LBNL – ####

ERNEST ORLANDO LAWRENCE
BERKELEY NATIONAL LABORATORY

Accounting for Innovation in Energy Efficiency Regulation

Prepared by
K. Sydny Fujita and Margaret Taylor

For the European Council for an Energy Efficient Economy

March 11, 2013

This work was supported by the Assistant Secretary for Energy Efficiency and Renewable Energy, Office of Building Technology, State, and Community Programs, of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

Disclaimer

This document was prepared as an account of work sponsored by the United States Government. While this document is believed to contain correct information, neither the United States Government nor any agency thereof, nor The Regents of the University of California, nor any of their employees, makes any warranty, express or implied, or assumes any legal responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by its trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof, or The Regents of the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof, or The Regents of the University of California.

Ernest Orlando Lawrence Berkeley National Laboratory is an equal opportunity employer.

Accounting for Innovation in Energy Efficiency Regulation

Margaret Taylor
Lawrence Berkeley National Laboratory
Environmental Energy Technologies Division
Mail Stop 90R2002 Bldg. 90, Room 2024K
Berkeley, CA 94720
Email: mtaylor@lbl.gov

Sydny Fujita
Lawrence Berkeley National Laboratory
Environmental Energy Technologies Division
One Cyclotron Road
Mail Stop 904000
Berkeley, CA 94720
Email: ktfujita@lbl.gov

Abstract

U.S. federal energy efficiency standards are, by law, set at the maximum level of energy efficiency that is technically feasible and economically justified, as gauged through regulatory impact analysis (RIA). RIA has been an integral part of the government policy-making process for over 20 years in the U.S., but in February 2011, the main government oversight agency issued guidance to all agency heads for the first time that the "best-available techniques" for RIA include those that identify "changing future compliance costs that might result from technological innovation or anticipated behavioral changes." This paper explores the question of what techniques are currently in use by regulators to account for innovation in RIA, and what makes those available techniques "best." The paper focuses on two major examples of the use of learning curves in regulatory cost estimates. The first was pioneered by the Environmental Protection Agency (EPA) in vehicle emissions control at least as early as the late 1970s, and later adopted for vehicle fuel economy regulation by the National Highway Safety Administration (NHTSA) in the mid-2000s. The second was implemented by the Department of Energy (DOE) in 2011 to help set appliance energy efficiency standards. The paper: (1) provides an overview of some of the major findings of the academic literature on learning curves in order to inform an assessment of a "best" approach to a learning curve-based RIA cost adjustment technique; (2) describes the EPA-NHTSA and DOE approaches to this technique; and (3) assesses these approaches against the criteria of alignment with economic theory and of administrative sustainability (i.e., fit with existing laws and institutional arrangements, including relationships between regulators and regulated industries).

Introduction

Regulatory impact analysis (RIA) plays an important role in the setting of U.S. federal minimum efficiency performance standards (MEPS), as it does in many other areas of U.S. regulation. RIA is a tool to promote good governance that is formally required in a growing number of nations (see Figure 1; for more information, see OECD (2009)) in order to help governments weigh the costs and benefits of proposed regulatory actions, particularly as they compare to those of alternative actions and other government priorities (e.g., fiscal health, small business support, the operation of free markets, the welfare of various economic and social groups, etc.). RIAs promise government transparency and regulatory accountability. To fulfill this promise, the Organization for Economic Cooperation and Development (OECD) has established a number of "best practices" for the conduct of RIAs. These practices include: the use of transparent and consistent data, assumptions, and models; the assessment of an appropriate range of alternatives to the proposed intervention, including clear and consistent baseline assumptions for the world without the intervention; consideration for discounting the future; sensitivity analysis; attention to non-monetizable/non-quantifiable aspects of a policy and its effects; and estimates of benefits and costs that are grounded in economic theory (Harrington, Heinzerling et al. 2009). Such practices guide the RIA approach employed to assess, for each proposed MEPS for a regulated product, the legally

required maximum level of energy efficiency that is technically feasible and economically justified, as well as the impacts on consumers and on national energy savings of a proposed MEPS. The grounding of cost estimates in economic theory is perhaps one of the trickier of the OECD best practices for RIAs, as economic theory is complex and can change over time.

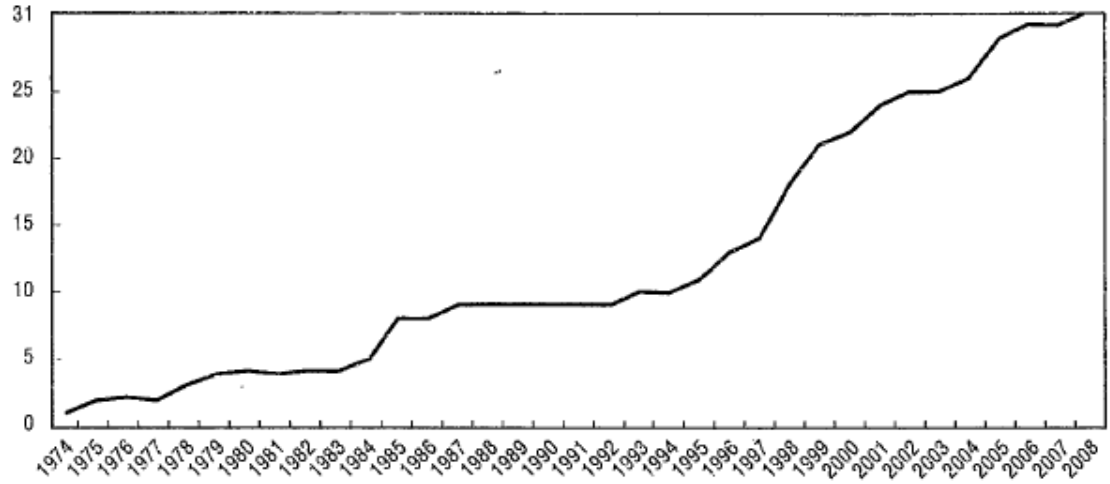

Figure 1: Number of OECD jurisdictions with a formal requirement for regulatory impact analysis. Source: (OECD 2009)

One aspect of economic theory that appears to have been overlooked in designing many RIA frameworks, for example, is how to think about the compliance cost changes that should be expected over the time a regulation is in effect due to factors related to technological innovation. Traditionally, many U.S. RIAs have taken current estimates of the costs of technologies that comply with a proposed regulation and projected that those costs would remain unchanged over the full regulatory time frame. This would appear to conflict with economic theory, as changes in technology costs can be expected to result from many factors over time, including the successful completion of research and development (R&D) projects, economies of scale, exogenous technological change, direct and indirect labor learning effects, and others. There are a number of indications in the literature that ignoring technological change is an important contributor to the often-observed tendency for RIA cost projections to seriously over-estimate the costs of compliance with new environmental, health, safety, and energy efficiency regulations (see, e.g., Harrington, Morgenstern et al. 2000; Hwang and Peak 2006; Dale, Antinori et al. 2009; Taylor 2012).[1] These indications have helped spur recent influential calls for the reform of RIAs, which have described the appropriate modeling of technological change as an element of "an ideal cost-benefit analysis," although with the expressed caveat that this modeling is "at the frontier of economic research" (Harrington, Heinzerling et al. 2009).

