No Limits
Energy and Technology

Peter W. Huber and Mark P. Mills

Derived from research relating to the Digital Power Report and from Huber-Mills forthcoming book "The Logic of Power: Energy in the Digital Millennium."

Colonel Edwin Drake was prepared to go quite a bit deeper when he turned on his steam engine in Crawford County, Pennsylvania, in 1859, but his drill struck oil at 69 feet. Today's oil companies drill as much as six miles for their crude – a first vertical leg through deep water and rock, and then significant horizontal distances beyond that.

That we now drill miles for oil, rather than feet, is hardly surprising. We pump the easy oil first, so tomorrow's crude is farther away than yesterday's. What's surprising is this: Over the long term, the price of oil holds remarkably steady. The six-mile oil costs less than the 60-feet oil did, and about the same as one-mile oil did a decade ago. There have been price spikes and sags, but tied to political and regulatory instabilities, not discovery and extraction costs. (**Figure** 1)

Figure 1. Run Out of Resources?

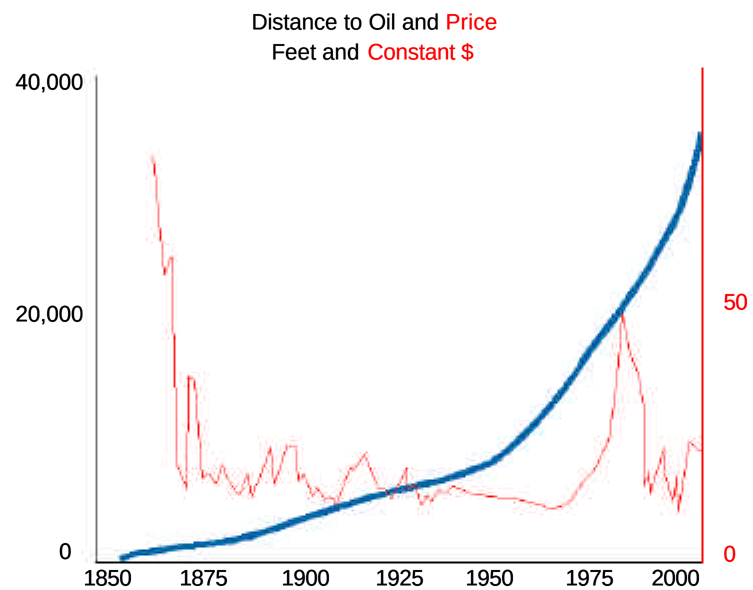

Source: Energy Information Association, ExxonMobil, and J. Ray McDermott Inc.

The amount of oil pumped out of each individual production rig is rising, too, because horizontal drilling brings more oil to the surface through a single spigot. Getting more oil out of fewer rigs costs less, because so much of the expense is in the equipment at the surface, not in the length of the bore hole.

These trends are mirrored in all subsequent stages of energy conversion and consumption, as well. More oil per derrick is followed by more horsepower per pound of car engine, more thrust per pound of jet engine, more power per solid-state switch, more photons per light-emitting diode and laser. At every stage of extraction, conversion, and final use, more energy gets squeezed out of, or packed into, smaller, lighter rigs, systems, engines, power plants, and emitters.

Many technological factors have contributed to these changes, but the most important has been the rise of power semiconductors. The materials and quantum phenomena that brought us digital information are now ushering in a new age of digital power. They extract, process, and use energy in altogether new ways. They pack far more power, into far less space. They control high-power streams of electrons and photons at speeds thousands to millions of times faster than the old technologies they are rapidly supplanting. They transform electricity into light and light into electricity with efficiencies never approached by what came before. Every facet of energy technology is now changing as a result, and faster than at any time before in human history – faster than in 1765, when James Watt invented his steam engine; faster than in 1876, when Nickolaus Otto invented the internal combustion engine; faster than in 1879, when Edison patented his light bulb.

Efficient though they are, these new technologies don't reduce demand for energy, they expand it. And they don't exhaust supply – they expand it still more.

Burn Rate

The pursuit of more power in less space began centuries ago, with raw fuels at the very bottom of the energy food chain. A century ago we burnt mainly carbohydrates. Today, the United States consumes about 100 quadrillion BTUs (Quads) per year of dense, high-grade, fossil and nuclear fuel.

A Quad's worth of wood is a huge forest – beautiful to behold, but bulky and heavy. Pound for pound, coal stores about twice as much heat. On average, oil beats coal by that much again. And a gram of uranium-235 is worth about three tons of coal. The proponents of solar, wind, biomass, and other low-energy-density "renewables" are thus pushing back against a powerful historical trend. Left to its own devices, the market has not pursued thin, low-density fuels, however inexpensive – it has paid hefty premiums for fuels that supply more energy in less space.

Above the primary fuels, all the rest of our energy economy can be mapped out on a similar pyramid, in which lower-grade forms of power give way to higher. Of the 100 Quads of thermal energy we consume, less than 33 Quads are used in the form of heat itself. All the rest are used in combustion engines to spin shafts – raw heat giving way to kinetic energy. About half of that shaft power is used for transportation; the other half is immediately transformed again – it spins generators to produce electricity. And some share of the electricity – about 6 Quads worth, if we continue to track everything by the amount of thermal fuel used at the outset – goes through additional stages of refinement, to yield the extremely well-ordered power required to drive such things as microprocessors and lasers.

The price of power rises sharply with each step up this pyramid. Relatively unreliable grid power retails for 10 cents per kWh. The same amount of thermal energy locked up in raw coal costs about 1/3 of a cent. Computer-grade power, backed up by uninterruptible power sources and many additional layers of power-conditioning electronics, costs $3 or more.

From primary fuels to laser light, the overall trajectory of energy technology is thus defined by a single overarching trend: rising power density. Oil derricks, car engines, and microprocessors all run much faster than they used to, and handle more power in less space. The coal furnace of an old-fashioned steam locomotive is large and heavy, and it turns over slowly; the gas turbine under the wing of a jumbo jet burns much better fuel, much faster, in much smaller casing. The same, powerful trends are evident across the board, in all technologies that move energy up the staircase, from heat to motion to electricity to light.

More compact, higher-speed devices run more efficiently – a point to which we return shortly. But even so, they end up burning more energy, not less. If nothing else changed, less weight and higher efficiency would lower fuel consumption. But other things do change – we find new things to do with faster, more compact, more efficient power-processing engines. The <u>more</u> and the <u>faster</u> invariably boost demand for energy more than size and efficiency improvements lower it.

