PARALLEL REAL ASSET MANAGEMENT WITH ENVIRONMENTAL REGULATION: INTEGER PROGRAMMING AND APPROXIMATE DYNAMIC PROGRAMMING APPROACHES

A Dissertation

Presented to the Faculty of the Graduate School

of Cornell University

in Partial Fulfillment of the Requirements for the Degree of

Doctor of Philosophy

by

Timon Herrick Stasko

August 2011

PARALLEL REAL ASSET MANAGEMENT WITH ENVIRONMENTAL REGULATION: INTEGER PROGRAMMING AND APPROXIMATE DYNAMIC PROGRAMMING APPROACHES

Timon Herrick Stasko, Ph.D.

Cornell University 2011

This dissertation presents a pair of models designed to assist in the management of multiple deteriorating real assets, given financial and environmental concerns. Whether the assets are buildings or vehicles or machines, their purchase and upkeep can be costly, making optimal management policies valuable. The models presented build upon a strong literature. They incorporate numerous factors which have been modeled previously, though generally not together. These include technological change, linked decisions for multiple assets, and non-steady-state demand. They stand out from previous literature due to their ability to model retrofits, as well as repairs and replacements. These retrofits can have initial as well as ongoing costs, and can impact externalities, making them relatively general. The integer program model is fast and well suited to analysis requiring large numbers of runs, such as the comparison of a wide range of regulatory alternatives. The approximate dynamic program, while slower, is able to handle stochastic asset failures and repair costs for large asset portfolios, something which previous models have struggled to accomplish without strong simplifying assumptions. A customized value iteration approach produces good solutions within a few hours for sample problems involving a fleet of well over a thousand vehicles subject to clean diesel regulation.

BIOGRAPHICAL SKETCH

Timon Stasko was born on June 24th, 1985 to Jennifer and Joseph Stasko. For the majority of his childhood, Timon lived with his family on the campus of The Masters School in Dobbs Ferry, New York. In 2003, Timon received his high school diploma from The Masters School. In 2007, Timon received his B.S. in Civil Engineering from Cornell University. He continued his studies at Cornell by joining Prof. Gao's research group as a Ph.D. student in January of 2007. As a Ph.D. student, he worked on several projects involving transportation and the environment. In his free time, Timon enjoys the outdoors through hiking, biking, kayaking and photography.

ACKNOWLEDGMENTS

I thank my advisor, Oliver Gao, for making graduate school an incredibly rewarding experience for me. You were a constant source of enthusiasm and ideas, while remaining open to new directions proposed by students.

I thank my committee members, Mark Turnquist and Jon Conrad, for their feedback and guidance. Your suggestions inspired several aspects of the models presented. I also thank other professors who have provided valuable insights, including Linda Nozick and David Shmoys.

I don't know where to begin when thanking my parents. You worked so hard to provide me with opportunities. You inspire me to take full advantage of my good fortune. I thank my brother, Cory, both for his helpful questions and his ability to help me relax and recharge.

I thank Masha Pitiranggon for being my constant companion, and for helping me to look at the same problems from new perspectives.

I greatly appreciate all the support I received from my fellow graduate students: Andrew Buck, Yohannes Kesete, Meredith Legg, Brian Levine, Anna Li, Allison Reilly, Soham Shah, Darrell Sonntag, Xun (Richard) Wang and Russell Winans. Our conversations, both those directly related to my research and those with no apparent connection, helped me to keep moving forward. Dianne Dietrich, you count too.

I thank New York Metropolitan Transportation Council (NYMTC), New York State Department of Transportation (NYS DOT), U.S. Department of Transportation (US DOT) and U.S. Environmental Protection Agency (EPA) for their financial support, advice, and data, all of which were essential to developing practical models. NYMTC provided financial support through the September 11th Memorial program for Regional Transportation Planning. NYS DOT funded this work through research project C-07-12. US DOT support came in the form of an Eisenhower Graduate Fellowship. The EPA provided funding through STAR Fellowship Assistance Agreement no. 91717401-0.

The contents of this dissertation have not been formally reviewed by any of the project sponsors. The views expressed in this dissertation are solely those of Timon Stasko and project sponsors do not necessarily endorse any products or commercial services mentioned in this dissertation.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

LIST OF ABBREVIATIONS

ADP	approximate dynamic program
ADPF	active diesel particulate filter
BART	best available retrofit technology
CARB	California Air Resources Board
CMAQ	Congestion Mitigation and Air Quality (Improvement Program)
CCVS	closed crankcase ventilation system
CNG	compressed natural gas
DMF	diesel multi-stage filter
DOC	diesel oxidation catalyst
DPF	diesel particulate filter
DSNY	Department of Sanitation of New York City
EGR	exhaust gas recirculation
ESW	Environmental Solutions Worldwide
FTF	flow through filter
IARC	International Agency for Research on Cancer
IP	integer program
LAO	(California) Legislative Analyst's Office
LNF	low NO_2 filter
LP	linear program
LXF	low NO_x filter
MAPE	mean absolute percentage error
M DEP	Massachusetts Department of Environmental Protection
MSAT	Mobile Source Air Toxics
NJ DEP	New Jersey Department of Environmental Protection
NV DMV	Nevada Department of Motor Vehicles
NYC DOT	New York City Department of Transportation
NYCRR	New York Codes, Rules, and Regulations
NYS DOT	New York State Department of Transportation
NYS DEC	New York State Department of Environmental Conservation
OPWA	Oakland Public Works Agency

PDPF	passive diesel particulate filter
PM	particulate matter
PAH	polycyclic aromatic hydrocarbon
ULSD	ultra-low-sulfur diesel
US CDC	U.S. Centers for Disease Control and Prevention
US DHHS	U.S. Department of Health and Human Services
US DOE	U.S. Department of Energy
US DOJ	U.S. Department of Justice
US EPA	U.S. Environmental Protection Agency
VTL	New York Vehicle and Traffic Law

CHAPTER 1

THE STUDY OF ASSET REPLACEMENT

Many organizations, from private corporations to government agencies, depend on large numbers of capital assets to accomplish their objectives. Whether the assets are in the form of a fleet of vehicles or a group of buildings, sizable capital investments and operational expenses are involved. There is considerable interest in cost-minimizing asset maintenance and replacement strategies, with an increasing emphasis on emissions reduction. Deciding when assets should be bought, sold, repaired, and retrofitted is no simple task, especially given the uncertainty surrounding when assets or components will break down. Fortunately, there is a sound foundation of literature on which to build decision support systems.

It speaks to the importance of asset replacement that Richard Bellman chose it as one of the earliest applications for dynamic programming, a computational tool which is now extremely widely used. In Bellman's 1955 paper "Equipment Replacement Policy," he examined the question of when to replace a single piece of aging equipment, given deterministically increasing upkeep costs and deterministically decreasing output. He stated that with enough simplifying assumptions, "the problem may be resolved quite easily."