In February 2011, this issue was raised by the U.S. government agency tasked with establishing and enforcing best practices in RIAs – the Office of Information and Regulatory Affairs (OIRA) in the Executive Office of Management and Budget – when it provided guidance on presidential Executive Order 13563, the latest in the series of Executive Orders that have institutionalized the principle that the social benefits of new federal regulations should exceed their private costs. At that time, OIRA called for federal agencies "to use the best available techniques to quantify anticipated present and future benefits and costs as accurately as possible," which include techniques that identify "changing future compliance costs that might result from technological innovation or anticipated behavioral changes." (OMB 2011). But OIRA, the OECD, and the academic literature do not currently provide detailed guidance on these "best available techniques" for accounting for innovation in RIAs.

This paper seeks to identify the best technique that is currently available to incorporate technological innovation into the compliance cost estimates for new regulations. We reason that two criteria must inform the assessment

[1] Harrington et al. (2000) is the most comprehensive of these studies. The paper reviews ex ante analyses of 28 regulations across a variety of fields, finding that approximately half overestimate unit cost and total cost by more than 25%. Note that both exogenous and induced technological change (the latter due to increased investments directed toward compliance) are frequently cited in the literature as relevant to RIA cost overestimates. Other explanations include lack of commitment by firms (in providing accurate cost estimates for possible future regulation) as well as errors by regulators who have asymmetric information regarding firm resources and dynamic capabilities.

of a “best” technique. First, it must be grounded in economic theory, in keeping with OECD best practices. Second, it must be administratively sustainable (i.e., it must fit with existing legal and institutional precedent, including relationships between regulators and regulated industries). Meanwhile, we find that only one technique can currently be considered to be widely available: learning curve-based adjustments to projected costs. The Environmental Protection Agency (EPA), the National Highway Transportation Safety Administration (NHTSA), and the DOE have all incorporated learning curve-based cost adjustments in numerous RIAs. The key difference between the approaches used by these agencies is the degree of their focus on the cost of the components that enable compliance with a regulation, as opposed to the cost of the full product that has to meet the new standard. The EPA and NHTSA take a component-based learning curve approach (referred to here as the “EPA-NHTSA” approach), while the DOE takes a full product -based approach.

The paper is divided into four major sections. In the second section, we help to define the economic theory that an RIA learning curve technique should align with by providing an overview of some of the major findings of the academic literature on learning curves. In the third section, we describe the learning curve technique as it has been used to date in the EPA-NHTSA approach and in the DOE approach. In the fourth section, we provide recommendations for the ongoing practice of RIAs and some future research directions.

Background on Learning Curves

The focus of this section is on understanding the economic theory underlying the learning curve phenomenon that agencies are interested in marshalling in order to adjust their RIA cost estimates. It should be stated up-front, however, that several disciplines have contributed to the crafting of the current economic understanding of learning curves. Learning curves originally emerged from individual learning and psychology studies in which it was repeatedly demonstrated that the more times a task is performed by an individual, the less time is required for the individual to perform each subsequent iteration of that task. The concept was later extended to organizations and industries in studies focused on productivity in manufacturing, such as the production of Liberty Ships and aircraft (see, e.g., Wright 1936; Alchian 1963; Rapping 1965). In these and later management studies, economists, psychologists, and sociologists in the field of organizational behavior united behind the empirical observation that as more units of a good are produced, the less it costs to manufacture each subsequent unit. Note that the initial focus of the voluminous management literature on learning curves was on documenting their existence, magnitude, and shape in many different empirical settings. Over the past two decades, however, the management literature has focused primarily on explaining the factors that underlie the learning curve, in an effort to make the phenomenon of more practical use to managers. Meanwhile, the energy policy literature adopted the learning curve concept at least as far back as the early 1990s, initially in order to assist in modeling technological change endogenously in long-term integrated assessment models of global climate change. Beyond this application, the energy policy literature has also worked to empirically describe learning curves in energy supply and demand technologies, with a particular focus on their implications for the costs of emerging renewable energy technologies. Note that as the energy policy literature on learning curves tends to focus on determining functional form for use in modeling, rather than on the economic factors underlying the learning curve phenomenon, this literature is not a primary focus of this review.

Learning Curves in the Academic Literature

At its heart, a “learning curve” is a mathematical representation of the concept that with continuous production, production knowledge and experience grow and result in economies. The most commonly used and most robust functional form of an organizational learning curve is the power function presented in Equation 1 (Kantor and Zangwill 1991). The dependent variable here is unit costs, while the independent variable here, as it is in most learning curve formulations, is the cumulative number of units produced, which is a proxy for knowledge acquired through production until the current period (this is sometimes expressed as lagged cumulative output, since the current period’s output does not go into the accumulated knowledge stock). Note that it is important to have a sense of how to bound the count of units (i.e., understand which unit is the first unit, both temporally and spatially). In interpreting this equation, it is understood that if unit costs decrease as a function of the cumulative number of units produced, *ceteris paribus*, organizational learning is said to occur (Argote and Epple 1990).

Equation 1: The Traditional Functional Form of an Organizational Learning Curve

$$C(X) = C_0 * X^{-b}$$

where:

$C(X)$ = the cost of the Xth item made
C_0 = the cost of the first item made
X = the cumulative number of units produced through the previous period

b = a parameter that measures the rate costs are reduced as cumulative output increases

It is simpler for estimation purposes if the power function in Equation 1 is expressed on a logarithmic scale so that the data more closely resemble a straight line. The general form of that logarithmic equation is expressed in Equation 2, although additional predictor variables – already in log form – are usually included in this equation (Levin 2000) If the coefficient of Equation 2 is statistically significant when the equation is estimated with appropriate control variables, learning is said to occur.

Equation 2: The Estimation Form of the Traditional Organizational Learning Curve

$$\ln C(X) = \ln C_0 - b \ln X$$

The effects of the learning curve are generally discussed by using the language of Equations 3 and 4. Equation 3 defines the progress ratio (PR) as the percentage of the initial unit cost that is left after cumulative output doubles, while Equation 4 defines the learning rate (LR) as the percentage reduction in unit cost that is associated with each doubling of cumulative output. For example, a progress ratio of 80% means that each doubling of cumulative output leads to a learning rate of 20% reduction in unit cost (with the value for b in Equation 1 of 0.322).