Our consumption of fossil fuels has been rising since the dawn of the thermal engine, since the days of James Watt, Nickolaus Otto, Rudolf Diesel, and Henry Ford. Our consumption of electricity has risen since the days of Thomas Edison, George Westinghouse, and Nikola Tesla. Americans consumed about 7 Quads of primary fuel in 1910, about 35 Quads in 1950, and about 100 Quads per year today. It will reach 150 Quads long before the middle of this century. We consumed about 117 billion kWh in 1900, about 334 billion kWh in 1950, and about 3,500 billion kWh per year today. By 2020, U.S. energy consumption will reach 130 Quads, and 5,000 billion kWh. (**Figure** 2)

Figure 2. More Energy
U.S. Annual Energy Consumption

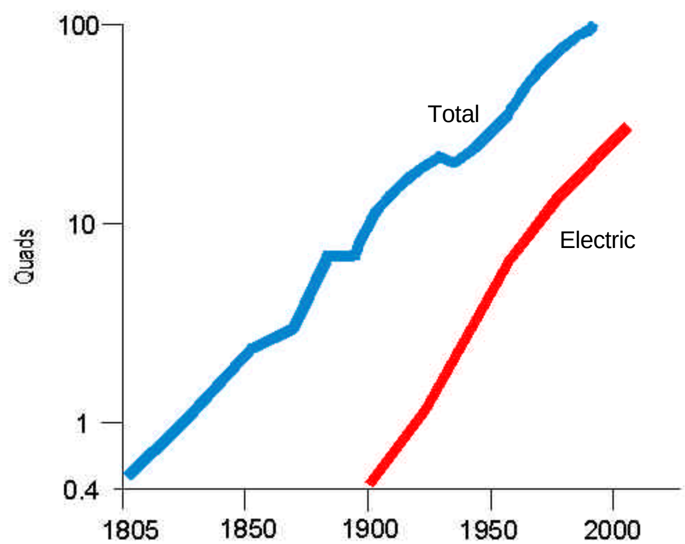

Source: Energy Information Association and Historical Statistics of the United States, U.S. Dept. of Commerce.

Predictions to the contrary – predictions of exhaustion, stasis, and lethargy – are based on wishful thinking about efficiency, that cannot be reconciled with actual experience. Or else they are all based on static views of technology, at a time when energy technology is anything but static. They overlook energy's capacity to find and capture energy itself. The more energy you have, the more you get. And the more you use, the more you want. Energy begets more energy.

Expanding Resources

The first "deep water" wells stood in 100 feet of water in 1954. Today, they reach through 10,000 feet of water, 20,000 feet of vertical rock, and another 30,000 feet of horizontal rock. Yet production costs in the hostile waters of the Statfjord oil fields of the North Sea are not very different from costs at the historic Spindletop fields of southeast Texas a century ago. An industry-wide collaborative effort (Deepstar) is now developing ways to drill and work reliably below 10,000 foot depths. In the past five years, deep-water rigs have yielded some 5.4 billion barrels of oil. The forecast for the next five is 20 billion. What happened? (**Figure** 3)

Figure 3. Unlimited Resources

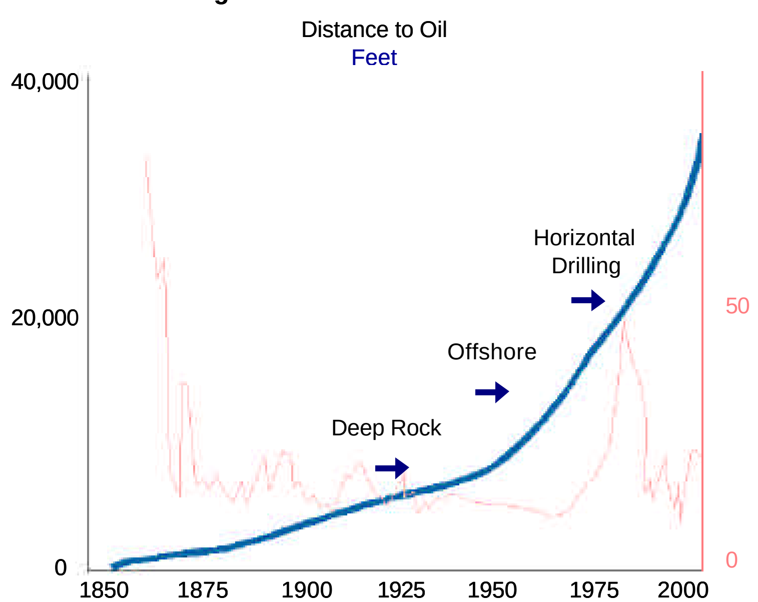

Source: Energy Information Association, ExxonMobil, and J. Ray McDermott Inc.

To get energy out of the earth you first have to project power into it – first to find out where the deposits lie, then to bring them to the surface. When James Watt developed his coal-fired steam engine in 1765, his main objective was to improve the process of pumping water out of coal mines. Colonel Drake used coal-fired steam to drill for oil. Oil is now used to extract oil – huge diesel engines power today's drills. And strange though it may sound, it is electricity, from the top of the energy pyramid, that now makes today's oil companies so much more efficient at extracting raw fuel from the bottom. Thus, it has been the increasingly intelligent use of energy itself that has continuously expanded our energy supplies, from James Watt to the present day.

Seeing comes first in the oil industry, because it's the dry holes that drive up the cost of drilling astronomically. And seeing begins with the projection of power, in the form of light, radar, and sound. Land-based oil production starts with satellite imaging to locate promising geological areas, followed by seismic (low-frequency acoustic) imaging, that can look through rock, salt, and sand in much the same way as ultrasound discerns a fetus in a woman's womb. On the seabed it's acoustic imaging from the get-go. Long wavelength pulses are generated by an electric thumper (a giant low-frequency loudspeaker) or compressed water/air guns; a fanned array of detectors (hydrophones) pick up what bounces back; computers then generate three-dimensional images of what lies up to six miles into the earth.

What the images generally reveal in a rich field is a jumble of isolated pools of oil scattered across a several-mile area deep underground. The most cost-effective way to get to all the pools is to build a hub-and-spoke network of pipes on the seabed. Many separate bore holes bring the crude out of the rock; the network then channels it to a central node, which feeds the oil up to a production platform on the surface. This allows fewer production platforms – the most expensive component – to serve the whole field. But it requires elaborate engineering, piping, networking, pumping, and repairing down on the seabed.

For many years, exploitation of the North Sea fields therefore depended on human divers, who worked (limited) depths in huge steel pressure suits and small submersibles. At the peak, some 1,400 intrepid divers were employed to keep the oil and gas flowing. Today, the offshore oil industry relies instead on Remotely Operated Vehicles (ROVs), which are powered and controlled via a tether that leads up to a mother ship. The ROVs carry complex arrays of sensors – fiber-optic gyroscopes, ground-penetrating sonar, acoustic imaging systems, low-light digital cameras, along with integrated arrays of light-emitting diode lamps for illumination; and high-power scanning laser systems. They search for oil, inspect equipment, move and connect gear, pull and connect pipes, and actuate valves. They dig trenches and bury cables, align the drill with the bore hole, and provide maintenance.