Bellman acknowledged that factors like technological improvement could complicate the problem, but he likely had no idea how much research would be done on the countless complicating factors which are important in different situations. Thousands of models have been published (Wang, 2002) and categorized by numerous reviews. Wang (2002) provides a broad review of single-asset models, with limited coverage of multi-asset models. Cho and Parlar (1991) review multi-asset models with economic or stochastic dependence, while Dekker et al. (1997) focus on multi-asset maintenance models with economic interdependence. Pham and Wang (1996) review models which incorporate the concept of imperfect maintenance.

The features, level of sophistication, and objectives of published models are largely a function of the applications for which the models were designed. Often, the obvious objective is to minimize financial costs. In some cases, particularly when dealing with a system of interdependent components, the objective is to maximize the reliability (Moghaddam and Usher, 2011). Where system uptime is of great importance, the rate at which repairs are completed and the inventory of spare components can become important model features (Cho and Parlar, 1991).

Even given the seemingly straightforward objective of minimizing costs, there are numerous ways of quantifying the objective. The appropriate method depends on the scope of the model and the perspective of the modeler. When there is no discounting and a finite time horizon, total cost can be minimized. Without discounting, costs over an infinite time horizon are likely to be infinite, making total cost per period a more computationally appealing metric. When there is discounting, either the net present cost or the annualized cost can be minimized. This dissertation will allow for discounting, and will generally use net present cost.

Just as the objectives vary across models, so do the policy structures. For a single asset, the policy might come in the form of a uniform retirement age (Bellman, 1955) or set of limits on how much money or time to spend repairing it, possibly depending on its age or other factors (Drinkwater and Hastings, 1967; Nguyen and Murthy, 1981). For multiple assets, one can still apply single-asset policies such as replacing only at complete failure or replacing at a fixed retirement age. Alternative rules have been developed to take advantage of economies of scale, such as replacing assets at failure, while also replacing all assets every x periods, independent of age.

Berg and Epstein (1978) compared these relatively simple policies and proposed a rule for choosing which policy to follow based on the costs involved. Unsurprisingly, replacing all units every x periods is more costly unless there are noticeable economies of scale in replacement costs. Naturally, various alternative policies for grouping replacements have been developed (Cho and Parlar, 1991). Other researchers have avoided recommending simple policy rules for the multi-asset case, favoring a full mathematical programming approach.

This dissertation will focus on groups of assets that are used to accomplish combined goals, but which can function independently. For example, a group of trucks may be used together to plow a road network under a common budget, but if one truck breaks down it does not mean that other trucks must also stop operating. In another example, a portfolio of investment properties may be managed together to maximize financial returns while minimizing financial risk and environmental damage, but if one building requires a repairs it does not necessarily mean that other buildings do as well. As a contrasting example, an entire assembly line may have to stop operating if one stage is not operational. Such an example typically requires a different analysis approach.

If the connections between the assets are relatively limited, the problem can be decomposed and solved for individual assets. In the simplest case, where all costs are deterministic, one can determine a single cost-minimizing lifespan. Given that the optimal policy is defined by a single variable, generally no sophisticated optimization algorithm is required. In many cases, it makes the most sense to quickly compute the equivalent uniform annual costs for a range of lifespans in a spreadsheet, and simply pick the cheapest (Newnan et al., 2002). Yatsenko and Hritonenko (2008) extended the problem by allowing technology change which impacts both purchase and maintenance costs, though both are assumed to be known. They developed a nonlinear integral equation for the optimal service life, which can vary over time. Selecting an optimal lifetime becomes slightly more complicated if the asset can fail randomly. Howard (1960) addressed this problem with dynamic programming, while Ghellinck and Eppen (1967) used a linear program. Both still dealt only with expected operational costs (as opposed to a distribution) and produced policies based on asset age.

Repair costs are often far from deterministic, however. If repair costs are stochastic, a simple retirement age prescription will not necessarily minimize costs. Multiple researchers have found “repair limits” to be a superior policy to a fixed retirement age (Drinkwater and Hastings, 1967; Love et al., 1982). Hastings (1968) showed that repair limits are the optimal policy structure, given certain assumptions such as the policy maker being unable to influence asset utilization levels. The concept of a repair limit is quite simple. If a damaged asset requires more than this limit to repair, it should be replaced. Otherwise, it should be repaired and kept. Repair limits are typically permitted to vary with the asset’s age (and sometimes with other variables). This makes repair limits more difficult to solve for than a simple retirement age.

Drinkwater and Hastings (1967) presented an iterative method of computing successively improved repair limits as well as an analytical solution for a special case with a Poisson distribution for repair frequency and an exponential distribution for the repair costs. The analytical solution derived an equation which was solved with the method of steepest descents. Hastings (1968) presented a dynamic programming approach to solving for repair limits. Much of the early work on repair limits did not consider the impact of discounting, but later research found that repair limit policies can be sensitive to the discount rate (Love et al, 1982).

The parameters over which repair limits are allowed to vary play an important role in determining the difficulty of the problem. Repair limits typically change with a single “aging” parameter which could be the asset’s actual age in years or measure of usage such as total mileage or hours of operation. This setup carries an implicit assumption that future repair cost distributions are independent of past repair costs. Love et al. (1982) found low serial correlation in repair cost data on Postal Canada vehicles, supporting this assumption. This assumption is powerful in its ability to control the size of the problem. It allows the problem to be modeled as a Markov process, meaning that the probability of an event is dependent only on the current state of the asset, not on its history. Even when this assumption holds, it is sometimes necessary to allow repair limits (and the related residual asset values) to vary with multiple parameters. For example, in non-steady-state situations, they could vary with time as technologies, demands or regulations change.

Various authors have extended the single-asset replacement problem without describing the optimal strategy with a simple rule such as a retirement age or set of repair limits. Rust (1987) examined engine replacements at a bus company, and estimated the impacts of various factors, such as the potential loss of goodwill with riders due to breakdowns. Kim and Makis (2009) developed a steady-state policy iteration algorithm for a single asset which was subject to multiple kinds of failures. In addition to choosing when to repair or replace the asset, the decision maker could also choose the level of repair to conduct.

When the connections between replacement decisions are strong, the decisions must be made simultaneously, which can greatly complicate the problem. This can be the case with fleets of vehicles. In a fleet setting, replacements are often not as tidy as a single new vehicle replacing a single retiring vehicle. The number of purchases in a given period may not equal the number of retirements. Purchase decisions can be linked by a common budget, or by economies of scale in vehicle purchases. Simms et al. (1984) tackled fleet vehicle replacement in a deterministic setting. Their non-linear objective, integer variables, and non-convex feasible region led them to a dynamic programming approach, despite concerns about the size of the state space. The size of their state space was particularly problematic because they solved their dynamic program by enumerating all feasible states in all years. Karabakal et al. (1994) presented an alternate methodology for replacing multiple assets under shared budgets, in a deterministic setting. They used a branch-and-bound algorithm to solve an integer program, including a Lagrangian relaxation of budget constraints.