Equation 3: The Progress Ratio

$$PR = PR = 2^{-b}$$

Equation 4: The Learning Rate

$$LR = 1 - 2^{-b}$$

Improvement correlated with production experience by organizations has been documented in myriad empirical studies over the past 75-plus years, with improvement typically measured by the unit cost of production, but sometimes by other performance or quality indicators, such as reduction in complaints or industrial accidents. Variables are not interchangeable in derivations of "learning curves," however. For example, calendar time is not as good a predictor of an organization's improvement with experience as is its own cumulative output (Rapping 1965; Lieberman 1987; Argote and Epple 1990). Similarly, using price as an outcome metric can be misleading. Whereas cost changes occur over time due to such things as changes in input prices and production efficiency, price changes can occur from several other factors, including joint determination with output, which can lead to learning curves that are statistically unidentified, without enough observations. To be most useful in learning curve estimation, price-cost margins need to be constant over time. This, unfortunately, requires a number of unlikely things to stay constant in an industry, such as the number of producers, the elasticity of demand, and the return on equity for producers. For a useful discussion, see Papineau (2006).

Organizations vary considerably in their rates of improvement (see, e.g., Yelle 1979; Hayes and Clark 1986), as documented most famously in Dutton and Thomas (1984), which reviewed the progress ratios (i.e., the percentage of the initial unit cost found with each doubling of cumulative output) observed in 22 field studies of 108 manufactured items in a wide range of industries. In the frequency distribution of progress ratios illustrated by Dutton and Thomas (1984) and displayed in Figure 2, the average progress ratio is 81-82% (i.e., an 18-19% unit cost reduction per doubling of cumulative output), the maximum progress ratio is 55% (i.e., a 45% improvement per doubling), the minimum positive progress ratio is 95% (i.e., a 5% improvement per doubling), and in one instance, a negative progress ratio of 107% is observed (i.e., a 7% unit cost *increase* per doubling). Similar results have been found in recent reviews of the Dutton and Thomas (1984) data as well as many energy supply and demand learning curve studies, which reveal within-technology variation in progress ratios often as large as between-technology variation, but approximately normally distributed around a mean of 82 ± 9% (i.e., an 18% improvement per doubling; see Ferioli, Schoots et al. 2009; Weiss, Junginger et al. 2010).

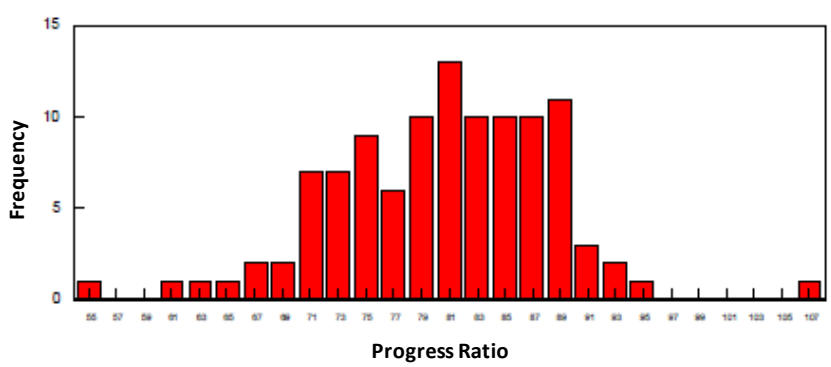

Figure 2: Distribution of progress ratios documented in the studies reviewed in Dutton and Thomas (1984). Source: Dutton and Thomas (1984)

The original learning curve concept – as production continues, production knowledge and experience grow and result in economies – was influenced by experimental psychology research which typically finds that the more times an individual performs a task, the better that individual is able to perform each subsequent task (Argote and Epple 1990). In the organizational setting, however, individual learning by workers performing tasks (i.e., direct and indirect labor learning in the context of a given set of capital goods) is often only one of several factors that help to explain progress ratios. For example, labor learning was only one of the four main categories of factors Dutton and Thomas (1984) identified as explaining, either singly or in combination, the progress ratios in the studies they reviewed. The others were: (1) effects of technological change in the production environment due to capital investment, which contribute to an organization's knowledge and experience (see, e.g., Arrow 1962; Sheshinski 1967; Joskow and Rozanski 1979); (2) local industry and firm characteristics, particularly with regard to the degree of mechanization, the distinction between assembly and machining, the length of cycle times, and the type of manufacturing process (see, e.g., Hirsch 1952; Hirsch 1956; Adler and Clark 1991); and (3) effects which result from changing production techniques in order to absorb indivisible costs and exploit other economies that emerge as an organization anticipates increased production (note that because of the aggregate nature of the "learning curve," there have long been instances of cost reductions due to scale economies being misattributed as the effects of organizational knowledge or experience – see, e.g., Conway & Schultz (1959), Wright (1936) – which is why Dutton and Thomas (1984) includes scale effects as a factor in explaining progress ratios).

Beyond those factors discussed in Dutton and Thomas (1984), factors that explain variation in progress ratios include: organizational forgetting (i.e., knowledge depreciation; Argote and Epple (1990)); employee turnover (ibid.); transfer of knowledge across products, business lines, and organizations (ibid.);[2] level of interdependence between organizational units (e.g., through vertical integration, see Sorenson 2003); management practices such as debriefing activities, coaching behavior, use of formal procedures, and use of cross-functional communication (Pisano, Bohmer et al. 2001); and volatility in the competitive environment (Sorenson 2003).

To simplify this list of factors, Dutton and Thomas (1984) laid the foundations for a long-standing approach to distinguishing between factors that can be "induced" by management versus developed "autonomously" by labor, with both sets influenced by endogenous as well as exogenous sources of innovation (see, e.g., von Hippel 1976). Induced learning (also known as second-order learning – see Adler and Clark (1991)) stems from the formal knowledge generation and transfer practices of management (e.g., R&D investments; employee training programs; organizational practices, procedures, and rules; etc.), while autonomous learning – which is most commonly referred to by the labels of learning-by-doing, learning-by-experience, first-order learning, etc. – stems from informal experiments and process refinements by employees in order to execute tasks more cost-efficiently (see Adler and Clark 1991; Wiersma 2007). Autonomous learning often occurs in stable organizational contexts in which the opportunity to develop experience is present, while induced learning tends to occur when firms redefine their strategy, often in response to competitive opportunities and threats (Fiol and Lyles 1985). Argote et al. (2003), meanwhile, suggests that the list of explanatory factors can be simplified by focusing on the degree to which employees and/or organizations have the ability, the opportunity, and the motivation to learn in a given context.

Note that most, but not all, of the academic literature on learning curves since the early 1990s focuses on explaining observed variation in firms that adopt a new technology, start producing new products, or open new plants or production lines (for a discussion, see Wiersma (2007); useful references include Adler and Clark (1991), Argote et al. (1990), Pisano et al. (2001), Sorenson (2003), and Reagans (2005)). Firms that provide established products and/or use more mature technologies can still maintain positive learning curves, however, as well as experience occasional sudden reductions in average cost if, for example, a new technology is implemented (Wiersma 2007).

Considerations for the Predictive Use of Learning Curves

Managers and policy-makers have long been interested in marshalling the learning curve as a practical tool, rather than simply a descriptive one. The learning curve has been used in a variety of management settings, informing organizations in their efforts to: formulate manufacturing strategy, develop production schedules, establish pricing and marketing approaches, direct employee training, predict the costs confronting competitors, and decide on the utility of and approach to subcontracting production Argote and Epple (1990).