And because you can't run a combustion engine for any length of time under water, the ROVs depend for their power on electricity. It's generated on the mother ship at the surface, boosted to thousands of volts for more efficient transmission, and channeled down a monstrous tether that can run as much as five miles long, and weigh up to 16 tons. Power electronics on the ROV then step down the voltage and control the current to light the sensors, power the electric motors that provide the thrust for the ROV, and provide high-pressure hydraulics for the manipulators and the grabber arms.

The seabed is just one of many lines where advancing technology meets retreating resources. Electricity pursues oil. It also pursues sunlight – we use high-grade electricity to purify silicon, and then dope it, to end up with photovoltaic cells that can generate still more electricity. High-power lasers are used to enrich uranium, which is then used to generate more electricity, which can produce more laser light. Power pursues and captures more energy, which produces more power.

Energy itself is abundant; how much we can channel to our own ends is determined not by "what's out there" but by how good we are at finding and capturing it. For the first two centuries of industrial history, the technologies used to find and extract fuels improved much faster than the horizon of supply receded. All that is different today is that technology is gaining ground even faster. Viewed in that light, the planet's resources aren't contracting, they're expanding.

The Efficiency Paradox

Can't technology cut demand as readily as it boosts supply? It seems obvious that rising efficiency in cars, furnaces, lights, and lawn mowers could, in the aggregate, achieve as much as rising efficiency achieves in the offshore oil industry. So obvious, indeed, that the pursuit of more efficiency has become a cornerstone of national energy policy. Efficiency is said to be cheaper than oil, more reliable than Arabia, and much better for the planet. If we take the opportunity seriously, efficiency will curtail demand even faster than it expands supply. Or will it?

Higher efficiency is certainly coming. The ROV is made possible by power conversion technologies that are extremely fast, compact, and efficient. These same technologies can boost efficiency in lathes, drill presses, SUVs, and refrigerators. And they are doing so. Within a decade, the power train of a Buick, for example, will look a lot more like the power train of mother-ship-to-ROV. Most of the power generated by a car's internal combustion engine will be converted directly into electricity, and then conveyed to direct-drive electric motors that run the fans, control steering and braking, operate the engine valves, and ultimately drive the wheels. Such technologies will replace hundreds of pounds of click-click bang-bang mechanical and hydraulic hardware found today under the hood of almost every car and truck. They're already being rolled out commercially at the top end of the automotive market, because they deliver better performance, even as they lower weight and boost overall efficiency.

Lighting efficiency is improving at least as fast, because electron-to-photon transitions are migrating from hot filaments to quantum junctions. Light-emitting diodes (LEDs) now offer far better performance than either incandescent or fluorescent bulbs. Solid-state LEDs shrink the "bulb" from the size of a pear to the size of a poppy seed. Per unit of area and of energy used, semiconductor lights are much more compact, efficient, and cool. Silicon carbide LEDs have reached a stunning 28 percent electron-to-photon conversion efficiency; incandescent bulbs typically run in single digits. It is now reasonable to project that solid-state lights will completely displace Edison's filaments within the next few decades.

The laser diode takes solid-state light another major step further. And lasers can be far more efficient than thermal and electrical technologies when they are enlisted to heat-treat materials, move ink in printers, etch silicon and metal, solder opto-electronic chips, cure epoxies, bond polymers, weld, sinter, burn hair, cauterize tissue, and reshape the surface of the eye. They can be aimed and focused with unequaled precision, which means that they illuminate, move, or heat more payload and less extraneous real estate. A fiber-optic system requires far less power to transmit bits than an electric wire. Microwaves can heat just the water in the soup, not all the air and stovetop around it; lasers do the same, only more so.

In his 1958 book *Energy for Man: From Windmills to Nuclear Power*, Hans Thirring surveyed the 250-year history of efficiency in these kinds of engines. So far as all the major trends go, nothing fundamental has changed since. Efficiencies have been rising for as long as there have been thermal engines. (**Figure** 4) Two centuries ago, they couldn't beat 10 percent. By raising boiler temperatures and pressures, engineers pushed that to about 20 percent by the turn of the twentieth century. By mid-century, they were up to approximately 40 percent. Today, the best thermal plants routinely hit 50 percent efficiency.

Figure 4. History of Engine Efficiency

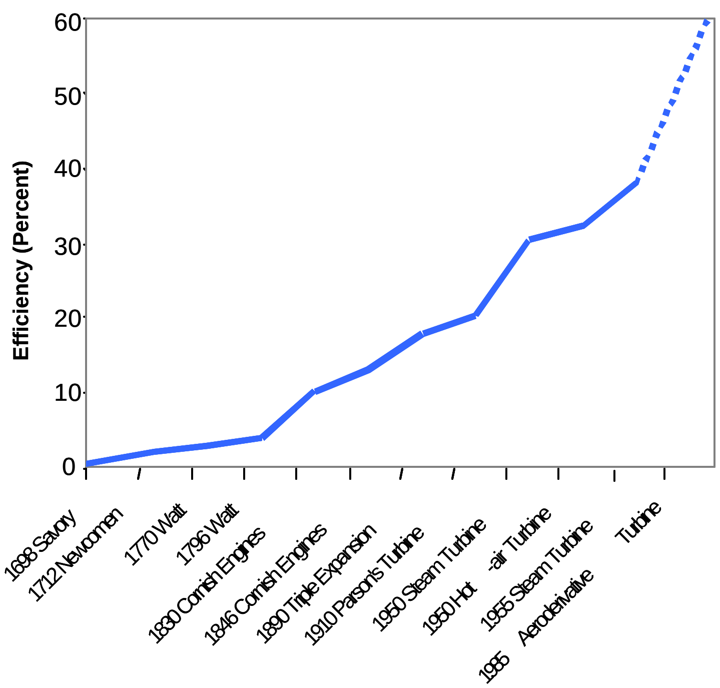

Source: Thirring, Hans, *Energy For Man: From Windmills to Nuclear Power* (1958).

Efficiency gains this large ought to have had a dramatic impact on supply and demand. And indeed they have. The long-term price of miles in cars or planes, and kilowatt-hours from the grid, has fallen steadily. And the total amount of fuel consumed in cars, planes, and electric power plants has risen – not fallen – equally fast. (**Figure** 5)

Efficiency does undoubtedly curtail demand, in the short term, and for the specific task at hand. But its long-term impact is precisely the opposite. When steam power plants, jet turbines, and car engines were less efficient than today, they also consumed less total energy. Light bulbs, electric motors, air conditioners, and computers have all grown more efficient, year after year. But the total amount of electricity they consume has risen apace.