Multiple authors have examined the effects of stochastic demand while keeping maintenance costs deterministic. Hartman (2004) did so, while allowing the decision maker to simultaneously exercise some control over asset usage. The problem was solved using dynamic programming, and despite efforts to trim the state space using symmetry Hartman expressed concerns about scalability. Much of the analysis focused on the two-asset case. Hsu et al. (2011) used a dynamic program for optimizing an airplane purchase, leasing, and disposal decisions in a fleet setting. The dynamic program was solved using backward induction, which typically involves enumerating all states in all periods.

Love et al. (1982) presented a dynamic programming approach for determining repair limits for multiple assets in the steady-state. It was a revised version of the general policy improvement routine of Howard (1960). Repair costs were stochastic and multiple replacement decisions were linked by a common capital budget, but all existing assets and replacements were considered identical in characteristics and in usage, restricting the model to steady-state applications.

Enormous state spaces can make multi-asset replacement problems extremely difficult to solve, even without stochastic costs. Numerous researchers have sought to reduce the size of the state space by making simplifying assumptions. Jones et al. (1991) demonstrated that the size of the problem could be dramatically reduced using two theorems. The first, known as the no-splitting rule, states that there is an optimal strategy in which all assets of the same age are treated the same way in any given period. The second, known as the older cluster replacement rule, states that it is only optimal to replace an asset if all older assets have been replaced. Childress and Durango-Cohen (2005) extended adaptations of these rules to the stochastic case, given comparable assumptions. The older cluster replacement rule requires several assumptions about cost structure. In particular, the assumption that the sum of maintenance cost, operating cost and salvage value is nondecreasing with respect to vehicle age is questionable, as relatively new vehicles often exhibit rapid declines in salvage value (McClurg and Chand, 2001). While the no-splitting rule does not require such assumptions about the cost structure, it falls apart in the presence of a binding budget constraint. Nonetheless, a significant portion of the literature on multi-asset replacement uses these or similar assumptions to reduce problem size (Jones et al., 1991; Chen, 1997; Jin and Kite-Powell, 2000; McClurg and Chand, 2002; Childress and Durango-Cohen, 2005).

This dissertation will present a pair of multi-asset replacement models. While they were both designed for fleets of vehicles, the concepts can be applied more broadly. The first model, presented in Chapter 4, is an integer program capable of handling large numbers of parallel assets, technology change and regulation. It is fast to solve, but is limited to purely deterministic problems. The second model, presented in Chapter 5, adds the capacity to model stochastic costs and stochastic asset failures. Instead of relying on strong simplifying assumptions to produce solutions, it uses an approximate dynamic programming approach which is well suited to dealing with enormous state spaces. Both the multi-period integer program and the approximate dynamic program are capable of solving for optimal policies in steady-state and non-steady-state periods. Given the dynamic evolution of demands, prices, technologies, and regulations, this is a valuable capability. In addition, both models include the possibility for retrofits, which are rarely discussed in the asset replacement literature. Some previous work has included narrowly defined retrofits under the name of "imperfect maintenance." Imperfect maintenance returns an asset to a younger state, without making it equivalent to a brand new asset Pham and Wang (1996). The retrofits in the models presented in this dissertation are more general in their abilities.

CHAPTER 2

BACKGROUND ON DIESEL VEHICLES AND EMISSIONS

2.1 Health Impacts of Diesel Emissions

From transporting freight to bringing children to school, diesel engines serve important purposes throughout our economy, but there is considerable evidence that diesel exhaust can be harmful to human health. Diesel exhaust is a general term used to describe a combination of many different gaseous compounds and heterogeneous particles produced by diesel engines. The quantity and composition of diesel exhaust can vary considerably depending on the vehicle, the fuel, and the driving conditions. Some research examines the health impacts of diesel exhaust as a whole, while a considerable body of work focuses on specific components found in diesel exhaust, as well as exhaust from other sources. The potential health impacts of diesel exhaust exposure range from frustrating to life threatening. A 2005 report by the Clear Air Task Force estimated that diesel fine particles shortened over 2,700 lives in the New York metropolitan area in 1999, more than in any of the other 39 regions studied (Clean Air Task Force, 2005).

Multiple studies have found that exposure to components of diesel exhaust is associated with higher risk of stroke (Tsai et al. 2003; Hong et al. 2002). Tsai et al. (2003) compared hospital admissions for primary intracerebral hemorrhage and ischemic stroke with pollutant concentrations on the same day, while Hong et al. (2002) compared stroke mortality data with pollutant concentrations. Both studies took place over a period of roughly 4 years, but in different Asian cities. Both studies took into account other factors such as temperature. The studies found that PM_{10} (particulate matter under 10 micrometers in diameter) and NO_2 concentrations were significantly associated with hospital admissions and mortality. Results for other pollutants were more mixed (Tsai et al. 2003; Hong et al. 2002).

Numerous studies have examined the connection between exposure to diesel exhaust (or its components) and cancer risk. The strength of the scientific evidence linking diesel exhaust with cancer has led multiple respected bodies to formally recognize the probability of diesel exhaust acting as a carcinogen. Based on human exposure studies and animal testing, diesel exhaust was declared a probable human carcinogen by the International Agency for Research on Cancer (IARC, 1998). In 1996, the World Health Organization found that diesel exhaust is probably carcinogenic (CARB, 2008a). The California EPA found a causal link between diesel exhaust and lung cancer in 1998, and by 2000, the U.S. Department of Health and Human Services' National Toxicology Program listed diesel exhaust as reasonably anticipated to be a human carcinogen (CARB, 2008a; US DHHS, 2005).

Several components of diesel exhaust have been identified as regulatory targets because of their health impacts. In a 2001 rule, the EPA designated "diesel particulate matter and diesel exhaust organic gases" as one of the 6 priority Mobile Source Air Toxics (MSAT) categories (Carr et al., 2007; Claggett and Houk, 2006). Other components of diesel exhaust are also considered priority MSATs, including benzene and formaldehyde (US OSHA, 2009). The EPA considers benzene to be a known human carcinogen, while it considers formaldehyde to be a probable human carcinogen. Acute exposure to benzene is associated with drowsiness, dizziness, headaches and unconsciousness. Chronic inhalation of benzene can affect bone marrow, causing blood disorders including aplastic anemia, excessive bleeding, and damage to the immune system (e.g. loss of white blood cells). Exposure to formaldehyde has been associated with respiratory symptoms, as well as eye, nose, and throat irritation (US EPA, 2009a).

