Technology Policy and Climate Change

Abstract

There is a strong foundation in theoretical and empirical research in economics for the proposition that efficient climate policy must include both carbon-price policy and technology policy. Even the most modest projections of Greenhouse Gas ("GHG") reductions needed to moderate climate change imply very large reductions in the carbon-intensity of the world economy, something in excess of a 60% reduction by 2050. This is a greater proportionate reduction than has occurred in the petroleum intensity of world GDP since 1970, despite a six-fold increase in the price of oil. This illustrates how unlikely it is that the needed economic transformation could be brought about by price-based policy instruments alone. There is no good historical analogue to the needed transformation, but the closest parallels all involved major roles for technology policy. Increased public funding of research and training is a necessary but not sufficient component of such policy. Historical experience with technological transformation in other sectors suggests that government support for purchases of low-carbon technologies will be needed. Unfortunately, we do not have good evidence on efficient design of such programs. We need systematic evaluation of different policy instruments designed to accelerate the transformation of basic technologies into large-scale commercial products. We have the "technology" to do this kind of systematic evaluation, but it is not generally used.

Adam B. Jaffe
Fred C. Hecht Professor in Economics, Brandeis University
Research Associate, National Bureau of Economic Research

Prepared for:

The Next Round of Climate Economics & Policy Research
Washington, D.C. October 2728, 2011

Revised January 2012

I thank Lawrence Goulder, Michael J. Holland, Ian Parry and conference participants for helpful comments. I alone am responsible for the views expressed and any remaining errors.

Introduction

I believe the evidence is overwhelmingly convincing that the climate is now undergoing significant changes, that these changes are caused at least in part by human activity, and that the probability distribution on the consequences of these changes includes significant mass at points where the social cost would be very large. I therefore work from the premise that it is a social objective to mitigate and adapt to these changes, and therefore the design of public policy to achieve such mitigation and adaptation in a cost-effective matter is an important concern.

I also accept the premise that the primary policy response to this challenge must be to raise the price perceived by private agents to be associated with the emission of so-called greenhouse gases ("GHG") through taxation or regulation, in order to create the appropriate economic incentives for such agents to align all of their economic activities—production, consumption, and investment of all kinds—with the social objective. Other papers in the conference will no doubt explore aspects of these policies to "get the price of carbon right" in some detail.

In this paper, I will discuss the extent to which such "carbon-price policies" should be complemented by other policies in order to achieve an overall cost-effective policy stance. The other policies that I consider are those designed to foster the creation, improvement and diffusion of new low-GHG technologies by pathways other than the incentives created by a higher price on GHG emissions, which I will call, loosely, "technology policies." Since it is my understanding that the purpose of this conference is to foster discussion on fruitful areas of research, I will be generally provocative rather than cautious in my characterization of what is known and not known about the likely effects of technology policies in this arena.

In summary, I will argue:

- The magnitude of reduction in the carbon-intensity of the economy that appears to be necessary is large, and cannot be achieved without significant reliance on technologies not in use today, unless dramatic social and cultural changes are contemplated.
- It is somewhere between unlikely and inconceivable that the necessary technological change could be brought forth by carbon-price policy alone.
- There is no close historical precedent for the kind of technological transformation we need, but the closest ones we can identify were, in fact, supported by significant technology policies.
- There is a strong foundation in theoretical and empirical research in economics for the proposition that efficient climate policy must include both carbon-price policy and technology policy.
- Public funding of research and training is a necessary but not sufficient component of such policy.
- We do have some experience with what works well and what doesn't with respect to public policies designed to support the commercial improvement and diffusion of new technologies, but the record is murky.
- We need systematic evaluation of different policy instruments designed to accelerate the transformation of basic technologies into large-scale commercial products. We have the "technology" to do this kind of systematic evaluation, but it is not generally used.

The Problem is Big

I will not dwell on this discussion, but simply use a simplistic straw man to get started. If you think this strawman is way off base, you can skip the rest of the paper. My strawman is an assumed stabilization of global GHG emissions at current levels by 2050, accompanied by "business as usual" growth in world GDP over that same time period. Most serious analysts would say greater reductions are necessary, but it doesn't matter for my purpose. [1] If we assume a tepid 2.5% annual growth in world GDP, this means that the carbon intensity of world GDP needs to fall by about 60% over 40 years. If you accept that GHG emissions have to be stabilized at a lower level than this, the reduction in GHG intensity of world GDP would have to be even greater.

There are a lot of modelers who have explored in detail how this might come about, but let's look at it crudely. What kind of precedent might there be for a transformation of this magnitude? In the 40-year period since 1970, the world has been trying to reduce its dependence on petroleum. Indeed, Figure One shows that the oil intensity of the world economy has fallen by about 40% over this period. One factor driving that reduction has been an approximately 6-fold increase in the price of oil in real terms.

It is not my purpose to analyze how oil intensity was reduced, or the specific role played by the price of oil in that reduction. But I do think the gross picture is instructive. Over four decades, a six-fold increase in the price of oil was associated with a roughly 40% decrease in oil intensity. Of course, the price increase was not steady or predictable—if the price had risen 6-fold in the 1970s and stayed there, oil intensity would surely be somewhat lower than it is today. But on the other hand, a 60% reduction is a lot bigger than a 40% reduction, and the long-run elasticity of demand for carbon has to be lower than the elasticity of demand for oil, since oil is

---

[1] Note that I am talking about stabilization of *global* emissions. Most likely, this would require significant reductions in emissions in industrialized countries in order to make room for increased emissions in currently less developed countries.

one sub-category of carbon, and its intensity can be (and was, to a significant extent) reduced by simply switching to other fossil fuels.

Without trying to establish a precise number or time path, I submit that this experience suggests that to bring about the needed reduction in carbon intensity relying on price-based incentives, the effective or perceived price of carbon would have to increase by more than six-fold. It goes without saying that that won't happen in the near future, but I think it highly unlikely it will happen at all, barring the kind of cataclysmic climate collapse that policy seeks to avoid.