Figure 5. Efficiency & Demand

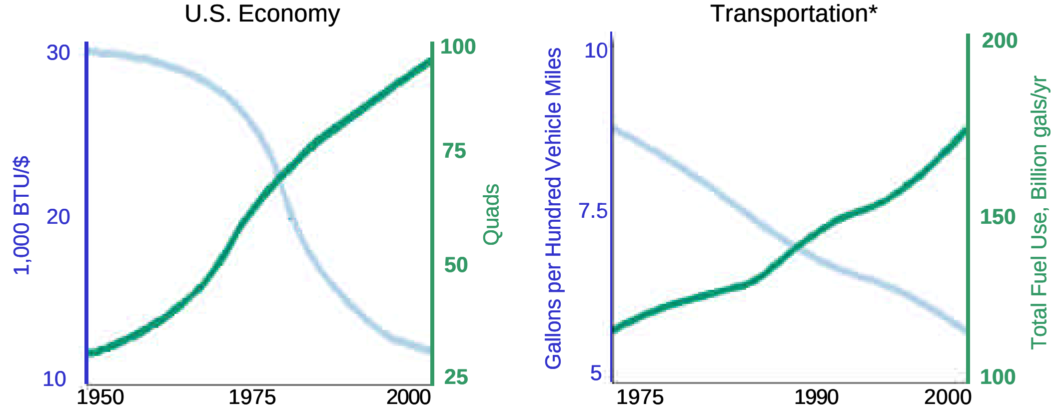

Source: Energy Information Association and Dept. of Transportation *All modes

Pyramids of Demand

Efficiency fails to curb demand because it lets more people do more, faster – and the more/more/faster completely swamp the efficiency gains. In cars, jets, printers, phone lines, and computers, more speed requires more power, and more total energy; the task at hand gets completed faster; which gives us time to do more. (**Figure** 6)

Figure 6. Travel Distance

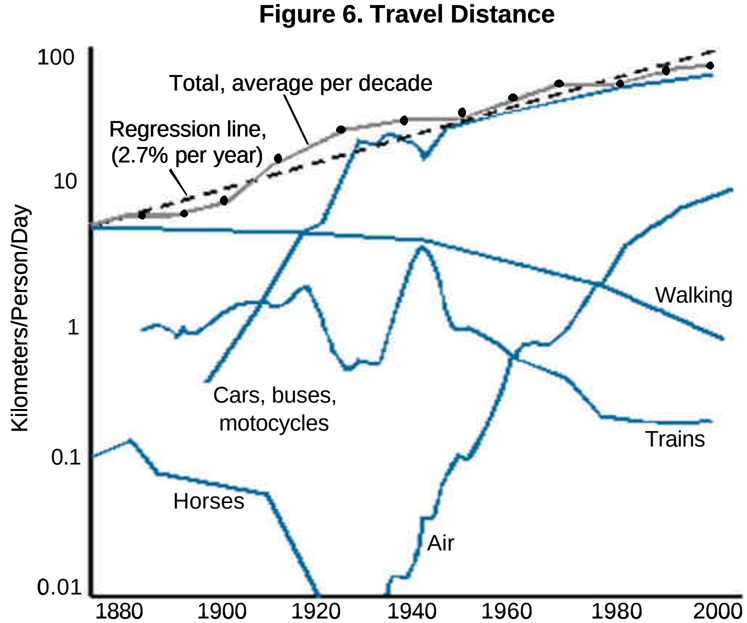

Source: Ausubel and Marchetti, *The Industrial Physicist*, April/May 2001.

We live in bigger houses, and wash more dishes, and as our cars get faster, easier to drive, and more comfortable, we drive more, and we fly more, too. Over the long term, our total travel distances move up almost 3 percent per year. The car didn't "substitute for" the horse – most of the miles we travel today would never have been traveled at all on horse. The jet didn't substitute for the train – air-travel miles are mainly new miles, too.

Microprocessors, those wonders of ever-rising efficiency, tell the whole story in a sliver of silicon. The power consumed by a single gate depends on the gate's size; the gates keep shrinking; and so the electrical energy required to process a single logic instruction is cut in half about every 14 months. But the smaller the gates, the faster they run – that's at least half the reason for making them smaller in the first place. And the number of gates per chip rises as fast as the gates themselves shrink, and faster still, as chips get bigger. So overall, the number of bits processed rises much faster than bit efficiencies improve – and the amount of power consumed by the chip doubles about every 36 months. (**Figure** 7) And the number of chips in use rises faster still. A Nintendo is far more efficient than the original ENIAC computer, which consumed 174,000 Watts of electrical power to light its 18,000 vacuum tubes in 1946. But one Nintendo per teenager adds up to a lot more total power than one ENIAC per planet.

Figure 7. Power

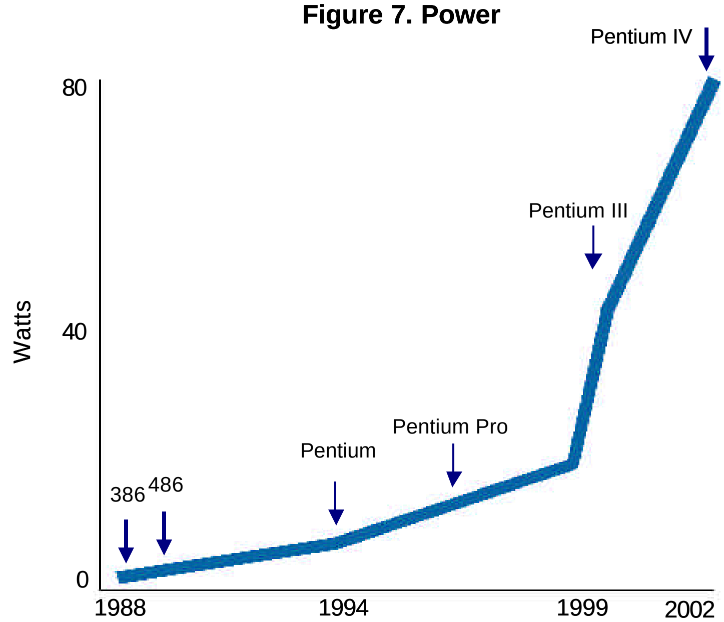

Source: Intel.

The ENIAC-to-Nintendo trajectory illustrates a larger point. When you develop faster, more compact technologies, you don't just rebuild the SUV. You also begin building Segway scooters, skidoos, leaf blowers, and pilotless Predators to fly reconnaissance over Afghanistan. New miles, new types of demand multiply and take hold faster than the new efficiencies do. The Segway scooter has been touted as a replacement for cars. What's far more likely, is that it will displace walking or biking. The new scooter will get loaded into the back of the old SUV to save the effort of trudging around the mall.