Diesel exhaust contains quite a few other gaseous components which have health and/or environmental consequences, despite not being MSATs. Examples include carbon dioxide, carbon monoxide, and nitrogen oxides. Carbon dioxide is well known for its contribution to global climate change. Carbon monoxide reduces oxygen delivery to the body's organs, with those who suffer from heart disease facing the greatest risk of cardiovascular effects (US EPA, 2011a). Currie et al. (2007) found that carbon monoxide exposure significantly increased school absences, even when levels were below federal air quality standards. Short-term exposure to nitrogen dioxide has been linked to adverse respiratory effects (US EPA, 2009b). Nitrogen oxides are also known for their complicated role in influencing concentrations of ozone, which has its own negative health impacts such as inflammation of airways and decrements in lung function (US EPA, 2011b). More minor impacts of diesel exhaust include eye and nose irritation (Rudell et al., 1996) and increased susceptibility to allergic materials (Wargo et al., 2002).

Asthma is a major concern for both children and adults in the United States. It is one of the most common long-term diseases in children (US CDC, 2009) and the CDC estimated that 7.7 percent of the U.S. population had asthma in 2007 (US CDC, 2008). Diesel exhaust exposure has been linked to exacerbated childhood asthma (Tolbert et al., 2000; Slaughter et al., 2003). The development of asthma in children who tend to play outdoors in areas with high levels of air pollution may be linked to ozone exposure (McConnell et al., 2002). Nadeau et al. (2010) studied several groups of children in Fresno and Stanford, California. They found that increased exposure to ambient air pollution is associated with impaired regulatory T-cell function, increasing asthma morbidity (regulatory T-cells suppress immune responses).

Asthma imposes significant costs on society. In 2002, asthma medications were estimated to cost $500 per child per year (Wargo et al., 2002). A 2003 study based on survey data from California found that the total cost of adult asthma averaged $4912 per person per year, with pharmaceuticals accounting for the largest portion, $1605. Other major components were lost work time and hospital visits. Cost figures varied substantially across individuals, depending on the severity of the asthma (Cistemas et al., 2003). The CDC estimated that asthma costs in the United States totaled more than $30 billion in 2007 (US CDC, 2008). Many of asthma's impacts are difficult to quantify in dollar terms. In 2007, asthma accounted for 3,447 deaths in the United States (US CDC, 2008). Historically, approximately half of the people who die due to asthma are 65 or older (US CDC, 2007). One study found children with asthma were more likely to have learning disabilities, even after adjustment for demographic factors (Fowler et al., 1992).

The health impacts of particulate matter are especially difficult to study because particles are highly heterogeneous, both in terms of physical properties (e.g. size and density) and chemical composition. At the same time, concentrations are commonly measured with rather coarse metrics, such as total mass or mass of particles under a single specified aerodynamic diameter (usually 2.5 or 10 micrometers). Li et al. (2003) found that ultrafine particles (<0.1 micrometers) were the most potent at inducing cellular heme oxygenase-1 expression (a sensitive marker for oxidative stress) and depleting intracellular glutathione, when compared to fine particles (<2.5 micrometers) and coarse particles (2.5-10 micrometers). Ultrafine particles are able to penetrate tissue more effectively, due to their small size. Electron microscopy revealed subcellular penetration and mitochondrial damage from ultrafine particles, and to a lesser extent, fine particles. Heme oxygenase-1 expression was also correlated with the high organic carbon and polycyclic aromatic hydrocarbon (PAH) content of the ultrafine particles, indicating that their smaller size may not be the only explanation for their more severe biological impacts (Li et al., 2003). Earlier work by Nemmer et al. (2002) found that radioactively labeled ultrafine carbon particles passed rapidly into the bloodstream after inhalation, with detectable levels appearing after 1 minute. These types of findings have led numerous scientists, including the authors of Li et al. (2003) to suggest considering regulatory standards based on particle number instead of mass, in order to increase the focus on smaller particles.

Naturally, there is some disagreement on the health impacts of diesel exhaust, as well as the amount of exposure required to generate the impacts. Some dissenting voices come from the vehicle manufacturing industry. After reviewing animal testing research, Hesterberg et al. (2005) concluded that chronic inhalation of high concentrations of diesel exhaust caused lung tumors in rats but not in mice or Syrian hamsters. The authors argued that the results in rats could be species specific, and therefore not applicable to humans. At the time the review was published, Hesterberg was an employee of International Truck and Engine Corporation. Though Heserberg et al.'s arguments are particularly bold and optimistic, industry is not alone in acknowledging uncertainty exists regarding health effects, especially when it comes to quantifying risks and acceptable concentrations. The EPA concedes its estimates of acute reference concentrations (RfCs) have uncertainty spanning as much as an order of magnitude (US EPA, 2009c). Acute reference concentrations (RfCs) are levels of acute (meaning for 24 hours or less) continuous inhalation exposure that is likely to be without an appreciable risk of deleterious effects in humans during a lifetime.

Overall, it is well established that diesel exhaust and its components can have numerous serious harmful health impacts, but the exact consequences of a given concentration of diesel exhaust or its components remain difficult to accurately quantify.

2.2 Clean Diesel Regulation

2.2.1 New Vehicle Emission Standards

In response to growing concerns over the health impacts of emissions, the U.S. EPA has repeatedly tightened standards for new vehicles. Figures 1 and 2 provide simplified histories of PM and NO_x emission standards for new trucks, as presented in (US EPA, 2003a). The units are grams per brake-horsepower-hour, where brake-horsepower-hours is a unit of energy. Defining standards in this manner allows them to apply to a range of vehicle sizes. Figures 1 and 2 are simplifications because EPA standards are complicated by sophisticated phase in schemes which ease transitions. The 0.2g/bhp-hr NO_x standard, for example, does not take full effect until 2010 (US EPA, 2001). Nonetheless, the general trends in Figures 1and 2 hold true. Between the late 1980s and 2010, both PM and NO_x emission rate caps have been lowered by over 98%.