So the challenge is large, and unlikely to met through price-induced changes alone. It will require profound changes in our energy-economic system. We can't say today exactly what form those changes will take, but in terms of magnitude they will have to be comparable to moving an almost entirely carbon-based transportation system to a largely non-carbon base, and largely replacing coal and oil in the generation of electricity (currently about half). I don't know how to make a quantitative comparison, but I submit that digital computation and communication have been improved over the last four decades in a way that is qualitatively comparable to the change we need in our carbon system. And I think the analogy is instructive. We do not calculate or communicate today with improved versions of the instruments that were available for these purposes in 1970. We use a system whose backbone infrastructure and individual components did not exist, and in important aspects were not imagined, in 1970. If we are going to meet the climate challenge, we are going to have to effectuate a comparably broad and deep reconstruction of our energy and industrial systems.

The Economic Case for Technology Policy[2]

This is well-trod ground so I will be brief. There are market failures relating to technological advance that are logically and empirically distinct from the environmental externality addressed by carbon-price policy. Theory says if you have more than one market failure, you need more than one policy instrument.[3]

First and foremost, creators of new knowledge, new products, or new ways of doing things do not reap all of the social returns to these creations, because others can copy and build on them. This "appropriability problem" corresponds to a positive externality that leads to under-investment in technology creation, in the same way that the negative environmental externality leads to over-use of the environment.[4] The positive externality is largest for creations that transform or launch an entire new technology trajectory, along which many parties will earn rewards that are in part dependent on the fundamental development, but will not typically share those rewards.[5] Thus the kind of radical transformative technological change we seek will be underprovided by the marketplace, even in the presence of optimal carbon-price policy.

---

[2] For an extensive discussion of technology and the environment, see Popp, *et al*, 2010.

[3] There is a significant literature modeling the joint implementation of carbon-price and technology policy at a theoretical level. See, e.g. Goulder and Schneider (1999); Popp (2006); Grimaud and Lafforgue (2008).

[4] For more extensive discussion of the issues in this section, see Jaffe (1998).

[5] Indeed, for incremental innovations, the net technology externality can be negative. Depending on market structure and intellectual property rules, the inventor of an incremental improvement on an important existing technology may be able to take the entire market away from the incumbent, earning thereby profits that exceed the incremental value of the improvement.

There is a second positive externality associated with technological change that is less-widely discussed and somewhat more controversial as to its theoretical and empirical significance. There is a large empirical literature on learning-curve effects, whereby the unit cost of production for a good or component falls with its cumulative industry production. One interpretation of this empirical regularity is that there is a base of knowledge about production methods that grows through the process of production itself (Thompson, 2010). Like any other body of knowledge, it has public good properties, including the difficult of excluding parties from its use that may not have contributed to its generation. Those who produce early versions of a technology create a benefit, in terms of knowledge about the production process itself, that they may not be able to capture if for some reason they are displaced from the market and others come later to dominate production, benefitting from reduced costs based on knowledge produced through early production. Early producers of a product may struggle to compete because their production cost makes the product uncompetitive with an older technology; if they received compensation for the learning curve knowledge they produce with each unit they make, this would make such new technologies economically attractive sooner, and hence speed their development and diffusion.

Another aspect of the technology development process that may lead to a socially sub-optimal rate of introduction of new products is the key role played by supplier-user interactions in improving products after they are introduced. Since many of the key improvements in products are made based on feedback from users (Von Hippel, 2010), holders of a potential new technology that cannot anticipate what those improvements will be may fail to bring a potentially valuable product to the market in the first place. This may result in delay or failure in the introduction of products whose eventual social value (once further improved based on user interactions) would be high.

Investments in the development of new technology frequently struggle against imperfections in the market for capital. In particular, information about the potential of a new technology will always be held asymmetrically, making parties that invest in others' efforts at technology development subject to the problems of adverse selection ("lemons") and the winner's curse. These problems raise the cost of capital for financing technology development, and in some cases may make it impossible to secure financing for projects that have a positive expected present value at the appropriate discount rate. (Hall and Lerner, 2010)

Finally, evolutionary models of technological change emphasize the path dependence of technological progress, and the importance of specific transformative events in generating or diverting technological trajectories (Dosi and Nelson, 2010). These transformative events may be exogenous or endogenous, but even if endogenous they are unlikely to respond to economic incentives in any kind of smooth or predictable way, suggesting that carbon-price policy alone is unlikely to bring them forth, certainly not in any way that can claimed to be efficient.

These theoretical arguments for under-investment in technological change are generally confirmed by empirical research. The evidence is that social rates of return to research and development are high, and they are higher than the private rate of return. Probably the most frequently cited study is that of Mansfield and his students, who found for 17 important manufacturing innovations that the median private rate of return was 25 percent while the median social rate of return was 56 percent (Mansfield, et al, 1977). Such case-study analyses suffer from the possibility of selection bias. A survey by Griliches of econometric analyses found "surprisingly uniform" results regarding rates of return: estimates of the private rate of return are generally in the range of 15-30%, with an excess of social returns over the private rate of 50 to 100% of the private rate. Griliches concludes "R&D spillovers are present, their magnitude may be quite large, and social rates of return remain significantly above private rates" (Griliches, 1992, Table 2 and text at page 24).

There is much less empirical evidence as to *which* of the potential market failures discussed above are of greatest quantitative significance. In terms of consequences, knowledge spillovers from research, learning curve spillovers, and knowledge spillovers from customers all suggest that the social rate of return exceeds the private rate of return to investments in new technology. Capital market failures and (possibly) evolutionary models imply that the private rate of return to these investments is greater than the private rate of return to other forms of investment. The numbers cited in the previous paragraph from the Griliches survey suggest that both of these are true—the private return to technology seems to be higher than the private return to other investments, and the social return is even higher. But these results do not effectively distinguish among the different mechanisms that may be contributing to these effects.