Almost all of the new demand for power emerges at the top of the energy pyramid – it is demand for high-quality power. And it takes a lot of low-grade power to produce much smaller quantities of the high-grade. Digital devices, for example, require layer upon layer of power conditioning electronics, batteries, and power backup systems, along with a lot of air conditioning, too. To get a kilowatt-hour of ultra-pure, ultra-reliable electricity into a microprocessor at the very top of the energy-consumption pyramid you have to start with a hundred times as much energy in the raw fuel that gets burnt in the power plant. (**Figure** 8)

Figure 8. CPU → Coal Mine
@ 2 GHz

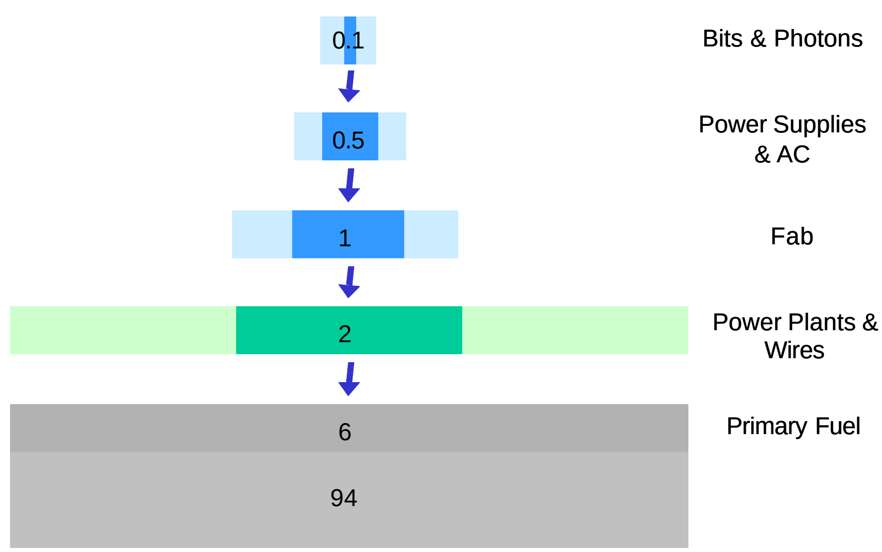

Such pyramids of demand are universal, and they invariably have similar demand-amplifying structures. The Mom-and-kids in an SUV are microprocessors too (so to speak) and they require equally high-quality power. They get it from a big engine, so they can get places fast and accelerate out of harm's way when they need to, along with a big, heavy frame, to support the big engine, and to make their ride smooth, and shield them from what happens when high-speed energy fails unexpectedly – when the SUV hits a pot-hole, say, or a lamp post. As a result, only a tiny fraction of the energy that starts its journey in an oil pool two miles under the Gulf of Mexico ends up being used to propel two-hundred pounds of mom and the kids – the real human payload – two miles to the soccer field. (**Figure** 9)

Figure 9. Soccer Mom → Oil Well

@ 60 mph

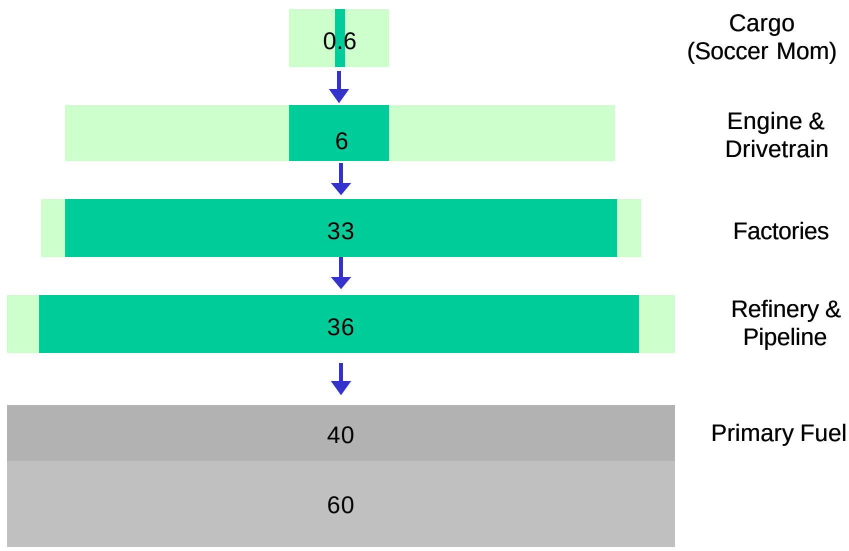

With the microprocessor and the SUV alike, demand gravitates to the very highest quality power because that kind of power is far better than any alternative. Get your power to the very top, and you can do things very cleanly and efficiently at the very end of the line, where those qualities matter the most. You can also do a lot of new, improved things – spreadsheets and laser surgery, for example – that just can't be done at all with lower-grade fuel. When the light bulb gives way to the solid-state laser, you don't get a better reading lamp to help your tired old eyes, you get laser surgery to give you back the eyes of a twenty-year old. Every major step up in the quality of power spawns a slew of new and improved uses for it.

All the focus on efficiency is aimed at reducing the losses in intermediate tiers. But the efficiency opportunities arise only after the new tiers come into existence, and get well established. No one could improve gas mileage until there were gasoline engines, and no one even bothered to try until Detroit was building tens of millions of them. And even then, efficiency is only improved at the margin, when old cars, light bulbs, and motors wear out and get replaced by new ones. At the top of the pyramid, by contrast, everything is new, and we keep discovering new ways to use these new forms of power.

The real story of the "new economy" isn't that it moves a lot of bits. The new economy moves a lot of very-high-quality power, some of which we call bits. Power that good can move an enormous range of other things, too – from Predators to Segways. And it does. Efficiency will never overtake demand until we stop finding new ways to use better power at the top of the pyramid. And we won't ever do that.

Digital Demand

Unless, perhaps, the bits themselves displace a lot of heavy lifting. One theory holds that they do. In the new Harry-Potter economy we won't have to drive to the mall to pick up the latest kids' bestseller; we'll order it from Amazon. And Amazon's business model requires less energy than Barnes and Noble's. An economy that moves bits, rather than atoms, just doesn't need as much energy.

Joseph Romm makes the case in a 1999 book, *Cool Companies,* and in related articles. GDP historically grew in lock step with energy consumption, but it no longer does. Since about 1996, "energy intensity" – the energy consumed per dollar of economic output – has been declining in all sectors of the economy, due to "improvements in technology, large expenditures by businesses, the network effect, and the growth in the labor force of a generation raised on electronic interaction." "(T)he fundamental relationship between energy use and economic growth … has been changed permanently by the spread of New Economy technology." Online stores substitute "clicks for bricks." Delivery trucks are more efficient that car trips to malls. Wired supply chains reduce inventories, cut overproduction, reduce unnecessary capital purchases, eliminate paper transactions, reduce mistaken orders, and thus save energy all around. The Internet could also "have a large positive impact" on reducing energy consumption in the transportation sector.