[2] There is often more variation across organizations or organizational units producing the same product than within organizations producing different products Argote and Epple (1990).

The implied framework for application is to use past experience to predict future progress ratios that can be used for various purposes, including estimating the improvement possible with the manipulation of cumulative output. But in order to predict progress ratios accurately based on the past requires the progress ratio to be reliable – i.e. subject to the same variation over time and space – and consistently explainable (Dutton and Thomas 1984). Given the very long list of explanatory factors underlying learning curves, it is perhaps not surprising that large errors have been observed when attempting to project future progress ratios based either on an industry's historic progress ratio (derived from industry-wide learning curves, which are typically called "experience curves") or on a firm's own progress ratio. In a very helpful discussion of this topic in Dutton and Thomas (1984), the authors point out that Alchian (1963) found mean prediction errors in both cases of 22-25%, and other studies echoed this unreliability (see, e.g., Hirsch 1952; Hirsch 1956; Conway and Schultz 1959; Billon 1966). Indeed, firms that have used the "learning curve" concept for planning purposes have been found to achieve smaller-than-expected profits (see, e.g., Porter 1980; Kieschel 1981).

Thus, although the "learning curve" effect is widely observed, it cannot be used for prediction without caution. Given the standard deviation of about 9%-points around the average progress ratio, Weiss (2010) warns that projecting beyond three or four doublings of cumulative production results in substantial under- or over-estimation of production costs. Accuracy of estimation is also influenced by the way that doubling of cumulative production is calculated – i.e., by organizational unit, organization, industry, nation, or region (see Nemet 2009); note that these units of analysis also appear to matter to the recurring, but currently unpredictable, phenomenon of knowledge depreciation (for more on the phenomenon of organizational forgetting, see Argote 1999). An additional caution for prediction on the basis of learning curves is based on the evidence that the degree of maturity of a product or a production technique makes a difference to the shape of the learning curve. When a technology is new, there are significant degrees of freedom in redesigning processes to make them more efficient; the learning curve at this stage is often steep (Wiersma 2007). When a technology is mature and its production is fixed, however, autonomous learning tends to dominate induced learning, and at this stage, the curve generally starts to level out such that additional cumulative output has less and less impact on the improvement metric (see Adler and Clark 1991; Wiersma 2007). Older studies within single manufacturing facilities have found learning curves to become essentially flat after approximately two years (Baloff 1966; Hall and Howell 1985).

Use of Learning Curves in Regulatory Impact Analyses

The Environmental Protection Agency (EPA), the Department of Transportation's National Highway Transportation Safety Administration (NHTSA), and the Department of Energy (DOE) have all incorporated learning curves into RIAs for various policies. The EPA pioneered the technique for all U.S. RIAs in its cost estimation of vehicle emissions standards for traditional air pollutants (e.g., nitrogen oxides, carbon monoxide, particulates, etc.), dating back at least as far as 1978 (see Lindgren 1978). The approach was formalized by EPA throughout the 1990s, and then generally followed – with some modifications – by NHTSA as it developed RIAs for Corporate Average Fuel Economy (CAFE) standards in the late 2000s that were designed to partner with EPA efforts to regulate vehicle greenhouse gas emissions. We use the term "the EPA-NHTSA approach" in order to highlight the similarities in the techniques used by these agencies. The DOE application of learning curves to its RIAs for appliance MEPS began in 2011 under the Obama administration and takes a very different approach. This section delves into some of the details of the EPA-NHTSA and DOE approaches, providing for each approach: the general RIA approach of the agency and a description of the general framing, derivation, and application of the approach.

Use of Learning Curves in Vehicle Regulations

This section summarizes the EPA-NHTSA approach to learning curves in RIAs as it has been formally incorporated into RIAs for more than 15 years in the area of vehicle emissions (clearly documented since at least 1997, although the practice dates back at least to 1978) and fuel economy regulation (clearly documented since 2008) (see, e.g., Lindgren 1978; E.H. Pechan & Associates 1994; U.S. Environmental Protection Agency 1997). The EPA-NHTSA approach is essentially a component-based learning curve approach (for more information, see Kantor and Zangwill 1991; Ferioli, Schoots et al. 2009), in that it focuses on using a learning rate(s) to appropriately adjust the cost of the various components that enable compliance with a new regulation, rather than adjust the cost of the full product as a technological system that is more-or-less a "black box."

General EPA-NHTSA RIA Approach

EPA's traditional RIA approach to assessing the economic impacts of its proposed regulations coordinates well with a component-based learning curve cost adjustment technique.[3] In its RIA framework, the EPA defines a set of features that can alter a vehicle's emissions (or fuel economy, in the case of NHTSA) in order to meet the

[3] Unless otherwise specified, the RIA framework described here for EPA applies to NHTSA as well.

proposed standard (i.e., a "technology package"). Cost projections are estimated for each component in the technology package based on product tear down studies,[4] such as those used by manufacturers to benchmark their products against competitors, as well as from confidential data either submitted by vehicle manufacturers or obtained through individual meetings with major original equipment manufacturers (OEMs) that account for most (~90%) of the vehicles produced for sale in the U.S. Besides estimating piece costs for all candidate technologies, the EPA estimates direct manufacturing costs and cost markups to account for manufacturers' indirect costs, considering known manufacturer practices, such as making major changes to model technology packages during a planned redesign cycle. Note that EPA models a baseline U.S. fleet of vehicles using estimates of current production volume based on EPA's vehicle emissions certification data, data purchased from Ward's Automotive Group, and CAFE certification data (U.S. Environmental Protection Agency and U.S. Department of Transportation 2011).

The cost difference between technology packages depends on the types of components included, the maturity of those components, the cost of their materials, and the cost of associated labor. Costs are categorized as either direct or indirect, long- or short-term. Table 1 provides a few examples of the types of components that can be included in a technology package and their estimated costs.

Table 1: Example Component Costs ($ per vehicle)

Technology	Relevant System	Small Car	Large Car
Cylinder deactivation	Base engine	--	$203
Camless valvetrain (electromagnetic)	Base engine	$336-673	$336-673
Diesel – Lean NOx trap	Base gasoline engine	$2790	--
Optimized E20-E30	Base gasoline engine	$713	$143

Source: EPA (2008)

EPA-NHTSA Framing of the Learning Curve Approach

In each RIA, the EPA frames its use of learning curves with an explicit call-out to the academic literature, and in particular, the aforementioned Dutton and Thomas (1984) and Argote and Epple (1990). The agency provides the distribution of progress ratios presented in Dutton and Thomas (1984) and reproduced in Figure 2, as well as cites other studies, such as Alchian (1963) and Benkard (2000), in support of the point that although there is a considerable amount of variation in the empirical observation of progress ratios, it is justified for the agency to select a progress ratio from other industries when trying to assess the costs of the components of a given technology package which are necessary to achieve compliance. The EPA acknowledges that learning curve effects can be lower in some industries (e.g., chemicals, nuclear power with approximately 11% savings; see Zimmerman 1982; Lieberman 1984), and difficult to decipher in others (e.g., the computer chip industry; see Gruber 1992).