Thus one is on firm theoretical and empirical ground in arguing that the private allocation of resources to technological change is far from optimal, and there is at least a clear possibility that public policy intervention could improve things. It is much harder, however, to say which interventions are likely to be most effective, and how far the government should go. At a conceptual level, one can map policy instruments onto particular market failures. Poor appropriability of research can be addressed through government funding of research, and/or subsidies to private research. Learning curve spillovers, and externalities between producers and users of new technologies can be addressed by government purchases of new technologies, subsidies to private parties to purchase new technologies, or regulatory mandates for such purchases. Capital market failures can be addressed through subsidies for the investment activities subject to such failures. It is less clear what specific policies are called for by the evolutionary model.

In the absence, however, of empirical evidence on the relative importance of the different market failures, there is little guidance on the relative value and importance of these different mechanisms. I believe that the best that we can do is to use judgment based on historical experience. I therefore turn now to a review of existing and historical policies that seem relevant.

Public Energy R&D

Figure Two shows Federal Non-Defense obligations for R&D by funding agency over time. We are currently spending about $60 Billion on non-defense R&D. By far the largest share of this goes to health. Historically, the overall total is only slightly higher than it was in the 1960s (in real terms), at which time NASA played the dominant role played today by NIH. Excluding health, there has been a steady and significant decline in Federal research as a share of GDP for decades. There was a slight blip for energy in the late 1970s, but prior to recent increases (not visible in Figure Two), energy R&D has otherwise been pretty constant over the decades, and therefore declining as a share of GDP along with other non-health categories.

Figure Three provides a more detailed look at federal energy R&D. These data come from the IEA, and are not limited to expenditures through the Department of Energy. They show a similar rise and fall around 1980, a gradual increase from about $10 billion/year in the 1990s to almost $15 billion in 2008, followed by a jump to about $22 billion in 2009, due to the Obama stimulus package, and then a return to trend in 2010. Total expenditure temporarily more than doubled in 2009, with much of the increase coming in the fossil fuels area (presumably concentrated on various cleaner fossil fuel technologies). Figure Four shows the data for all of the IEA, with the biggest difference being a noticeably larger overall share for nuclear research than in the U.S. Excluding 2009, the U.S. accounts for about a third of the IEA total. While I leave to others the question of whether the 2009 blip was useful counter-cyclical fiscal policy, it almost surely had no meaningful impact as technology policy, illustrating one of the problems with technology policy in "crisis" mode.

As noted above, it is difficult to say how much public support for energy research (or research in general) is "enough." We do know that in the 1960s through 1980s—when total spending as a share of GDP was higher than it is today—the social rate of return to *private* research investment remained very high. Investments in many specific government-sponsored technologies have been shown to have large ex post returns, although others have not. To my knowledge, no one has attempted an overall aggregation of government technology investments to determine some kind of overall return. As discussed below, studies seem to suggest a high return to at least some portions of the Federal investment in health research; since this investment is so large the likelihood of diminishing returns to research in any given area suggests that returns to energy research would be even higher, unless there are specific features of the two sectors that make the social/private gap larger in health. These seems very unlikely to me, even before considering the huge environmental externality associated with climate change. On this basis, I do think that a strong case can be made for a gradual significant increase in Federal energy research.

At a general equilibrium level, an important qualification to the argument for government support of technological progress is the relatively inelastic supply curve for scientists and engineers, particularly in the short run (Goolsbee, 1998). Technical personnel are a key input to all aspects of technological improvement. Unless their supply is increased, government efforts by any mechanism to accelerate technological improvement will be undercut and potentially stymied as increased demand merely bids up the wages of the fixed supply of workers rather than increasing the scale of the activity. While skilled labor markets would presumably adjust in the long run, the "short-run" effects could still last for decades. Hence it is important that any policy effort to accelerate technological advance include a component to ensure that the supply of appropriately trained technical workers be increased. The importance of labor-market adjustments to changes in Federal spending also suggests that rapid changes in the level of spending, such as those attempted in 2009, are likely to be highly inefficient if not actually counterproductive. Both of these points are discussed below in the context of experience with health-related research.

Experience with Technology Policy in Particular Sectors[6]

Manhattan and Apollo Projects

The Manhattan and Apollo projects demonstrate that the government can achieve an ambitious technological objective, where that objective is specific, can be isolated and executed outside the private sector, and cost is no object (Wilbanks, 2011). In today's dollars, the government spent about $28 billion on the Manhattan project, and $140 billion on the Apollo project. But even if the government were willing to spend comparable amounts on non-GHG technology, the relevance of these examples is unclear. What is needed for GHG reduction is not a specific, isolated technological objective, but a deep and pervasive transformation of the energy-economic system. It is possible that a Manhattan/Apollo style effort could be addressed to a specific technological objective that could *contribute* to the transformation needed for climate change, but it is unclear how that objective would be selected, and it is hard to imagine an objective that would be useful that did not incorporate economic criteria—we don't just need solar cells with certain technical parameters; we need cells with certain parameters at or below a given cost. And it

---

[6] In 2009, the National Bureau of Economic Research organized a symposium with the specific focus of drawing lessons in technology policy from other sectors that could be useful in technology policy motivated by climate change. This resulted in a recently published book (Henderson and Newell, 2011), upon which I draw for some of the discussion in this section.

is not at all clear that a Manhattan/Apollo-style project can achieve that kind of objective.

War on Cancer/NIH Budget Doubling from 1998-2007[7]

NIH support of research in the health sciences over the last several decades has had a major impact on the technology trajectory, measurably improving health outcomes and reducing mortality rates from specific diseases. This impact is not surprising, given the magnitude of the investments made. The research investment in health was and is hugely disproportionate, continuously more than half of all federal non-defense research expenditure since the 1990s. Indeed, for the foreseeable future, an attempt to ramp up federal support of research in energy-related science and technology to something resembling the historical commitment to health is hard to imagine.