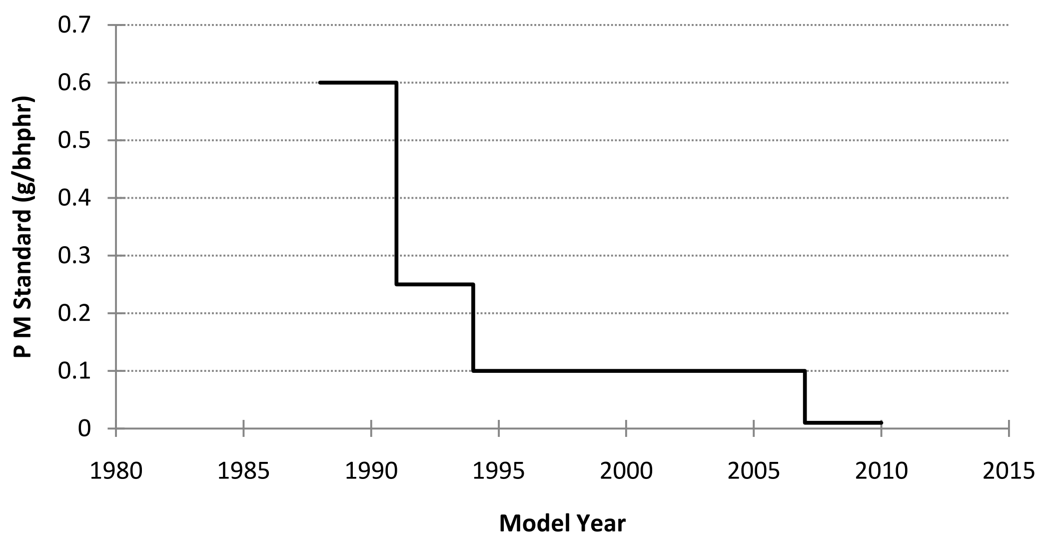

Figure 1. New Truck PM Standard

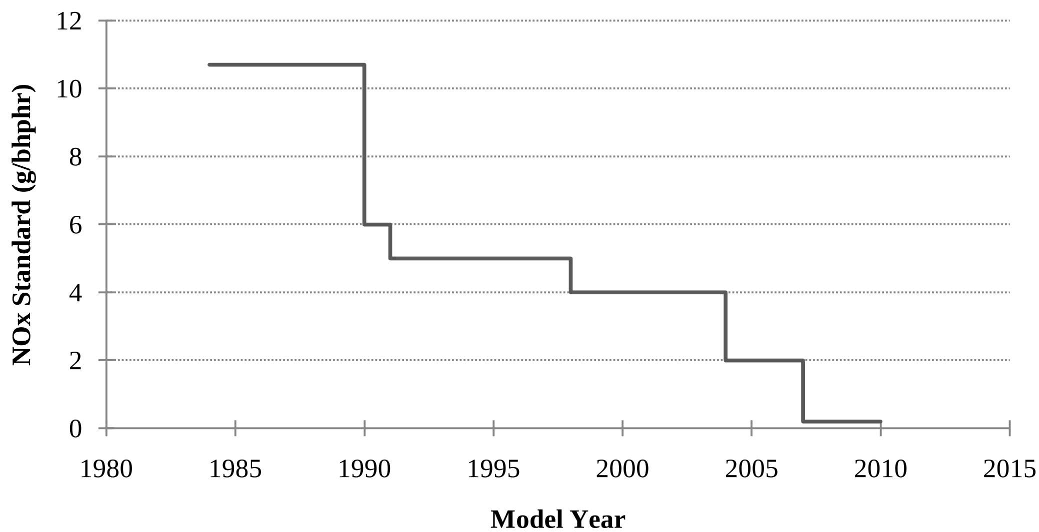

Figure 2. New Truck NO_x Standard

As dramatic as the changes in emission standards have been, their impacts have been slowed by the fact that many older diesel vehicles have remained in use. This fact has prompted interest in retrofitting and early replacement programs.

2.2.2 Federal Funding Programs

The Congestion Mitigation and Air Quality Improvement Program, which is jointly administered by the Federal Highway Administration (FHWA) and the Federal Transit Administration (FTA), has provided billions in authorizations for a wide range of projects, including but not limited to diesel retrofits. It was authorized in 1991, by the Intermodal Surface Transportation Efficiency Act (ISTEA), and continued under the 1998 Transportation Equity Act for the 21st Century (TEA-21) and the 2005 Safe, Accountable, Flexible, and Efficient Transportation Equity Act: A Legacy for Users (SAFETEA-LU), which increased the focus on air quality benefits and gave diesel retrofits higher priority. Funding is available for DOT and Metropolitan Planning

Organization (MPO) projects in nonattainment and maintenance areas, as determined by the National Ambient Air Quality Standards (US FHWA, 2011a). Funds are divided among nonattainment and maintenance areas based on their weighted populations, as computed using a statutory formula (US FHWA, 2011b).

The EPA is a major player in clean diesel funding. Since 2008, the EPA has awarded over 500 grants for diesel emission reduction projects (US EPA, 2011c). At the core of the EPA's efforts is the National Clean Diesel Funding Assistance Program, which is estimated to be $32 million in size in fiscal year 2011 (US EPA, 2011d). It provides grants for a wide range of EPA or CARB verified emission reduction technologies, including add-on emission reduction retrofits, idle reduction technologies, and cleaner fuels, as well as engine repowers and upgrades, and early vehicle replacement. Trucks, buses, trains, boats, and construction equipment can all qualify, but the emissions reductions achieved must not be mandated by federal, state, or local law. Applicants must be regional, state, local, tribal, or port agencies with jurisdiction over transportation or air quality, or transportation/air quality focused nonprofit organizations which provide pollution reduction or educational services. This can include school districts and metropolitan planning organizations. Half of the funds are reserved for public fleets or private fleets contracted or leased for public purpose, such as school buses or refuse haulers (US EPA, 2011e). The American Recovery and Reinvestment Act of 2009 included approximately $156 million in additional funding for the National Clean Diesel Funding Assistance Program, beyond the normal fiscal year 2009 appropriations (US EPA, 2009a).

The EPA administers multiple smaller programs which fund diesel retrofits. The Emerging Technologies Program is estimated to be $4 million in size for fiscal year 2011 (US EPA, 2011d). It funds projects which use technologies from the EPA's Emerging Technologies List. Such technologies are in the process of being verified

(US EPA, 2011f). As of January 12th, 2011, the majority of the list is comprised of urea-based selective catalytic reduction systems (US EPA, 2011g). The types of entities which are eligible to apply are the same as for the National Clean Diesel Funding Assistance Program, but these entities must partner with a device manufacturer (US EPA, 2011f). The Smartway Finance Program is estimated to be $6 million in size for fiscal year 2011. It helps establish low cost loans and financing to help fleets reduce diesel emissions (US EPA, 2011h).

The EPA also provides grants to states to implement their own clean diesel grant/loan programs. For fiscal year 2011, the estimated funding for the State Clean Diesel Grant Program is $18 million (US EPA, 2011d).

2.2.3 State Funding Programs

State programs which fund clean diesel projects range from complex interactions between state and local entities with a variety of funding sources and objectives to basic mechanisms for distributing funds from federal grants. When selecting projects to fund, states weigh a variety of objectives, but they nearly always consider the cost effectiveness, in terms of emissions prevented per dollar spent. Preference is often given to projects with local funding contributions, as well as those in areas with severe local air quality problems. Eligible vehicles are often required to spend some portion of their time operating in the grant-giving state (NDEQ, 2009; ODEQ, 2008; CARB, 2008b).

As is frequently the case in air quality issues, California has been a leader in developing systems for managing the funding clean diesel projects. In 1998, California established the Carl Moyer Memorial Air Quality Standards Attainment Program, which is administered by the California Air Resources Board (CARB). Serving as a key part of California's State Implementation Plan, the program issues grants to fund retrofits and replacements which result in cleaner than required on and off-road engines. In its first seven years, the program provided $170 million in funding to clean 7,500 engines, reducing nitrogen oxide emissions by about 24 tons per day, and particulate matter emissions by about one ton per day. Project selection is largely decentralized, with CARB distributing the bulk of the funds to districts, which have some flexibility to supplement the CARB guidelines with their own objectives. CARB can issue grants directly for multi-district projects, but such grants cannot exceed 10% of the available funds (CARB, 2008c). Local districts with a population exceeding one million are required by legislation to direct 50 percent of the program's funds toward projects in environmental justice areas. Funding sources for the Carl Moyer program are numerous, and include general budget appropriations, the smog check fee, the tire fee, and the motor vehicle registration fee (CARB, 2008c).