The first problem with this argument is trivial, but telling. Energy consumption *per unit of GDP* has indeed been falling steadily – not just since 1996, but for the last 20,000 years. Any different conclusion is based on trend lines that ignore the carbohydrate energy economy, which continued to play a large role even in industrialized countries until well into the twentieth century. The GDP of our hunter-gatherer ancestors was nothing but energy, in the form of food calories. Wealthier economies add less energy-intensive goods to the mix, so energy per unit of GDP invariably falls: it doesn't take much oil to make software or to run a symphony orchestra. But this "dematerialization" of the economy is entirely relative. Our consumption of knowledge-intensive goods grows faster than our consumption of energy-intensive goods – but both continue to grow. Adding a lot of energy-lite consumption is like adding artificial sweetener to a diet soda – it doesn't, in itself, lower total calories consumed.

And the information technologies that add so much economic value to GDP aren't diet sodas, either. To begin with, it takes a lot of energy to manufacture the hardware of the virtual economy. As a rough rule of thumb, the manufacturing of a digital box consumes as much power as a year of its operation. The quantum physics at the p-n junction is exquisitely sensitive to basic chemistry; the manufacture of semiconductor devices therefore requires one-in-a-thousand-trillion levels of purity, which require exorbitant amounts of electric power to achieve. The Pentium itself isn't very heavy, but you have to move a huge amount of material through furnaces and filters and so forth to get semiconductors to the levels of purity required to make them work. All in all, it takes about 800 kWh of electrical energy, for example, to manufacture one 200-mm semiconductor wafer, enough to power a typical household for about a month. An entire chip fab can use 30 to 50 megawatts of electrical capacity, enough to power a small city. And twice as much again is used on the premises of suppliers of ultrapure copper, iron, zinc, cadmium, tantalum, gold, mercury, and other materials that go into the fabrication of digital logic.

Once the semiconductors are manufactured, it takes a lot of power to power them. The United States now has an installed base of about 1 ½ computers per household (two-thirds of them outside the home), and many of them are left running around the clock, creating 10 to 100 Watts of background load. More chips are kept lit around the clock, because we want them to respond instantly on demand, and because much of their work involves continuous interaction with other equally tireless chips. A broadband connection adds another 15 Watts at the householder's end of the line – and still more in the mirror-image equipment on the premises of a service provider. A Palm Pilot recharging cradle runs between 2 Watts (when empty) to 12 Watts when the Palm is in it. To put these numbers in perspective, the average American currently consumes about 1,400 Watts, when all residential and commercial loads are averaged out per capita and across the twenty-four-hour day.

Finally, information networks have unique character: the bigger they get, the more active they become. Turning on a light in your own home does not impel your neighbors to turn on theirs, but wired computers and servers do have that kind of effect. With light bulbs, a handful per room saturate the potential for demand, but there is no such upper limit to demand for bits.

Figure 10. New Economy Infrastructure

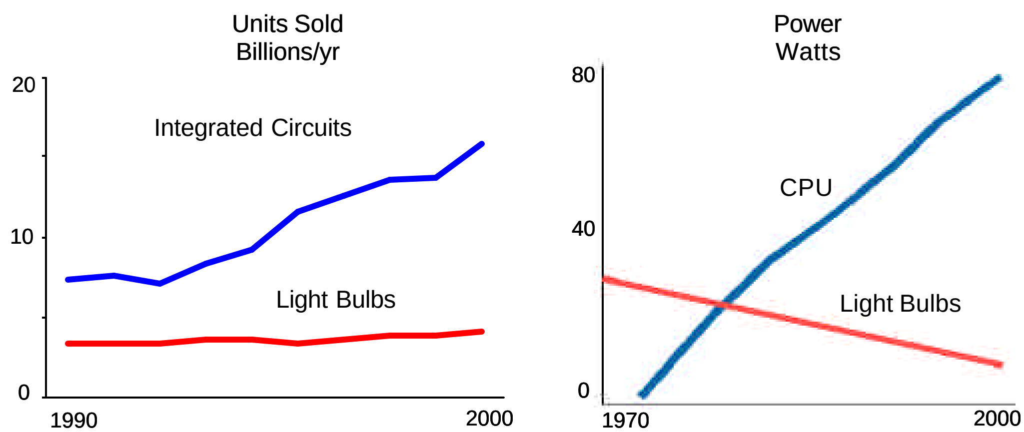

Source: U.S. Census Bureau, Semiconductor Industry Association, and Intel.

The electrical age began with Edison's light bulb – an information-gathering device, of sorts, in that it lets us see what's going on. The ascendant electric load of the 21st century is the silicon chip. The total number of bulbs sold per year in the United States is now holding fairly constant – about 40 billion bulbs were sold in the past decade, most of them as replacements. But U.S. vendors ship about 20 billion integrated circuits every year. The power consumed by the average light bulb is declining; the power consumed by the average microprocessor is rising. Demand for light bulbs plateaus once everything is well lit. But logic engines supply data, entertainment, connection, control, and in due course they will supply insight, understanding, knowledge, intelligence, and even wisdom. There is no reason to suppose that humanity will ever sate its hunger for any of that. (**Figure** 10)

Carbon Rationing

With global warming, the freeze-in-the-dark scenario of fuel exhaustion gives way to one of fuel supplies so abundant that burning them will (indirectly) warm the whole planet. The way things are headed now, carbon-limiting controls of one kind or another are likely to emerge, though probably later rather than sooner. Their almost certain effect, however, will be to consume more fossil fuel, less efficiently. Big Oil will prosper. So will Saudi Arabia.

Fossil fuels – hydrocarbons – are part hydrogen, part carbon. Both components burn well – charcoal barbecues burn pure carbon, the Space Shuttle's main engine burns hydrogen. Natural gas contains relatively more hydrogen and less carbon; over half of the heat from a methane-gas flame comes from the hydrogen, which is why environmental regulators like it more than other fossil fuels. About 80 percent of the heat in coal, by contrast, comes from the carbon.

One way to burn hydrocarbons without raising carbon dioxide concentrations in the atmosphere is to burn first, then suck the carbon back out of the air – by growing new trees, for example, or by seeding the oceans with iron to promote massive new growth of algae. If things go right, the new green plants then get buried, and in the fullness of geological time, turn into new coal – carbon recycling, the old-fashioned way. This is the cheapest approach, but it doesn't curb energy consumption – which means that it will never satisfy many environmental activists. "Giving society cheap abundant energy at this point would be equivalent to giving an idiot child a machine gun," Paul Ehrlich declared in 1991. That wing of the environmental lobby isn't likely to settle for anything as cheap and low-tech as carbon sequestration by way of new trees.

A second option is to strip the carbon out of the fuel before it's burnt. Most fuel cells depend on "reformers" that do just that. Most current reformers then turn around and dump the stripped-out carbon into the air as carbon dioxide, but other processes can leave the carbon behind as a solid residue. That takes care of carbon emissions all right – roughly speaking, such a fuel cell starts with methane, extracts the hydrogen, and then returns coal to the ground. But it also requires two to four times as much raw fuel to deliver the same amount of heat from just the hydrogen side of the hydrocarbon. (**Figure** 11)

Every other lower-carbon alternative entails a significant energy overhead, too. Scrubbing huge amounts of carbon dioxide out of smokestacks and tail pipes isn't easy, and can't be done without sharply reducing overall thermal efficiency. We already pay significant efficiency overhead for the scrubbers that remove comparatively tiny amounts of sulfur dioxide from coal-plant smokestacks, and for the catalytic converters that remove comparatively tiny amounts of nitrogen oxide pollutants from tailpipes.