The EPA also notes that areas involving direct labor and material cost savings are usually the source of the greatest savings; cited examples include a reduction in the number or complexity of component parts, improved component production, improved assembly speed and processes, reduced error rates, and improved manufacturing processes. These sorts of improvements all result in higher overall production, less scrappage of materials and products, and better overall quality. The agency is careful to provide statements related to these topics in its regulatory language because they help ground agency decisions on when it is appropriate and inappropriate to apply learning curves to compliance cost estimates.

EPA-NHTSA Derivation and Application of Learning Curves

The standard method and terminology used in the EPA-NHTSA approach today is described in detail in a 2008 staff technical report on the light-duty vehicle rulemakings for emissions of carbon dioxide and greenhouse gases, as well as for fuel economy (U.S. Environmental Protection Agency 2008; U.S. Environmental Protection Agency and U.S. Department of Transportation 2011).

The basic EPA-NHTSA approach has two parts. First, it selects a learning rate that is informed by a review of the literature (officially Dutton and Thomas 1984, although the fact that the EPA-NHTSA approach tends to use a learning rate of 20%, rather than 18-19%, implies a lingering influence of earlier government uses of learning rates in areas like antitrust policy and budgeting). Second, it applies that rate to the long-term direct costs (e.g., materials and labor) of the technology package that can alter a vehicle's emissions/fuel economy, as well as applies it to related consumer costs (i.e., vehicle purchase price rather than vehicle operating cost). The EPA-NHTSA approach explicitly does not apply the learning rate to the indirect costs (e.g., R&D, corporate

[4] Product teardown is the act of disassembling a product, such as a television set, to identify its component parts and functions.

operations, marketing, etc.; see EPA (2011)) of the technology package. In addition, the approach exercises some discretion as to whether or not to apply the learning curve adjustment to all components of a technology package, a situation which we describe in more detail below.

Derivation of Learning Curves

For many years – until NHTSA joined the EPA in applying learning curves in RIAs, starting in the mid-2000s – the EPA-NHTSA approach cited the literature as justifying the selection of a progress ratio that would result in a learning rate of 20% (i.e., a 20% reduction in cost for each doubling of production volume) for its technology package cost adjustment, unless it decided that no learning was likely to occur.

Once NHTSA joined the EPA in applying a learning curve-based cost adjustment approach, however, the EPA-NHTSA approach stopped using only a single 20% learning rate and instead selected a few learning rates that it deemed to be appropriate to more and less mature and widely adopted technologies. Figure 3 shows the distinction the EPA-NHTSA approach begins to make in the late 2000's between newer versus more mature technologies, in keeping with a major thread in the management literature on learning curves and the distinction found in that literature between phases in a technology's development in which induced and autonomous learning are more likely to play a role (see, e.g. Wiersma 2007). For newer technologies, the EPA-NHTSA approach explains that there is likely to be substantial learning in the near future – i.e., through induced learning – and it selects a "steep" learning rate of 20% for use in adjusting the relevant component costs. For more mature technologies that the agency judges to be widely available commercially, the EPA-NHTSA approach selects a "flat" learning rate (1-3%) or no learning rate (0%), in order to reflect the more limited learning opportunities associated with autonomous learning.

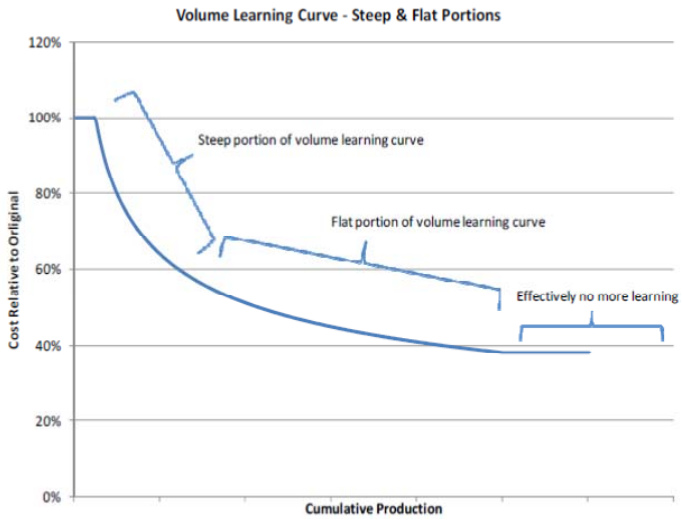

Figure 3: Tailoring the Learning Rate in the EPA-NHTSA Approach. Source: (U.S. Environmental Protection Agency and U.S. Department of Transportation 2011)

Application of Learning Curves

In applying the learning rate to adjust the direct costs of components of technology packages, the EPA-NHTSA approach has to make determinations about the maturity and diffusion of those components. It is aided in this by the agencies' access to product tear-down studies and confidential data from OEMs, derived either from submissions or meetings, as well as the agencies' access to vehicle emissions and CAFE certification data and data from Ward's Automotive Group. These data arrangements give the EPA and NHTSA an important understanding of industry practices, such as product redesign cycles, as well as access to relevant expertise that can inform determinations of new versus mature technologies.

When applying the learning rate to the direct costs of technology package components, the EPA-NHTSA approach makes the simplifying assumption that the production volume of the technology package is tied to the product cycle times of the automobile industry (shown to have been reduced from five years to four years by the end of the 1990s). For example, in many RIAs, the EPA-NHTSA approach assumes that a certain number of years after the imposition of a standard, the production volume of the technology package will double and after another like period of time, production will double a second time.[5] In years after this second doubling of

[5] For traditional air pollutant RIAs, the time period for production to double is usually two years, but this time period becomes longer for greenhouse gas RIAs, which assume a greater role for alternatives to the internal combustion engine.

production, the agencies assume that the technology package will become mature and the potential for further learning will be diminished as autonomous learning sets in. As a result, in many RIAs the EPA-NHTSA approach only applies the learning rate to the first two time periods (i.e. as production doubles the first two times) and does not adjust costs with a learning rate after this time.

As mentioned above, the EPA-NHTSA approach does not automatically adjust the direct costs of *every* component of a technology package by use of a learning rate; this was true even during the era of a uniform 20% learning rate or no learning rate, which was prevalent before NHTSA joined with EPA in applying learning curves to regulatory cost estimates. For example, in EPA's Tier 2/Tier 3 Non-Road Diesel Vehicles Emissions RIA in 1998, EPA chose not to apply a learning rate adjustment for such technology package components as: catalysts (due to the uncertainty of future precious metal prices), evaporative system costs (on the principle that these had already been well developed and anticipated system improvements were likely to be employed easily by manufacturers), and software costs (since they were deemed unlikely to align with manufacturing progress ratios). Note how these rationales for applying or not applying the learning rate reveal the agency's degree of familiarity with the state of the industries involved in manufacturing the relevant components.