The Carl Moyer Program has extensive guidelines which outline quite a few requirements for projects. Retrofits must generally be CARB certified, though in some cases EPA or International Maritime Organization certification is sufficient. In order to be eligible, non-boat vehicles must operate in California for at least 75% of the project life. There are provisions for monitoring and for preventing grantees from obtaining funds from multiple sources totaling more than the cost of the project (including from tax breaks and emission markets). The project must not cost more than $16,000 per weighted ton of NO_x, reactive organic gases, and PM_{10} reductions. PM_{10} is given a weight of 20, while NO_x and reactive organic gases are given weights of one (CARB, 2008b).

The Carl Moyer Program strives to pay the "incremental cost of cleaner technologies," but this number can be challenging to pin down, as it can be difficult to agree on what the alternative would have cost (or even what the alternative would have been). New guidelines produced in 2008 define the "incremental cost" as a percentage of the total project cost. This is much simpler than establishing a baseline cost for comparison (e.g. by obtaining a price quote for an engine rebuild in the case of engine repower projects or the purchase price of a conventional vehicle in the case of new clean vehicle purchase projects). For on-road vehicle projects, percentages are 100% for retrofits, 80% for repowering, and 25% for new vehicle purchase. For off-road projects, the maximum percentage is still 100% for retrofits, but the percentage for repowering depends of the emissions tier (CARB, 2008c).

In New York, NYSERDA administers the Clean Air School Bus Program, which is funded by the state's 1996 Clean Water-Clean Air Bond Act. In the first round it awarded $5 million to retrofit 2,200 school buses in 74 school districts. Awards can cover the entire cost of retrofits, including purchase and installation (NYSERDA, 2004a, 2007a). Retrofits must be EPA certified. School districts, municipalities and state agencies can all apply. Preference is given to projects which are cost effective, support emerging technologies, include co-funding, and are in areas in need of air quality improvement (NYSERDA, 2004b). For example, New York City's ozone and PM_{10} (in Manhattan) nonattainment status were cited as reasons why $1.25 million was awarded to the NYC Department of Education to retrofit 130 school buses with DPFs (NYSERDA, 2004a). Other grants funded DOCs, crankcase filters, new CNG buses, and CNG refueling stations, with co-funding coming from local and federal government sources, as well as a private energy company (NYSERDA, 2007a; 2007b).

Other state programs vary widely. Some are narrow in focus, perhaps applying only to school buses (KDAQ, 2008) or to projects which reduce a particular pollutant (TCEQ, 2009a), while other programs are much broader. Grants and tax credits are available (ODEQ, 2008), as are loans (UDEQ, 2009). Small retrofit grants can be linked with environmental education projects (Kansas Green Schools, 2009).

2.2.4 Regulatory Mandates

California has been a leader in developing clean diesel regulatory mandates, in addition to funding mechanisms. California has a broad public fleet rule which applies to nearly any city, county, public agency or utility that owns, leases, or operates on-road diesel vehicles from model years 1960-2006 over 14,000 lbs GVW. There are exemptions for military and emergency vehicles, as well as school buses, urban buses and garbage trucks. Non-exempted vehicles must apply the "best available control technology" to a growing percentage of their fleet by a sequence of deadlines. For most counties, the deadlines are December 31st of 2007, 2009, and 2011, but low population counties can use later deadlines. "Best available control technology" can mean an engine certified to 0.01g/bhp-hr PM or an engine with the highest level emission control strategy verified for that engine (CARB, 2006).

California's fleet rule has been the subject of some controversy. The implementation costs drew criticism from the Legislative Analyst's Office (LAO), California's nonpartisan fiscal and policy advisor for the legislature. CalTrans estimated it would cost $260 million to comply with four sets of state air quality regulations. Ninety percent of these costs were due to CARB's on-road and non-road diesel regulations. These estimates assume filters cost about $20,000 each (LAO, 2009). CalTrans told the LAO that it is unable to fit the filters onto some of its trucks, necessitating further modifications to the vehicles. CARB insists that no such modifications are required to be in compliance. Exemptions are available in some cases, but must be filed for each individual vehicle, unless CARB reevaluates its regulations. CARB advised the LAO that it may reevaluate its regulations and issue across-the-board exemptions (LAO, 2009). CARB's actions can serve as important precedent in other states and cities which develop similar regulation.

CARB has been considering expanding its regulatory mandate program to apply to any person, business, school district, or federal government agency which owns or operates covered vehicles in California. A wide range of diesel vehicles with gross vehicle weight over 14,000lbs would be covered, regardless of where they are registered, though exceptions would be included for authorized emergency vehicles, and a few other categories (CARB, 2008d). As of December 2010, CARB is considering a revised timeline which would require installation of PM reducing retrofits starting by January 1st 2012, and replacement of older trucks starting January 1st, 2015 (CARB, 2011).

New York City developed a set of laws regarding diesel emissions from public fleets and fleets operating under contract with the city. These local laws are quite similar to each other, but often apply only to vehicles filling specific functions. The general approach of the regulation mirrors that used in California's fleet rule.

NYC Local Law 77, adopted in 2003, requires the use of ULSD and "best available technology" for off-road diesel vehicles owned by the city, as well as private equipment used on city construction projects (M.J. Bradley & Associates & Gruzen Samton, 2004). NYC Local Law 39, adopted in 2005, requires that diesel vehicles owned or operated by city agencies use ULSD. Of these vehicles, those with a gross vehicle weight over 8,500 lbs must either use an engine certified as meeting the 2007 EPA particulate matter standard (0.01 g/bhp-hr) or use the "best available retrofit technology". A phase in scheme requires increasing percentages of vehicles to meet the mandate starting with 7% by January 1, 2007 and ending with 100% by July 1, 2012 (NYCC, 2005a). NYC Local Law 40, adopted in 2005, requires that solid waste or recyclable material contracts specify that diesel vehicles used to perform said contracts, and which operate primarily within NYC, use ULSD. All such vehicles must also use the "best available retrofit technology" by March 1, 2006 (NYCC, 2005b). NYC Local Law 42, adopted in 2005, requires the use of ULSD and "best available retrofit technology" for school buses serving public schools, even if the buses are not owned by the city (NYCC, 2005c). A portion of this dissertation research was funded by New York Metropolitan Transportation Council to evaluate cost effective means of complying with NYC school bus regulation.