A key lesson from the NIH approach is the importance of a focus on training of new scientists at the graduate and post-graduate levels, accomplished in part by making graduate and post-graduate training grants an important component of the overall public research funding approach. In this way the supply of scientists has been increased in tandem as the overall research enterprise was expanded, thereby mitigating what otherwise would have been a collision between growing public support and the supply of specialized human capital (Cockburn, Stern and Zausner, 2010).

A second lesson of the health technology experience is that the financing of health *care*, through both third-part insurance and Medicare/Medicaid, also played

---

[7] For an overview of the lessons of the NIH experience for energy technology policy, see Cockburn, Stern and Zausner (2011) in Henderson and Newell (2011).

a crucial role in the technological progress of the sector. If the government had supported the research, but the purchase of the services of the new technologies had been left to patients themselves, the advance would have been much smaller. The structure of energy markets is obviously different, but this experience suggests that some government policy on the purchase side of energy technology markets will be necessary as a complement to government support for the technology itself. These could include explicit government purchase of low-GHG products, subsidies for consumers' purchases (e.g. tax subsidies for low-emission vehicles), or purchase mandates (e.g. renewable energy portfolio standards for regulated utilities).

The final lesson from the NIH experience is the disastrous consequences of the doubling of NIH spending from 1998-2003. This rapid ramp-up in available funds led to a dramatic increase in the commitment of universities and other research institutions to facilities and staff devoted to NIH-funded research. Despite the fact that the leveling off the funding levels should have been predictable, when the growth stopped it resulted in significant over-capacity in the research system, from which we have still not recovered. Although the overall NIH funding level today remains significantly higher than it was in 1998, the funding probabilities for proposals remain extremely low, and scientists report spending large amounts of time preparing many proposals for each that is actually funded. It is clear from this experience that there is significant inertia and large adjustment costs in the research enterprise, so efforts to increase its scale should be gradual and long-term, rather than trying to bring about dramatic increases in short periods of time.

Computers, Communication and Semiconductors[8]

As noted above, the transformation in the social-economic system in recent decades deriving from advances in computation and communications provides a qualitative historical example of the kind of transformation that we need related to the generation of GHG. Therefore, the role played by technological policy in this transformation may be particularly salient.

In addition to supporting research in these areas through DARPA and other agencies, the government played a major role as a purchaser of components and systems, particularly in the early phases of technology development. Purchases were made of products meeting stated technical specifications, sometimes with little regard to cost, sometimes on a competitive basis with respect to cost, but typically with no maximum price set. These purchases helped moved products down the learning curve, eventually allowing civilian versions to be sold competitively. In addition to facilitating cost reduction by building cumulative volume, government procurement competition spurred private investment in R&D aimed at meeting the technical specifications requested by the government (Lichtenberg, 1988). While it is impossible to say what these industries would look like today had there been no government procurement role, there is no doubt that their advance would have been considerably slower.[9]

---

[8] For an overview of the lessons of the semiconductor and computer experience for energy technology policy, see Mowery (2011); for the internet, see Greenstein (2011).

[9] I believe that government procurement for both defense and the space program played a similar role in the evolution of the commercial aircraft industry, but I am not aware of a published study.

Again, the government's role in these markets is different from its role in energy markets, but the analogy does remain clear. The defense and space agencies advanced the development of these technologies by creating purchase programs in which they committed to buying quantities of specific items based on performance characteristics. The government could do the same for low- or zero-GHG emitting vehicles, space-conditioning systems and other investment goods that the government purchases.

Synfuels and Related Projects

There is, of course, a previous history of government technology policy in the energy sector itself. In the 1970s, the government undertook significant efforts aimed at demonstrating the commercial feasibility of a variety of technologies intended to replace petroleum. What distinguished this effort from those in electronics and health sciences were multi-billion dollar investments in commercial-scale implementation projects. These projects probably did develop solutions to some of the technical problems facing these technologies, but they did not result in significant commercially viable outcomes. This is in part due to the collapse in world oil prices in the 1980s, but it appears that the projects were also wastefully implemented, and crowded out rather than complementing private investment in the technologies (Cohen and Noll, 1991).

What Do We Know?

Based on theoretical and empirical research and the sectoral case studies, I summarize what we know as follows:

1. The rate of return to R&D is high. There is no way to know the "optimal" level of R&D related to GHG specifically, but the pre-2009 effort level seems low in proportion to the public importance of the issue. The kind of huge jump we saw in 2009, or even a "double over five years" kind of commitment is almost surely counterproductive, as we saw with the NIH. What we need is a slow but continuous increase.
2. Based on both the general modeling and the NIH experience, it is crucial that research funding be structured in such a way as to increase the supply of specific human capital in parallel with the increase in research support. Training grants at all levels should be an important component of the overall support.
3. Government purchases, or government policy that generates purchases (e.g. utility renewable mandates) are necessary to accelerate the commercialization of potential technologies. Policy should be based on performance characteristics (e.g. GHG emissions per unit of service), not specified technologies, in order to avoid premature lock-in to today's possibilities.
4. Both research and purchase policies should be designed to encourage and be complementary to private development efforts, to try to minimize crowding out of private development investment. Procurement competitions or public financial support of private-market purchases appear superior to government-financed demonstration projects.
5. "Success" will almost surely require technologies we don't know about today.
6. Nothing should be "off the table." Even with the best policy, the chance of failure is high enough that we need all options, including possibilities like geoengineering that appear highly problematic in various ways.

What We Don't Know (and How We Could Learn)

What we don't know is much at all about the relative effectiveness of different policy approaches and mechanism designs. Just as an example, if we want to use public funds to increase research, we can (1) hire a bunch of people and put them to work in a National Lab; (2) put out an RFP and then fund non-governmental researchers based on peer review of research proposals; or (3) offer a prize or prizes for research that achieves pre-specified scientific or technical milestones. Arguments can be made in favor and against each of these proposals, but we have made little effort to test empirically which works better in what circumstances.