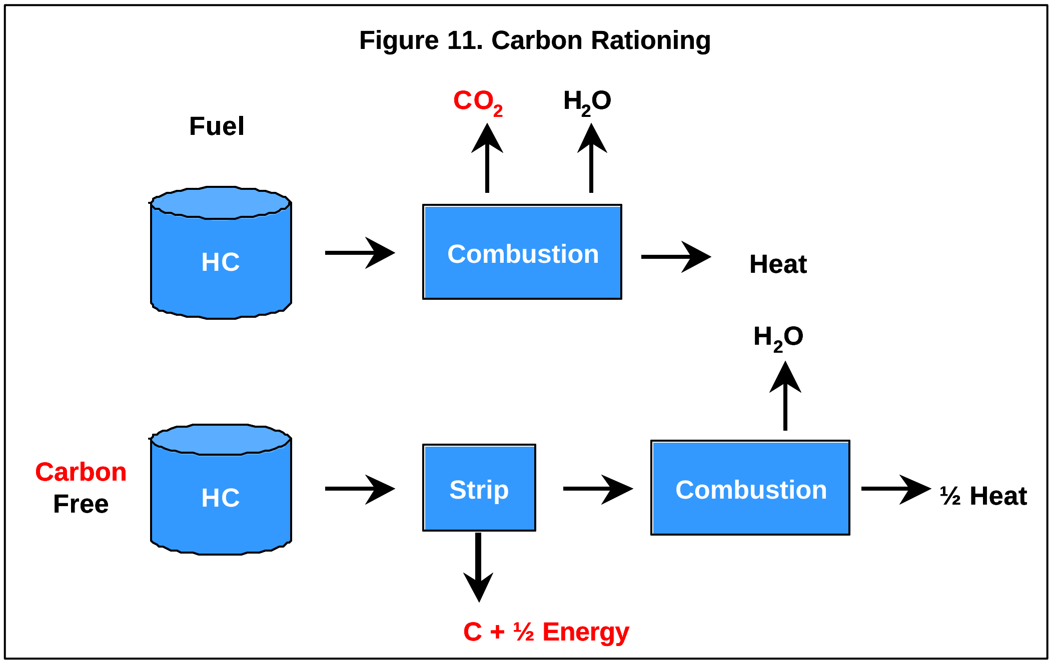

Some energy pundits hope that we can eventually get beyond hydrocarbons entirely. Hydrogen is abundant, they remind us, it's part of water. But water is what you end up with after you burn hydrogen. To turn water back into hydrogen fuel again, you have to un-burn it – you have to pump energy back into the water, usually in the form of electricity. Electricity supplied by coal or gas-fired plants won't do – that would defeat the whole carbon-reduction objective. Nuclear power plants could be used, but if there's one thing that many environmental activists hate even more than carbon, it's uranium.

So at the very end of the day, we probably will burn hydrogen, the most pristine of all fuels, just as many environmentalists hope. But we will also quietly accept that the only readily available supplies of hydrogen in un-burnt form are in gas, oil, and coal. Our national policy will be to burn more fossil fuel, less efficiently – to extract a lot less heat out of the same amount of fuel, even as we struggle to design new car engines that get more miles out of the same amount of heat. We will raise efficiency in every tier of the energy pipeline except at the source, where we actually burn fuel – back there, we'll reduce it, drastically, far more than we can raise it in any of the tiers above. While this may or may not save the icecaps, it certainly won't ruin the hydrocarbon companies.

Renewable Fuels and Alternative Technologies

Finally, there's the idea that one way or another, our energy destiny lies in renewables and "alternative" energy technologies – wind, solar, and fuel cells, among others.

The fuel cell is remarkably elegant in its basic operation. It goes from fuel to electricity in a single step, cutting out two of the intermediate stages (turbine and generator) that are used in generating electricity the old-fashioned way. What has attracted the most interest is cell design that runs relatively cool (at about 200ºF), and depends on a proton exchange membrane (PEM) impregnated with a platinum catalyst to promote the hydrogen-oxygen reaction. The recent surge of interest in fuel cells can be traced back to breakthroughs in platinum chemistry and solid electrolytes achieved only a decade or so ago.

A fuel cell can be remarkably clean, carbon-free, and compact. The alkali fuel cells used by NASA offer the highest power-to-weight ratio of any electrical generator ever devised, in a technically magnificent – though tricky and dangerous – package. But PEM-membrane fuel cells require exceptionally pure hydrogen as their fuel, because the membranes are quickly poisoned by carbon. And that's what's wrong with most fuel-cell scenarios – they assume a ready supply of the very highest-grade combustion fuel available. With perfect fuel, however, combustion engines run much better, too. The fuel cell won't prevail by beating gasoline engines; it has to beat out hydrogen-fueled combustion engines. As of now, combustion engines remain much cheaper and more reliable.

Solar power has even more intuitive appeal than the fuel cell. Begin with crystalline silicon, dope it just right with boron or phosphorous, build the interface with meticulous care, and you end up with a surface that can readily beat green plants and the rest of nature in extracting power from sunlight. Here, atomic-scale logic in the semiconductor replaces the entire sun-plant-fossil-furnace-turbine-generator sequence. We can now imagine a world in which the extremely high power densities on the square-centimeter silicon surfaces of the microprocessor are supplied by square meters of photovoltaic silicon a short distance away.

And if not photovoltaic technology, then something else much like it. The technologies that make possible the ROV also make possible the highly compact and efficient turbine in the modern windmill. Other comparable technologies can now tap into the enormous amounts of low-grade energy that wash through our environment – waves, thermal gradients in oceans and ponds, all the energy currently wasted in brake linings and shock absorbers; and much of the energy locked up in household trash, and agricultural waste. Extraction engines can now be built to pull high-grade power directly out of the most dilute and ubiquitous sources of energy on the planet, not just the most concentrated ones.

If this is good news, however, it is indiscriminately good. Sunlight is ubiquitous but dilute; so too are the heavy isotopes of uranium and hydrogen. Tar sands contain vast quantities of fossil fuel that we hardly touch at present. Even vaster quantities of methane hydrates are apparently locked up in sea-floor sediments and the arctic permafrost. There are, in short, still huge reservoirs of oil, gas, coal, tar sands, and other fossilized debris yet to be tapped. If advanced technology can squeeze useful energy out of a source as thin as sunlight, it can squeeze hydrogen out of much richer sources like coal beds or tar sands. And for the foreseeable future, the earth's fossils will remain much easier targets than the sun's photons.