Table 2 presents an example of the discretion employed by the EPA-NHTSA approach in applying different learning rates to different components of a technology package as part of the 2012 RIA for emissions of carbon dioxide and greenhouse gases, as well as for corporate average fuel economy, in light-duty vehicle models in 2017-2025. As an example of the more tailored approach to applying the learning rate in this post-NHTSA RIA, note the treatment of a less mature technology like air conditioner alternative refrigerants. This technology is considered to be initially subject to steep learning, with a consequent application of an initial 20% learning rate. After the initial period of steep learning (2016-2020), however, the EPA-NHTSA approach applies flat learning rates of 3% per year for five years, 2% per year for five years, then 1% per year for five years (U.S. Environmental Protection Agency 2012).

Table 2. Learning Rate by Technology Package Component, as cited in the Model Year 2017-2025 Light-Duty Vehicle Greenhouse Gas Emissions Regulations

Technology	Steep learning	Flat learning	No learning
Engine modifications to accommodate low friction lubes			2012-2025
Engine friction reduction – level 1 & 2			2012-2025
Low drag brakes			2012-2025
Secondary axle disconnect		2012-2025	
Cylinder deactivation		2012-2025	
Air conditioner alternative refrigerant	2016-2020	2021-2025	
Air conditioner related hardware		2012-2025	

Source: (U.S. Environmental Protection Agency and U.S. Department of Transportation 2011)

Note that the EPA-NHTSA approach includes sensitivity analyses to evaluate the impact of learning rate assumptions on projected prices. The basic sensitivity analysis approach is to bound the low, primary, and high side cost estimates via the use of a variety of learning rates.

Use of Learning Curves in Appliance MEPS

This section summarizes the DOE approach to incorporating learning curve cost adjustments in appliance minimum efficiency performance standards (MEPS) RIAs, which the agency began to do in 2011, independent of the EPA-NHTSA approach. Note that the DOE approach is not a component-based approach like the EPA-NHTSA approach, but instead considers the learning curve in the context of the price of the full appliance.

General DOE RIA Approach

The DOE has been tasked with consideration of Energy Conservation Standards – a form of MEPS – since 1979. The agency is required by statute to set forth MEPS that achieve the maximum improvement in energy efficiency given the constraints of technological feasibility and economic justification. The RIA approach addresses the requirement for economic justification through analyses of: life-cycle cost; economic impact on manufacturers; national benefits; impacts, if any, on utility companies; and impacts, if any, from lessening competition amongst manufacturers.

The DOE learning curve compliance cost adjustment approach has been a part of two of these analyses since 2011: (1) life-cycle cost and payback period analysis (LCC); and (2) national impact analysis (NIA). The LCC analysis focuses on the economic impacts of MEPS on individual consumers. It estimates the present value of costs incurred by consumers during the life of an appliance, including purchase and operating costs, discounting

future operating costs to the present (i.e., the analysis year). Purchase costs (also known in the RIA as the "total installed cost") are calculated by use of the following inputs: (1) the "baseline" manufacturer costs to meet the current MEPS; (2) the increases in manufacturing costs associated with meeting one of several proposed alternative efficiency standards; (3) the markups and sales tax increases associated with converting manufacturer costs to consumer product costs; and (4) the product installation cost to the consumer, which include labor, overhead, and miscellaneous materials and parts. Operating costs are estimated based on product efficiency level and energy consumption under usage conditions derived from data such as the Residential Energy Consumption Survey (RECS), as well as energy prices and trends, repair and maintenance costs, and assumed lifetime and discount rate distributions. Meanwhile, the NIA focuses on the national energy savings and the present value to the nation of the total consumer costs and savings projected for each potential standard level over the first thirty years after the appliance MEPS take effect.

Note that both the LCC and the NIA are linked to product cost-efficiency information derived from manufacturer interviews and product teardown analyses conducted by an independent organization. Three approaches are used by that organization to determine the cost-efficiency relationships of relevance to potential MEPS: (1) a design-option approach, which calculates the incremental costs of adding specific design options to a baseline model; (2) an efficiency-level approach, which calculates the relative costs of achieving increases in energy efficiency levels without regard to the particular design options used to achieve such increases; and/or (3) a reverse engineering or cost-assessment approach, which involves a "bottom-up" manufacturing cost assessment based on a detailed bill of materials derived from teardowns of the product being analyzed. The ultimate methodology is selected on a product-by-product basis, given the design options under study and any historical data that DOE can draw on.

DOE Framing of the Learning Curve Approach

Learning is discussed in several places in the technical support documents (TSDs) relating to each DOE rule. In the LCC overview of the inputs for establishing the total installed cost, the "learning rate" is defined as "the cost reduction factor associated with economies of scale and technology learning," implying that these two factors underlie long-term price drops in appliances, without distinguishing between them. The appendix associated with the LCC states that "examination of historical price data indicates that an assumption of constant real prices and costs may overestimate long-term price trends," and that "economic literature and historical data suggest that the real costs of those products may in fact trend downward over time according to learning, or experience, curves." In both of these instances, long-term price trends are equated with learning curves, as well as with experience curves. In the NIA, DOE frames the learning analysis in terms of "uncertainty in [DOE's] estimates of product price forecasts." The analysis is described as a "price-trend fit" to historical data, performed in order to capture potential long-term changes in price. Desroches, Garbesi et al (2013) fills in some of the additional theory underlying the approach, which is left out of the DOE RIA framework because of its traditionally succinct approach to providing descriptions of methodology.

DOE Application of Learning Curves

Derivation of Learning Curves

The DOE approach derives product-specific learning curves for use in its RIAs for proposed appliance MEPS relies on: (1) a time series proxy for historical cost; and (2) a time series proxy for historical production.

For the time series proxy for the historical cost of various appliances, the DOE approach turns to the Producer Price Indices (PPI) for the given appliance, since long time series of cost – and even price data – are not readily available for many appliances.[6] In some cases, product-specific PPI are not available and the PPI for a more aggregate category (e.g. household laundry equipment, miscellaneous appliances, etc.) is used instead (the disadvantage of this is that it assumes that product price trends within an aggregate category are roughly the same as the aggregate category as a whole). One of the advantages of the PPI as a proxy for cost over time is that the PPI "includes a quality adjustment, which attempts to factor out physical changes (such as capacity, premium features, government-mandated features, etc) in the product that affect the price"(Desroches, Garbesi et al. 2013).

For the time series proxy for production, the DOE approach uses historical U.S. shipments data as a measure of annual appliance production. Shipments data generally come from manufacturing trade associations (e.g., the Association of Home Appliance Manufacturers (AHAM), Gas Appliance Manufacturers Association (GAMA),

[6] The PPI is one of the oldest continuous systems of statistical data compiled by the U.S. federal government. In it, the Bureau of Labor Statistics (BLS) reports the average change in the selling prices received by domestic producers for their output. The prices included in the PPI are from the first commercial transaction for many products and some services.

the Air Conditioning, Heating, and Refrigeration Institute, etc.), either directly through data request or through trade association publications.