New York City Department of Environmental Protection tracks compliance with NYC retrofit laws. Their fiscal year 2009 report, the most recent report which could be obtained by Freedom of Information Act request, counts more than 3,300 retrofits including diesel oxidation catalysts, diesel particulate filters, and crankcase filters.

Several counties surrounding New York City passed their own clean diesel mandates, as did New York State (Nassau, 2009; Rockland, 2006; Westchester, 2010). The three suburban county mandates apply to a relatively small number of vehicles, primarily those owned or operated by the counties. The statewide regulation is much larger in scale, and will be the focus of the case studies presented in this dissertation. New York State Department of Transportation, which operates a large fleet covered by the regulation, funded a portion of the dissertation research.

New York State's clean diesel mandates are a result of the Diesel Emissions Reduction Act of 2006. Section 19-0323 (L. 2006, c.621) of New York's Environmental Conservation Law (ECL) mandates that NYS Department of Environmental Conservation (NYS DEC) promulgate regulations requiring the use of ULSD and Best Available Retrofit Technology (BART) for various state agency, state public authority, and regional public authority heavy duty vehicles, as well as heavy duty vehicles used on behalf of such agencies and authorities (NYS DEC, 2009a). The requirements of NYS DEC regulations naturally play a major role in determining which strategies fleet managers can pursue for reducing emissions. The remainder of this subsection provides a summary of 6 New York Codes, Rules, and Regulations (NYCRR) Part 248, which was adopted with attendant amendments to Part 200 in the summer of 2009 (NYS DEC, 2009b).

Part 248 requirements apply to all covered heavy duty vehicles owned by, leased by, operated by, or on behalf of regulated entities. Switching from an owning model to a leasing model, would therefore not allow regulated entities to avoid the requirements. Heavy duty vehicles include both on-road and off-road diesel vehicles with GVW over 8,500lbs. There are, however, numerous exceptions.

Authorized emergency vehicles are not included (NYS DEC, 2009c). Authorized emergency vehicles are defined by section 101 of New York Vehicle and Traffic Law (VTL). The definition includes “environmental emergency response vehicle” and “sanitation patrol vehicle.” Section 115-d specifies that environmental emergency response vehicles must be responding to the release, spill or leak of a hazardous substance, and section 141-a specifies that sanitation patrol vehicles only include those operated by the sanitation police of the NYC Department of Sanitation.

There is another category of exemptions which does appear to apply to the NYS DOT, however. Heavy duty vehicles does not include “road rollers, tractor cranes, truck cranes, power shovels, road building machines, snow plows, road sweepers, sand spreaders, … earth movers, which shall mean motor-driven vehicles in excess of eight feet in width equipped with pneumatic tires designed and constructed for moving or transporting earth and rock in connection with excavation and grading work.” (NYS DEC, 2009c) Slightly later, “farm type tractors and all terrain type vehicles used exclusively for agriculture or mowing purposes, or for snow plowing” are listed as not being included in heavy duty vehicles (NYS DEC, 2009c). These clauses appear to exclude many on-road and off-road NYS DOT vehicles. The exclusion of street sweepers might be especially reassuring, as NYS DOT staff have expressed concern that these vehicles might prove especially challenging to retrofit. As of January 1st, 2008, NYS DOT operated 35 street sweepers.

The exemption of snow plows would be of paramount importance to NYS DOT, if could be broadly applied. Well over a thousand NYS DOT diesel vehicles function as snow plows in the winter. It is not entirely clear from the regulation whether a truck which operates as a snow plow for a portion of the year and in another function in the summer would be considered a snow plow. Currently, roughly 30% of NYS DOT large dump trucks are used exclusively for snow and ice removal. NYS DEC did not respond to requests to clarify the definition of "snow plows," but NYS DOT staff indicated that only vehicles which can only be used for snow removal are exempt, such as snow blowers. This means that thousands of NYS DOT trucks are covered by the regulation.

Vehicles using an engine certified to meet the 2007 EPA PM standard (0.01 g/bhp-hr) are considered to be in compliance. Any heavy duty vehicle which has been retrofitted with an EPA or CARB approved conversion kit to enable it to run on a combination of CNG and ULSD is considered to be in compliance. Heavy duty vehicles retrofitted with an EPA or CARB verified device prior to February 12, 2007 are considered in compliance, so long as the device is maintained for the rest of the vehicle's life (NYS DEC, 2009c).

All covered vehicles owned by, leased by, operated by, or on behalf of regulated entities must use ULSD when running on diesel, unless a waiver is approved. Given the broad availability of ULSD resulting from EPA mandates and ULSD's relatively low marginal cost, this requirement should be relatively easy to satisfy.

Covered heavy duty diesel vehicles must also use the Best Available Retrofit Technology (BART). Originally, this requirement was phased according to the following schedule:

1. at least 33 percent of all such vehicles use BART by December 31, 2008
2. at least 66 percent of all such vehicles use BART by December 31, 2009
3. all such vehicles use BART by December 31, 2010.

The schedule was later changed substantially. The first two deadlines were dropped entirely, and the 100% BART requirement was delayed until December 31, 2012 (NYS DEC, 2011). This information was not available when many retrofit and replacement decisions were made. As a result, case studies are based on the original timeline.

There are three basic types of approaches for dealing with a vehicle which is not in compliance. First, the vehicle can simply be retired (and possibly replaced with a compliant vehicle). It appears that the vehicle can be kept for spare parts if the engine is removed. Second, the vehicle can be repowered. The vehicle will become compliant if its engine is replaced with an engine certified to meet the 2007 EPA PM standard (0.01 g/bhp-hr). It will also become compliant if its engine is replaced with one which operates on an approved alternative fuel, provided that model year 2004-2006 alternative fuel engines are certified to the optional, reduced emissions standards specified in title 13, California Code of Regulations, section 1956.8(a)(2)(A) (see Table 1, Section 200.9 of this Title). Approved alternative fuels include natural gas, propane, ethanol, methanol, gasoline (when used in hybrid electric vehicles only), hydrogen, electricity, fuel cells, or advanced technologies that do not rely solely on diesel fuel or a diesel/non-diesel fuel mixture. Note that bi-fuel CNG vehicles which use diesel as well as CNG are not on the list of alternative fuel vehicles, but they are considered in compliance if they were converted with an EPA or CARB approved conversion kit (as were NYS DOT's trial conversions). Third, the vehicle can go through the BART evaluation and selection process described in Part 248 (NYS DEC, 2009c).

All BART must be either EPA or CARB verified. There are three levels, corresponding to the three levels of CARB PM reduction verification. Level 1 means >=25% reduction; level 2 means >=50% reduction; level 3 means >=85% reduction or <=0.01g/bhp-hr. Regulated entities are expected to consider level 3 retrofits first. Level 2 retrofits can only be considered if no level 3 retrofits are compatible with the vehicle and application. Similarly, level 1 retrofits can only be considered if no level 2 or 3 retrofits are compatible with the vehicle and application. If multiple products of the same level are compatible with the vehicle and application, the product which offers the greatest NO_x reduction must be selected if it is less than or equal to 30% more expensive than the other options of the same level. None of the products selected can result in an increase in NO_x. The regulated entity can apply for a waiver of BART requirements if a vehicle and application are incompatible with all technologies of all levels (NYS DEC, 2009c).

