over time. That presupposition then gets flipped around into the even more dubious claim that just because something's new, it must be better than whatever it replaced. Move from there to specific examples, and all of a sudden it's necessary to deal with competing claims—if there are two hot new technologies on the market, is option A more progressive than option B, or vice versa? The answer, of course, is that whichever of them manages to elbow the other aside, by whatever means, will be retroactively awarded the coveted title of the next step in the march of progress.

That was exactly the process by which the appropriate tech of the 1970s was shoved aside and buried in the memory hole of our culture.¹² In its heyday, appropriate tech was as cutting edge as anything you care to name, a rapidly advancing field pushed forward by brilliant young engineers and innovative startups, and it saw itself (and presented itself to the world) as the wave of the

future. In the wake of the Reagan-Thatcher counterrevolution of the 1980s, though, it was retroactively stripped of its progressive label and consigned to the dustbin of the past. Technologies that had been lauded in the media as brilliantly innovative in 1978 were thus being condemned in the same media as Luddite throwbacks by 1988. If that abrupt act of redefinition reminds any of my readers of the way history got rewritten in George Orwell's 1984—"Oceania has never been allied with Eurasia" and the like—well, let's just say the parallel was noticed at the time, too.

The same process on a much smaller scale can be traced with equal clarity in the replacement of the safety razor and shaving soap with the disposable razor and squirt-can shaving foam. In what sense is the latter, which wastes more resources and generates more trash in the process of giving users a worse shave at a higher price, an advance (that metaphor again) on the former? Merely the fact that it's been awarded that title by advertising and the media. If razor companies could make more money by reintroducing the Roman habit of scraping beard hairs off the face with a chunk of pumice, no doubt that would quickly be proclaimed as the last word in cutting-edge, up-to-date hyper-modernity, too.

What Progress Actually Means

Behind the mythological image of the relentless and inevitable forward march of technology-in-the-singular in the grand cause of progress, in other words, lies a murky underworld of crass commercial motives and no-holds-barred struggles over which of the available technologies will get the funding and marketing that will define it as the next step in the march of progress. That's as true of major technological programs as it is of shaving supplies. Some of my readers are old enough, as I am, to remember when supersonic airliners and undersea habitats were the next great steps in progress, until all of a sudden they weren't, and we may

not be all that far from the point at which space travel and nuclear power will go the way of Sealab and the Concorde.

In today's industrial societies, we don't talk about that. It's practically taboo these days to mention the long, long list of waves of the future that abruptly stalled and rolled back out to sea without delivering on their promoters' overblown promises. Remind people that the same rhetoric currently being used to prop up faith in space travel, fusion power, or some other grand technological project was lavished just as thickly on these earlier failures, and you can expect to have that comment shouted down as an irrelevancy, if the other people in the conversation don't simply turn their backs and pretend that they never heard you say anything at all.

They have to do something of the sort, because the alternative is to admit that what we call "progress" isn't the impersonal, unstoppable force of nature that industrial culture's ideology insists it must be. Pay attention to the grand technological projects that failed, compare them with those that are failing now, and it's impossible to keep ignoring certain crucial if hugely unpopular points. To begin with, technological progress is a function of collective choices—do we fund Sealab or the Apollo program? Supersonic transports or urban light rail? Energy conservation and appropriate technology or an endless series of wars in the Middle East? No impersonal force makes those decisions; individuals and institutions make them, and then use the rhetoric of impersonal progress to cloak the political and financial agendas that guide the decision-making process.

What's more, even if the industrial world chooses to invest its resources in a project, the laws of physics and economics, not human preferences, determine whether the project is going to work. The Concorde is the poster child here, a white elephant that could never even cover its own operating costs. Like nuclear power, it was a technological success but an economic flop, only viable given huge and continuing government subsidies, and

since the strategic benefits Britain and France got from having Concorde in the air were nothing like so great as those they got from having an independent source of raw material for nuclear weapons, it's not hard to see why the subsidies went where they did.

That is to say, when something is being lauded as the next great step forward in the glorious march of progress leading humanity to a better world someday, those who haven't drunk themselves tipsy on industrial civilization's folk mythology need to keep three things in mind. The first is that the next great step (etc.) might not actually work when it's brought down out of the billowing clouds of overheated rhetoric into the cold hard world of everyday life. The second is that even if it does work, the next great step (etc.) may be an inferior product, and do a less effective job of meeting human needs than whatever it's supposed to replace. The third is that when it comes right down to it, to label something as the next great step (etc.) is just a sales pitch, an overblown and increasingly trite way of saying "You really ought to buy this."

That implies, in turn, that it's entirely reasonable to respond by saying, "No, I like the thing I'm already using better"—or even to utter the unmentionable and say, "No, I'm going to use this other technology from the past because it works better." Get past the thoughtstopping rhetoric that insists you can't turn back the clock—of course you can; most of us do it every autumn without a second thought when daylight saving time ends—and it becomes hard not to notice that "progress" is just a label for whatever choices happen to have been made by governments and corporations, with or without input from the rest of us, and that if we don't like the choices that have been made for us in the name of progress, we can choose something else.

Of course, it's possible to stuff that sort of thinking back into the straitjacket of progress and claim that progress is chugging

along just fine, and all we have to do is get it back on the proper track or what have you. This is a very common sort of argument and one that's been used over and over again by critics of this or that candidate for the next great step (etc.). The problem with that argument is that it may occasionally win battles but it consistently loses the war. By failing to challenge the folk mythology of progress and the unstated agendas that are enshrined by that mythology, it guarantees that no matter what technology or policy or program gets put into place, it'll end up leading to the same place as all the others before it.

That's the trap hardwired into the contemporary faith in progress. Once you buy into the notion that the specific choices made by industrial societies over the past three centuries or so are more than the projects that happened to win out in the struggle for wealth and power, once you let yourself believe that there's a teleology to it all—that is, that there's some objective goal called “progress” that these choices do a better or worse job of furthering—you've just made it much harder to ask the hard but necessary questions about where this thing called “progress” is going. The word “progress,” remember, means going further in the same direction, and it's precisely questions about the direction that industrial society is going that most need to be asked.

I'd like to suggest that going further in the direction we've been going isn't a particularly bright idea just now. Going further in the direction we've been going means trying to expand per capita energy consumption in an era when fossil fuel reserves are depleting fast and the global economy is creaking and shuddering under the burden of increasingly costly fuel extraction. It means dumping ever more waste into the biosphere when the consequences of previous dumping are already bidding fair to threaten the survival of entire nations. On a less global scale, it also means shoddier products with louder advertising in a race to the bottom that's already gone very far.

Look at a trend that affects your life right now, and extrapolate it out in a straight line; that's what going further in the same direction means. If that appeals to you, dear reader, then you're certainly welcome to it. I have to say it doesn't do much for me.

It's only from within the folk mythology of progress, though, that we have no choice but to accept the endless prolongation of current trends. Right now, as individuals, we can choose to shrug and walk away from the latest hypermodern trinkets, and do something else instead.