Application of Learning Curves

The DOE approach applies the learning curve cost-adjustment by way of a price factor index set equal to 1 in the analysis year. The price factor for future years is a function of the learning rate and the estimated cumulative production through that year, as based on forecasts of annual shipments. The price factor index is the key to estimating the future costs in each RIA. The price factor for a given year is calculated as follows:

Equation 5: Price Factor Index applied in the DOE Approach to Learning Curve-based Cost Adjustment

$$PF_i = PF_{i-1} \times \left(\frac{X_i}{X_{i-1}}\right)^{-b}$$

Where:

PF_i =	price factor in year *i* (*PF* = 1 in analysis year),
X_i =	cumulative production in year *i*,
b =	learning rate parameter

In the LCC, the price factor is used to estimate the consumer product cost in the year the MEPS will come into effect. Future year consumer product costs are calculated by applying the annual price factor to each product class at each considered efficiency level. In the NIA, price factors are applied in the calculation of the net present value (NPV) of the increased product cost and reduced operating cost associated with the difference between the base case and each potential standards case for the considered product.

Discussion

In February 2011, when OIRA called for federal agencies "to use the best available techniques to quantify anticipated present and future benefits and costs as accurately as possible," including "identifying changing future compliance costs that might result from technological innovation or anticipated behavioral changes," those best-available techniques were not obvious. Indeed, DOE's NODA shows the range of options that were considered available at the time:

(1) fitting an experience curve to available data on the cost trends for equipment or technologies that are components of a given efficiency design option in order to forecast the future cost of that design option;

(2) using experience curve cost trends for the analyzed product as a whole, in order to project both the price of that item as well as the price of the more efficient product options;

(3) applying an experience curve cost trend derived from an analogous product, equipment, or grouping of products or equipment (which includes the analyzed product) in order to project both the price of the analyzed product as well as the price of the more efficient product options; and

(4) using experience curve parameters or ranges of parameters drawn from published review articles that are applicable to certain classes or groups of products or equipment (which include the analyzed product) in order to project both the price of the analyzed product as well as the price of the more efficient product options (U.S. Department of Energy 2011).

When considering the "best available" options, a combination of the first and fourth option above can be considered to be most widely available, given the oeuvre of RIAs conducted by the U.S. federal government. This is because of the extensive track record of the EPA and later NHTSA in applying its learning curve-based adjustment to the direct compliance costs of the components of compliance technology packages. Access to significant data and expertise allows the EPA-NHTSA approach to consider each component's degree of diffusion in the marketplace and likely potential to see steep, flat, or no learning effects. The agencies are then able to apply associated learning rates, as derived from a classic review article, to estimates of the long-term direct costs (e.g., materials and labor) and related consumer costs (i.e., vehicle purchase price rather than vehicle operating cost) of achieving the proposed regulation (see USEPA and USDOT 2011). Note that the EPA-NHTSA approach applies these learning rates over only a limited time period, selected based on the assumption that the regulation creates the initial condition for producing a compliant vehicle, and that learning effects cannot be estimated with accuracy for more than a few doublings of cumulative production of that vehicle. It is important to recall that the indirect costs of compliance (e.g., R&D, corporate operations, marketing, etc.) are not subject to the learning rate adjustment.

For a variety of reasons related to data and fit with existing institutional arrangements, the DOE has, thus far, not adopted a similar component-based approach, and instead has employed a combination of the second and third

options listed above. It derives learning rates from historical price data for the full product to be subject to MEPS, and applies that product-specific learning rate to adjust the costs it projects in the LCC and NIA.

Although there does not yet appear to be a seminal retrospective review of the EPA approach, we were able to complete a first-order review of the improvements in accuracy the DOE approach provides, when compared to the traditional RIA approach of holding costs constant throughout the analysis period. Our review leveraged heavily a paper that investigated trends in retail appliance prices and compared these prices to DOE MEPS projections, finding that MEPS projections tended to overestimate retail prices (see Dale, Antinori et al. 2009). We performed a simple analysis comparing the same DOE price projections from previous rulemakings in the Dale, Antinori et al (2009) paper with the price estimates that would have been produced under a hypothetical scenario in which the current DOE approach to adjusting costs using a learning curve based approach was available at the time. Then we compared both of these estimates to the observed prices used in the Dale, Antinori et al (2009), with error bar estimates of 25%, following the approach to determining over- and under-estimation of the costs of compliance in RIAs followed by the seminal Harrington, Morgenstern, et al (2000) paper. Through this analysis, we observed that the current DOE approach to applying the learning curve-based cost adjustment in the RIAs does a better job at forecasting costs than the previous constant-cost assumption. But we believe that there is considerable room to refine the method, to bring the projections closer to the actual price and within the error bars of significant over- and under-estimation. We believe that we have much to learn still about the best available approach to adjusting RIAs to consider technological change, but that the learning curve approach holds great promise.

Acknowledgements

The authors would like to thank Larry Dale and Anna Spurlock for their substantive and strategic advice. In addition, we would like to thank Louis-Benoit Desroches, Steven Meyers, Alan Sanstad, Greg Rosenquist, and Alex Lekov for helpful comments.

References

Adler, P. S. and K. B. Clark (1991). "Behind the Learning Curve: A Sketch of the Learning Process." Management Science **37**(3): 267-281.

Alchian, A. (1963). "Reliability of Progress Curves in Airframe Production." Econometrica **31**(4): 679-693.

Argote, L. (1999). Organizational Learning: Creating, Retaining, and Transferring Knowledge. Norwell, MA, Kluwer Academic Publishers.

Argote, L., S. L. Beckman, et al. (1990). "The Persistence and Transfer of Learning in Industrial Settings." Management Science **36**: 140-154.

Argote, L. and D. Epple (1990). "Learning Curves in Manufacturing." Science **247**: 920-924.

Argote, L., W. McEvily, et al. (2003). "Managing Knowledge in Organizations: An Integrative Framework and Review of Emerging Themes." Management Science **49**(2003): 571-582.

Arrow, K. (1962). "The economic implications of learning by doing." Review of Economic Studies **29**: 155-73.

Baloff, N. (1966). "The learning curve-Some controversial issues." Journal of Industrial Economics **14**: 275-82.

Benkard, C. L. (2000). "Learning and Forgetting: The Dynamics of Aircraft Production." The American Economic Review **9**(4): 1034 -1054.

Billon, S. A. (1966). "Industrial learning curves and forecasting." Management International Review **6**(6): 65-96.

Conway, R. and A. Schultz (1959). "The Manufacturing Progress Function." Journal of Industrial Engineering **10**(1): 39-53.

Dale, L., C. Antinori, et al. (2009). "Retrospective evaluation of appliance price trends." Energy Policy **37**(2): 597-605.

Desroches, L.-B., K. Garbesi, et al. (2013). "Incorporating experience curves in appliance standards analysis." Energy Policy **52**(2013): 402-416.

Dutton, J. M. and A. Thomas (1984). "Treating Progress Functions as a Managerial Opportunity." The Academy of Management Review **9**(2): 235-247.

E.H. Pechan & Associates (1994). Analysis of Costs, Benefits, and Feasibility Regarding Implementation of OTC Petition on California Low Emission Vehicles.