That the renewables are "free" doesn't resolve anything; fossils are "free" too, the expensive part lies in finding and extracting. That renewables "won't run out" doesn't either – for all practical purposes the supplies of hydrocarbons are infinite, too, once we lower our sights to include thinner, harder-to-process supplies. The cost of energy has nothing to do with some distant theoretical date when we "run out." It is determined by the here-and-now practicalities of extracting high-grade energy from dispersed, low-grade sources. The renewables don't have to beat only the offshore oil platform that's probing through 10,000 feet of water. They have to beat, as well, all the next-generation hydrocarbon extraction technologies that will gasify coal or liquefy tar sands, unleash methane from seabed hydrates, and squeeze hydrogen out of hydrocarbons.

The most common mistake made by proponents of renewable fuels and alternative technologies is to assume that conventional fuels and technologies will stand still while the new arrivals advance. They won't. The coal, oil, and gas industries, together with all manufacturers of conventional thermal engines, have strong incentives to make their side of things more efficient and environmentally acceptable, too – and they have more resources at their disposal. Fuel-cell advocates assume a ready supply of a perfect fuel, and proceed from there. But they ignore the fact that perfect fuel makes lots of energy technologies – including conventional thermal engines – look very attractive. The solar cell advocates assume the perfect energy-scavenging technology, and assume that it will be sent out to scavenge sunlight. But they ignore the fact that advances in the technologies of energy capture and extraction will almost inevitably facilitate the capture of higher-grade energy supplies before the capture of lower.

And they ignore the fact that the technologies that expand supply also expand demand, and vice versa. Cars become more efficient; so do the "transportation" technologies that extract oil from deep under the surface of the earth. Semiconductor fabs manufacture photovoltaic cells; they also manufacture Pentiums. The PV cells capture perhaps 20 Watts per square meter, on a round-the-clock average. The Pentiums consume 20 Watts per square centimeter – which is to say, 10,000 times more. It is much easier to roll out PV cells by the meter than Pentiums by the centimeter – but demand for new acres of integrated circuits is growing much faster than demand for new acres of solar cells. The output of all the silicon-based solar panels on the planet today could not begin to meet the demand for power created by all the silicon-based integrated circuits.

Unlimited Supply, Insatiable Demand

Ours is a blue-whale energy economy, in which gargantuan quantities of high-grade mammalian power are supplied by vast but increasingly dispersed reservoirs of microscopic, low-grade plankton. There is a wide chasm in our energy economy, between demand centered on increasingly fast, compact, intense modes of consumption, and supply derived from increasingly distant and dilute forms of energy. It is getting steadily wider, too – as the oil industry's history reveals. But that doesn't stop us. The more energy we consume, the more we're able to find and capture. It's a chain-reaction process, and it spirals up, not down. It is, if you will, a perpetual motion machine.

Four billion years ago, life on Earth captured no solar energy at all, because there was no life. Life then got a foothold, and the capture and consumption of energy in the biosphere has been rising ever since. The thicker life grew on the surface of the planet, the more energy the biosphere managed to capture. And it used all that energy to create more life. Along the way it deposited huge amounts of biological debris underground.

Humans eventually arrived to incorporate that debris directly into their own energy cycle – we are the ultimate scavengers, we consume not just carrion but fossils, too. An organism called James Watt emerged from the biological cauldron with an idea about how to dig up the debris more efficiently. Today electricity extracts oil. When Enrico Fermi built the first fission reactor, the idea was to use one neutron emitted by a uranium atom to kick out two neutrons from other uranium atoms nearby.

None of these processes produces "perpetual motion" in the strict thermodynamic sense, of course – they all just improve on the process of grabbing energy from somewhere else. Living green plants still capture solar energy about three times as fast as we humans are able to dig up dead plants in the form of fossil fuels, but we'll overtake the rest of nature in the not too distant future. And perhaps some day we'll get to the point where we, too, can take much of our energy directly from the sun. There's certainly plenty of solar energy to spare – green plants currently capture only about one-three-thousandth of the solar energy that cascades onto the surface of the Earth.

But whether we catch our solar energy live, or dig it up in fossilized form, or dig up uranium instead, is really just a detail. The one near-certainty is that we will extract more energy from our environment, not less. In the grand scheme of things everything we think we know about "running out of energy" isn't just muddled – it's the exact opposite of the truth. The more energy we capture and put to use, the better we get at capturing still more.

Demand expands apace. There is an upward spiral here too – the more we get the more we want. This is because our main use of energy is to convert and refine energy itself. Energy consumes itself at every stage of its transit from a lump of coal to a beam of laser light, or from a pint of buried crude to a high-velocity soccer mom. The whole business reeks of a Ponzi scheme, with each successive tier of the pyramid feeding voraciously off the one beneath. Small wonder that so much of our energy economy is often characterized as wasteful. Small wonder that people who don't understand energy are easily convinced that there must be a better way.

But the huge pyramid of consumption, with its withering "losses" at every stage, doesn't reflect bad engineering, it reflects, in all its practical complexity and real-world tangle, one of the most fundamental laws of physics, the second law of thermodynamics. Energy doesn't want to accommodate us. It doesn't lounge about just waiting for the chance to propel moms and kids to soccer fields. Getting things to that point is a huge uphill battle. The remarkable thing isn't that our power-conversion technologies are inefficient, but that they work at all. Until there was a steam or a gasoline engine, none of the energy locked up in oil could be dissipated in an inefficient trip to a soccer field. If raising efficiency now presents an opportunity to curb demand, it is only because so much energy came to be funneled through inefficient engines to begin with.

Why bother with all this funneling? Because a lump of coal can't do laser surgery. Because a microprocessor can't run on power straight out of a wall socket. Because high-grade energy lets us defend ourselves from domestic terror, and fight distant wars in which we suffer almost no casualties. Ordered power feeds, moves, informs, and entertains us. It jets us across oceans, burns away cataracts, and excises cancerous tumors. It projects X-rays through luggage in airports, and it pursues terrorists into the recesses of distant caves. It dispatches rockets to explore other planets, moves bits down a strand of glass, and powers computers. It projects a gigahertz beam toward a cell phone and a pulse of electromagnetic power toward an enemy tank. It turns a lathe, and the wheels of a car. It cooks food and dries paint. It warms and cools a reactor vessel at Genentech, in which a polymerase chain reaction replicates strands of DNA. We will never sate our hunger for more logic, more memory, more vision, more range – all of which higher-grade energy supplies – because we are built to want more of these things, an unlimited more.

Civilization, like life itself, is a Sysiphian flight from chaos. The chaos will prevail in the end, but it is our mission to postpone that day for as long as we can, and to push things in the opposite direction, with all the ingenuity and determination we can muster. Energy isn't the problem. Energy is the solution.