Similar diverse policy options exist for other policy sub-objectives, such as increasing diffusion of low-GHG products in use.

We do have some *ex post* evaluations of various policy initiatives. But such *ex post* evaluations typically suffer from serious limitations. First, they are generally sponsored by the agency that undertook the initiative, and even the best-intentioned study of this kind is subject to bias. More fundamentally, in an evaluation based only on observations after the fact, it is impossible to know what the "but for" world would have looked like. If you fund a bunch of research, and good things come out of it, you can feel good—but without knowing what results would have been forthcoming in the absence of public funding, you don't really know anything about the effectiveness of your research funding program (Jaffe, 2002). Or, to come from the other side, if you provide loan guarantees to firms in the initial stages of commercial development of a GHG-reducing technology, and one firm goes bankrupt, does this mean the program was poorly conceived or designed?[10] After all, the reason loan guarantees are needed is that this is a highly uncertain business with a poorly functioning capital market; presumably your expectation going in is that some of these projects will fail and others will make great strides. To know if the failure rate is "too high" requires a standard for comparison.

The solution to this problem is to build program evaluation into the design of all government technology programs up front (Jaffe, 2002). This means, first, that

---

[10] DOE modifications of its loan guarantee program in light of a prominent default were in discussion as this paper went to press, see e.g. http://www.forbes.com/sites/toddwoody/2011/09/23/doe-rescinds-solar-loan-guarantees-in-wake-of-solyndra-bankruptcy/

instead of trying to decide now which approach is best, we should consciously fund a number of different approaches so that we can test their relative efficacy. Then, it is crucial that the data that will eventually be necessary to do a meaningful assessment of the different approaches *ex post* should be identified *ex ante*, and collected as the program unfolds. Importantly, this includes background information on the projects being funded, as well as on those that are not funded; possession of this kind of information allows the *ex post* evaluator to use econometric methods to construct pseudo but-for comparisons. The cost of doing the evaluation research should be considered part of the cost of the program, and should be built into the program budget.

I have emphasized that mitigating climate change is a very hard problem, but it is also a very long-term problem. If we structured all of our programmatic efforts over the next ten years in such a way as to yield meaningful data on the relative effectiveness of different approaches, we would still have decades in which to concentrate our resources in the approaches that prove to work best.

To end with the words of JFK's inaugural:

> All this will not be finished in the first one hundred days. Nor will it be finished in the first one thousand days; nor in the life of this Administration; nor even perhaps in our lifetime on this planet. But let us begin.

References

BP Statistical Review of World Energy 2010, http://www.bp.com/productlanding.do?categoryId=6929&contentId=7044622

Cohen, Linda R. and Roger G. Noll, 1991. The Technology Pork Barrel, Brookings Press

Cockburn, Iain, Scott Stern and Jack Zausner, 2011. "Finding the Endless Frontier: Lessons from the Life Sciences Innovation System for Energy R&D," in Henderson and Newell, *op cit*

Dosi, Giovanni and Richard R. Nelson, 2010. "Technological change and industrial dynamics as evolutionary processes," in Hall and Rosenberg, *op cit*

Goolsbee, Austin, 1998. "Does government R&D policy mainly benefit scientists and engineers?", *American Economic Review*

Goulder, Lawrence H. and Stephen H. Schneider, 1999. "Induced technological change and the attractiveness of CO2 abatement policies," *Resource and Energy Economics* 21 pp. 211-253.

Grimaud, Andre and Gilles Lafforgue, 2008. "Climate change mitigation policies: Are R&D subsidies preferable to a carbon tax?" *Revue d Economie Politique* 118:6 pp. 915-940 .

Griliches, Zvi, 1992. "The search for R&D spillovers," *The Scandinavian Journal of Economics*

Hall, Bronwyn H. and Josh Lerner, "The financing of R&D and innovation," in Hall and Rosenberg, *op cit*

Hall, Bronwyn H. and Nathan Rosenberg, eds., Handbook of the Economics of Innovation, North Holland

Henderson, Rebecca and Richard Newell, eds., 2011. Accelerating Energy Innovation: Insights from Multiple Sectors, University of Chicago Press

Jaffe, Adam B., 1998. "The Importance of 'Spillovers' in the Policy Mission of the Advanced Technology Program," *Journal of Technology Transfer***,** Summer

Jaffe, Adam B., 2002. "Building Programme Evaluation into the Design of Public Research-Support Programmes," *Oxford Review of Economic Policy*

Lichtenberg, Frank, 1988. “The Private R and D Investment Response to Federal Design and Technical Competitions,” *American Economic Review* 78:3, pp 550-559.

Mansfield, E., J.Rapoport, A.Romeo, S.Wagner and C.Beardsley, 1977. "Social and Private Rates of Return from Industrial Innovations," *Quarterly Journal of Economics* 77:221-240

Mowery, David C., 2011. “Federal Policy and the Development of Semiconductors, Computer Hardware, and Computer Software: A Policy Model for Climate-Change R&D?” in Henderson and Newell, *op cit*

Popp, David, 2006. "Comparison of climate policies in the ENTICE-BR model," *Energy Journal* Special Issue 1, pp. 163-174.