Ferioli, F., K. Schoots, et al. (2009). "Use and limitations of learning curves for energy technology policy: A component-learning hypothesis." Energy Policy **37**(2009): 2525-2535.

Fiol, C. M. and M. A. Lyles (1985). "Organizational Learning." Academy of Management Review **10**(4): 803-813.

Gruber, H. (1992). "The learning curve in the production of semiconductor memory chips." Applied Economics **1992**(24): 885-894.

Hall, G. and S. Howell (1985). "The experience curve from the economist's perspective." Strategic Management Journal **6**(3): 197-212.

Harrington, W., L. Heinzerling, et al., Eds. (2009). Reforming Regulatory Impact Analysis. Resources for the Future Report. Washington, D.C., Resources for the Future.

Harrington, W., L. Heinzerling, et al. (2009). Reforming Regulatory Impact Analysis, Resources for the Future.

Harrington, W., R. D. Morgenstern, et al. (2000). "On the Accuracy of Regulatory Cost Estimates." Journal of Policy Analysis and Management **19**(2): 297-322.

Hayes, R. H. and K. B. Clark (1986). "Why Some Factories are More Productive than Others." Harvard Business Review **64**: 66-73.

Hirsch, W. Z. (1952). "Manufacturing Progress Functions." Review of Economics and Statistics **34**: 143-155.

Hirsch, W. Z. (1956). "Firm Progress Ratios." Econometrica **2**: 136-143.

Hwang, R. and M. Peak (2006). Innovation and Regulation in the Automobile Sector: Lessons Learned and Implications for California's CO2 Standards.

Joskow, P. L. and G. A. Rozanski (1979). "The effects of learning by doing on nuclear plant operating reliability." Review of Economics and Statistics **61**: 161-168.

Kantor, P. B. and W. I. Zangwill (1991). "Theoretical Foundation for a Learning Rate Budget." Management Science **37**(3): 315-330.

Kieschel, W. (1981). "The decline of the experience curve." Fortune **1981**: 139-145.

Levin, D. Z. (2000). "Organizational Learning and the Transfer of Knowledge: An Investigation of Quality Improvement." Organization Science

11(6): 630–47.

Lieberman, M. B. (1984). "The Learning Curve and Pricing in the Chemical Processing Industries." The RAND Journal of Economics **15**(2): 213-228.

Lieberman, M. B. (1987). "The learning curve, diffusion, and competitive strategy." Strategic Management Journal **8**: 441-452.

Lindgren, L. (1978). Cost Estimations for Emission Control Related Components/Systems and Cost Methodology Description. U. S. E. P. Agency.

Nemet, G. F. (2009). "Interim monitoring of cost dynamics for publicly supported energy technologies." Energy Policy **37**(2009): 825-835.

OECD (2009). Regulatory Impact Analysis: A Tool for Policy Coherence.

Organisation for Economic Co-operation and Development (2009). Regulatory Impact Analysis: A Tool for Policy Coherence.

Papineau, M. (2006). "An economic perspective on experience curves and dynamic economies in renewable energy technologies." Energy Policy **34**(2006): 422-432.

Pisano, G. P., R. M. J. Bohmer, et al. (2001). "Organizational differences in rates of learning: Evidence from the adoption of minimally invasive cardiac surgery." Management Science **47**(6): 752-768.

Porter, M. E. (1980). Competitive strategy: Techniques for analyzing industries and competitors. New York, Free Press.

Rapping, L. (1965). "Learning and World War II production functions." Review of Economics and Statistics **47**: 81-86.

Reagans, R., L. Argote, et al. (2005). "Individual Experience and Experience Working Together: Predicting Learning Rates from Knowing Who Knows What and Knowing How to Work Together." Management Science **51**(6): 869-881.

Sheshinski, E. (1967). "Tests of the Learning by Doing Hypothesis." Review of Economics and Statistics **49**(4): 568-578.

Sorenson, O. (2003). "Organizational Learning and the Effect of Integration." Management Science **49**(4): 446-463.

Taylor, M. R. (2012). "Innovation under Cap and Trade Programs." Proceedings of the National Academy of Sciences **Early Edition**.

U.S. Department of Energy (2011). Equipment Price Forecasting in Energy Conservation Standards Analysis: Notice of Data Availability; Request for Comment. . O. o. E. E. a. R. Energy, Federal Register. **76:** 9696-9700.

U.S. Environmental Protection Agency (1997). Final Regulatory Impact Analysis: Control of Emissions of Air Pollution from Highway Heavy-Duty Engines.

U.S. Environmental Protection Agency (2008). EPA Staff Technical Report: Cost and Effectiveness Estimates of Technologies Used to Reduce Light-duty Vehicle Carbon Dioxide Emissions. **EPA-420-R-08-008**.

U.S. Environmental Protection Agency (2012). Regulatory Impact Analysis: Final Rulemaking for 2017-2025 Light-Duty Vehicle Greenhouse Gas Emission Standards and Corporate Average Fuel Economy Standards. **EPA-420-R-12-016**.

U.S. Environmental Protection Agency and U.S. Department of Transportation (2011). Draft Joint Technical Support Document: Proposed Rulemaking for 2017-2025 Light-Duty Vehicle Greenhouse Gas Emission Standards and Corporate Average Fuel Economy Standards. **EPA-420-D-11-901**.

U.S. Environmental Protection Agency and U.S. Department of Transportation (2011). Final Rulemaking to Establish Greenhouse Gas Emissions Standards and Fuel Efficiency Standards for Medium- and Heavy-Duty Engines and Vehicles Regulatory Impact Analysis. **EPA-420-R-11-901**.

U.S. Office of Management and Budget (2011). Memorandum for the Heads of Executive Departments and Agencies and of Independent Regulatory Agencies on Executive Order 13563, "Improved Regulation and Regulatory Review". Office of Information and Regulatory Affairs. Washington, D.C. **M-11-10**.

von Hippel, E. A. (1976). "The Dominant Role of Users in the Scientific Instrument Innovation Process." Research Policy **5**(3): 212-239.

Weiss, M., M. Junginger, et al. (2010). "A review of experience curve analyses for energy demand technologies." Technological Forecasting & Social Change **77**: 411-428.

Weiss, M., M. K. Patel, et al. (2010). "Analyzing price and efficiency dynamics of large appliances with the experience curve approach." Energy Policy **38**(2): 770-783.

Wiersma, E. (2007). "Conditions That Shape the Learning Curve: Factors That Increase the Ability and Opportunity to Learn." Management Science **53**(12): 1903-1915.

Wright, T. P. (1936). "Factors Affecting the Cost of Airplanes." Journal of Aeronautical Science **3**: 122-128.

Yelle, L. E. (1979). "The Learning Curve: Historical Review and Comprehensive Survey." Decision Sciences **10**(2): 302-328.

Zimmerman, M. B. (1982). "Learning Effects and the Commercialization of the New Energy Technologies." Bell Journal of Economics **13**: 297-310.