Popp, David, Richard Newell and Adam Jaffe, 2010. “Energy, the environment, and technological change,” in Hall and Rosenberg, *op cit*

Thompson, Peter, 2010. “Learning by doing,” in Hall and Rosenberg

US Dept of Agriculture, World GDP Historical Data, http://www.ers.usda.gov/Data/Macroeconomics/Data/HistoricalRealGDPValues.xls

U.S. National Science Foundation, “Federal obligations for total research and development, by major agency and performer”, Multiple tables: 1951- 2002; 2001-2003; 2004-2006; 2007-2009.” Science and Engineering Statistics, Federal Funds for R&D, http://www.nsf.gov/statistics/fedfunds/

Von Hippel, Eric, 2010. “Democratizing innovation: The evolving phenomenon of user innovation,” in Hall and Rosenberg, *op cit*

Wilbanks, Thomas J, 2011. “Inducing Transformational Energy Technological Change”, *Energy Economics* 33:4

Figure One

Historical Oil Intensity of the World Economy

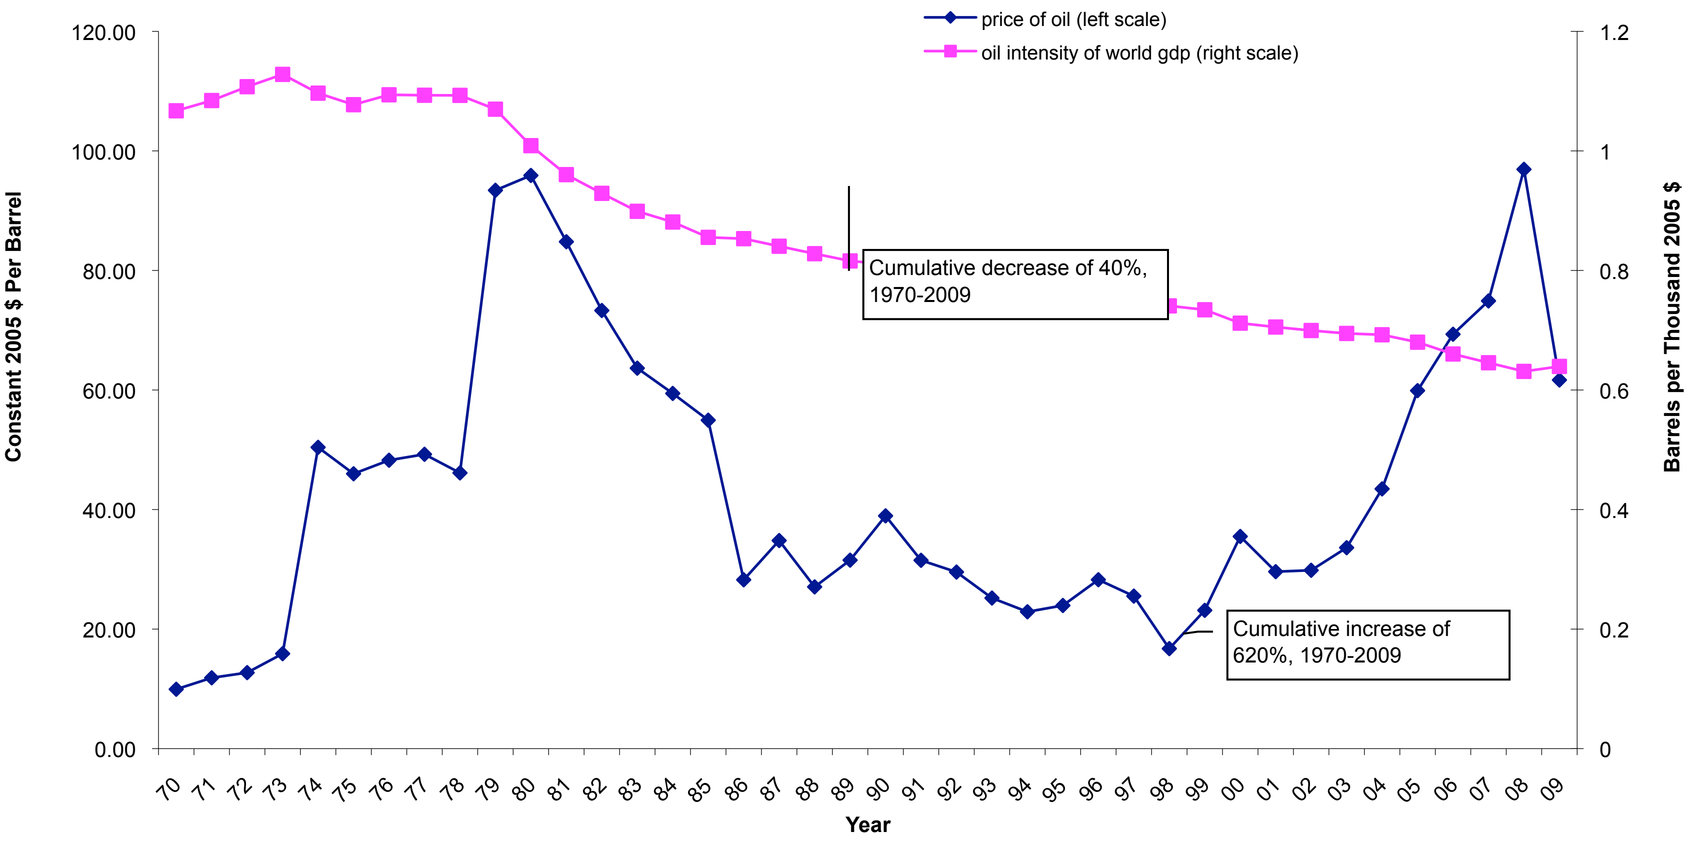

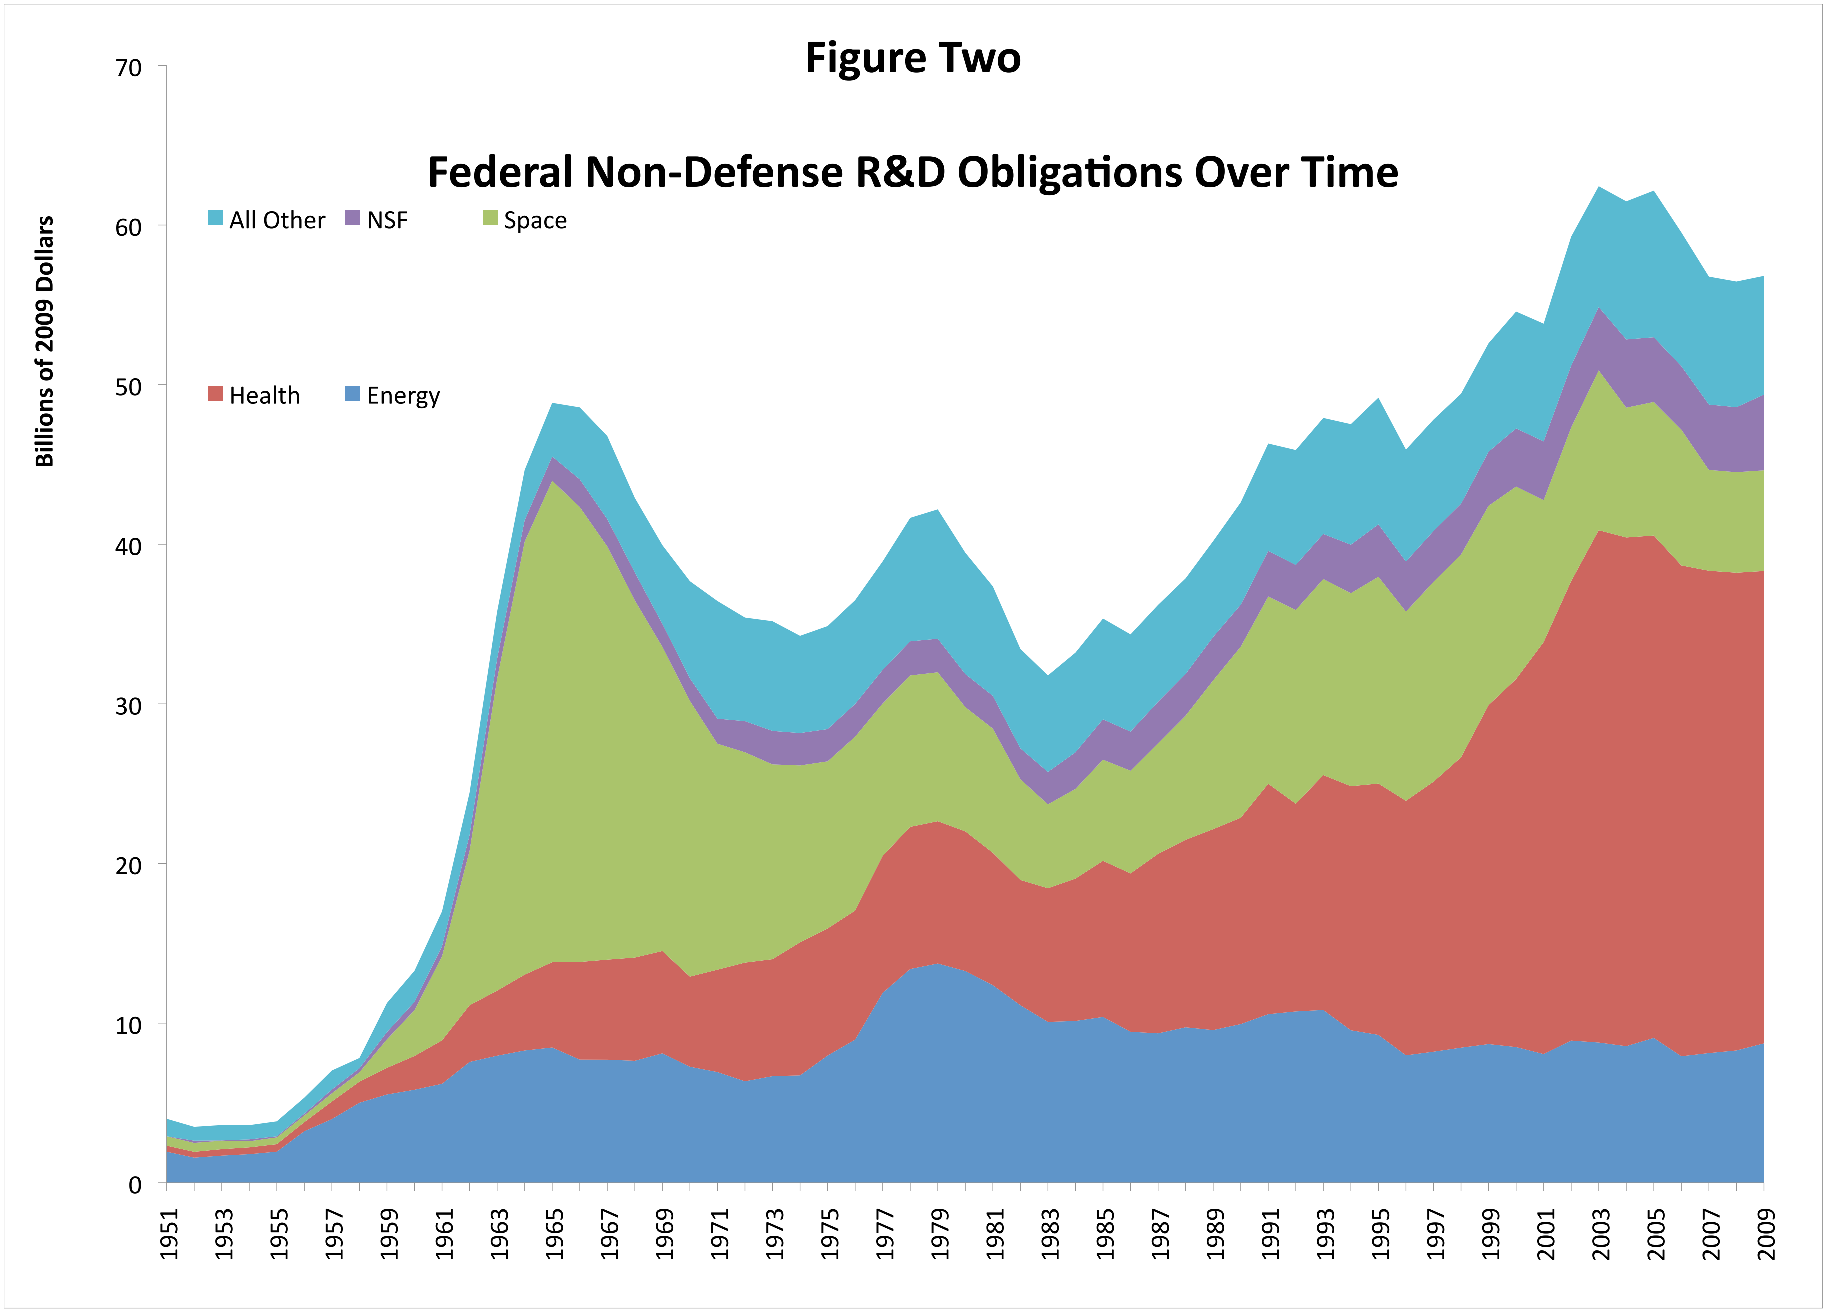

Figure Two

Federal Non-Defense R&D Obligations Over Time

Figure Three

U.S. Energy R&D

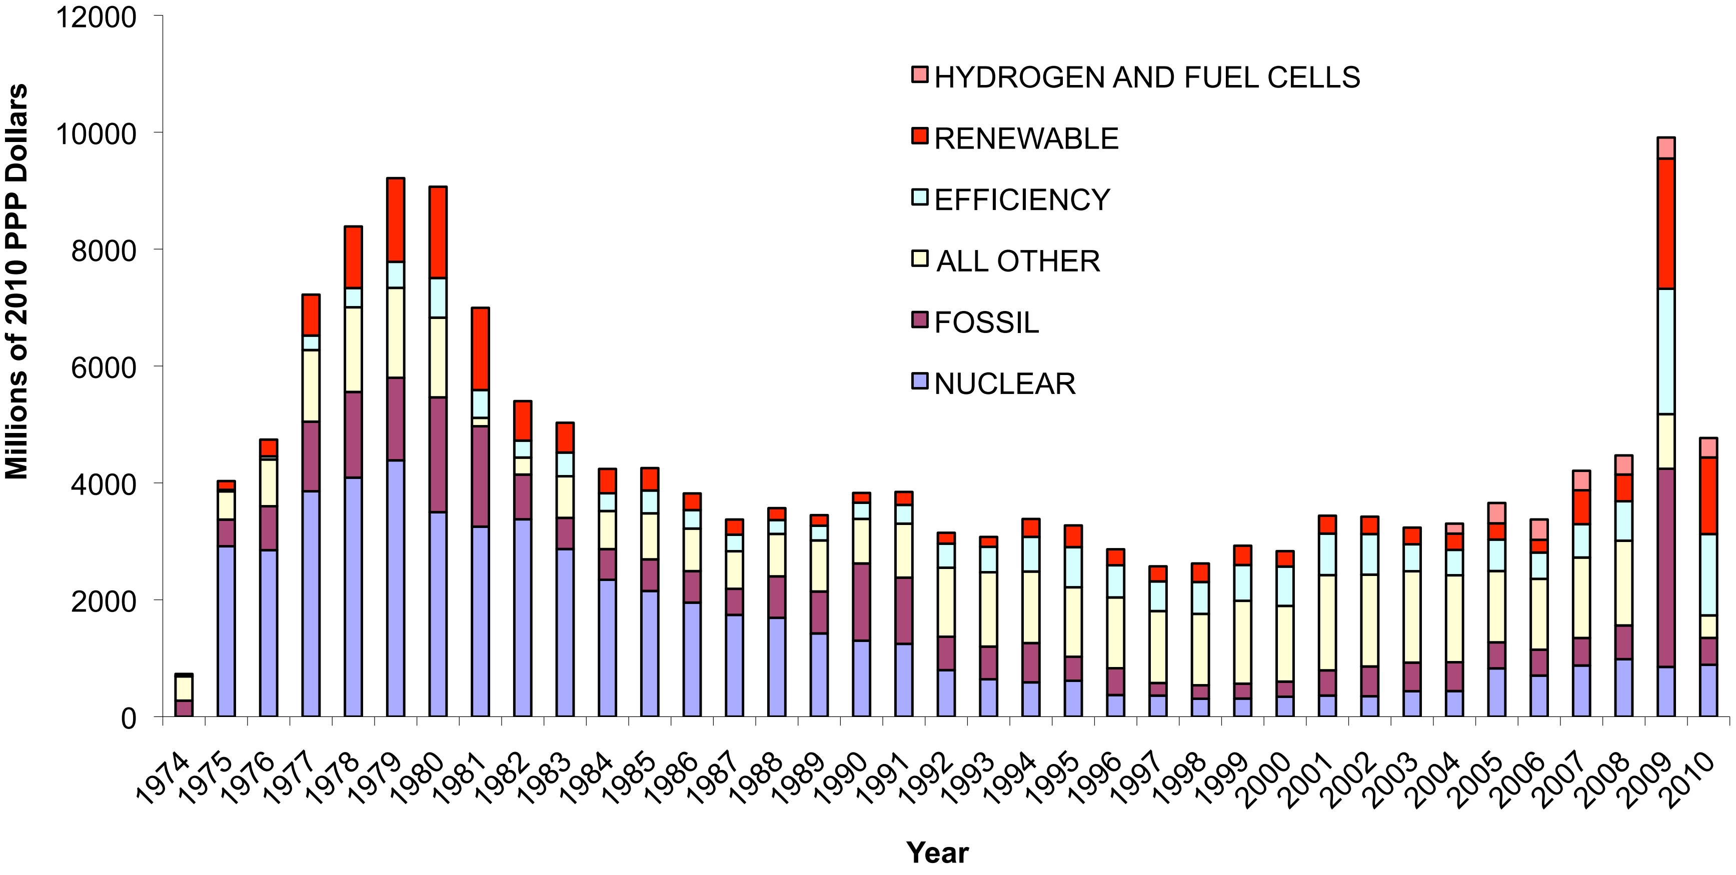

Figure Four

IEA Countries Energy R&D