Any vehicles subject to the ULSD or BART requirements, and to the consent decree, must have approved low NO_x rebuild kits installed prior to any BART device (NYS DEC, 2009c). This requirement is connected with a 1998 court settlement. In the mid-1990s, the U.S. Department of Justice (US DOJ), EPA, and CARB discovered that seven major engine manufacturers had designed model year 1993-1998 heavy duty diesel engines to operate differently when cruising steadily, as opposed to when speed patterns resembled emissions testing duty cycles. The steady highway cruising mode of operation improved fuel economy, but also caused excessive NO_x emissions. The resulting court settlement required the manufacturers to provide dealers with modified software, called low NO_x rebuild kits or chip reflash. These kits must be installed free of charge when conducting an engine rebuild, or upon owner/operator request (CTC & Associates, 2006). The companies were Cummins, Volvo Truck, Detroit Diesel, Mack Trucks, Caterpillar, Navistar International and Renault Vehicules Industriels (US DOJ, 1998).

Part 248 also has numerous record keeping, reporting and labeling requirements, such as annual inventories and low NO_x rebuild labels on engines which had low NO_x rebuild kits installed (NYS DEC, 2009c).

2.3 Retrofit Technologies

Emissions can be produced from both the crankcase and the tailpipe of a diesel vehicle, and retrofits are available for both. This section will outline the basics of how the technologies function, without going into the details of differences between brands and models. Specific models, and their applications to the NYS DOT fleet, are examined in Chapter 3.

Crankcase emissions are caused by smoke excretions necessary to eliminate high pressure buildup in the crankcase (Hill et al., 2005). These excretions release fine particulate matter (Hill et al., 2005). According to Cummins, crankcase emissions can constitute up to 25% of total emissions (Cummins, 2009a). Closed crankcase ventilation systems (CCVS) are used to reduce these emissions by rerouting engine blow-by back into the engine intake, filtering out particulate matter, and recombusting air toxics. Emissions which are not combusted on second pass through the engine can still be treated by an emissions control device in the exhaust system. Both the EPA and CARB verify CCVS only when used with an emission control device in the exhaust system. EPA grant funds cannot be used to install a CCVS by itself (US EPA, 2009d).

Once emissions have reached the exhaust system, two major types of reduction technologies are used: diesel particulate filters (DPFs) and diesel oxidation catalysts (DOCs). Both work by oxidizing hazardous gases and particulates into less harmful chemical compounds, but the particulate filter includes a physical ceramic filter (US EPA, 2007a). Some devices use less intense filters than others. This results in a spectrum of retrofit devices, rather than two strictly distinct categories. Devices with less intense filters are sometimes referred to as "flow through filters."

Diesel oxidation catalysts are fairly broadly compatible. Their exhaust temperature requirements are relatively low, typically around 150˚C (US EPA, 2007b). The primary disadvantage of DOCs is that much of the particulate mass flows through. As more intense filters are added, a much larger fraction of the particulate mass is stopped, but this comes with the challenge of disposing of the buildup. Regenerating filters burn off the particulates stopped by the filters, but they can be sensitive to temperature. If temperatures are too low to support regeneration for a long period, the buildup can burn at too high a temperature when finally ignited. The resulting temperature gradients can be damaging to the exhaust system (van Setten, 2001). A typical passive diesel particulate filter might require that the exhaust temperature be at least 240˚C at the filter inlet for 40% of the duty cycle (US EPA, 2007c). Active diesel particulate filters attempt to solve the temperature problem by providing additional heat for regeneration, at additional cost. Huss, a filter manufacturer whose active filters are being used in a trial project on NYC school buses, claims that some of its filters have no minimum exhaust temperature (Huss, 2008). Using ultra-low-sulfur diesel fuel tends to help with buildup problems, and is typically required for particle filters. Though not required for diesel oxidation catalysts, the EPA believes ultra-low-sulfur diesel can prevent diesel oxidation catalysts from increasing ultrafine particulate matter (US EPA, 2007d). ULSD can also extend the useful life of all retrofits, by reducing levels of sulfuric acid in the exhaust system (M DEP, 2008).

Catalysts can cause soot to oxidize through direct physical contact, or they can catalyze the formation of a gaseous molecule (such as NO_2 from NO and O_2), which is more reactive than O_2 itself. The reactive gas molecule (e.g. NO_2) can then oxidize soot, or other gaseous molecules such as CO. In doing so, NO_2 would revert to NO, and the process could repeat in a cycle (van Setten, 2001). Concern over the potential for catalyzed retrofits to increase NO_2 emissions led the EPA to issue limits on NO_2 increases for all retrofits on its verified technology list, following a similar move by CARB (US EPA, 2007e). As of January 1st, 2009, all retrofits on the EPA's verified technology list must not increase NO_2 emissions by more than 20 percent (US EPA, 2007e). Increased NO_2 emissions, sometimes referred to as NO_2 slip, can pose problems because NO_2 is a potent oxidizer (NJ DEP, 2006). Rim et al. (2008) found the addition of a diesel oxidation catalyst did not appear to increase the in-cabin NO_2 concentration. Another study tested tailpipe nitrogen oxides when high performance diesel oxidation catalysts (also referred to as "flow-through filters") were deployed on school buses in New Jersey. No significant overall reduction in NO_x was found (none was expected), and post-retrofit tests actually yielded a slightly lower NO_2 to NO ratio (NJ DEP, 2006). Retrofits sometimes employ multiple stages of catalysts. In particular, oxidation catalysts can be employed upstream of NO_2 reduction catalysts (Johnson Matthey, 2009).

The degrees to which emission control technologies prevent PM, CO, and hydrocarbon emissions have been studied by numerous authors. Hill et al. (2005) found crankcase emissions to be an extremely strong source of $PM_{2.5}$ inside a school bus, and that a Donaldson Spiracle closed-crankcase filtration device eliminated this form of self-pollution. Clark et al. (2002) reported that tests on a 1992 model year refuse truck revealed a PM reduction of 24% and a CO reduction of 8.3%, due to a catalyst. Herner et al. (2009) tested four heavy-duty and medium-duty diesel vehicles with four different particulate filters. They found the filters realized more than 95% PM mass reduction on both duty cycles. They found a catalyzed filter removed 99% of hydrocarbons and 94% of CO, but an uncatalyzed filter did not produce such benefits. Low exhaust temperatures reduced the effectiveness of the catalyzed filter at controlling hydrocarbons, and nearly eliminated its ability to control CO.

There has been some concern that emission control technologies which reduce PM mass may actually increase nanoparticle number concentrations (Holmén and Ayala, 2002). Kittleson et al. (2006) found that one type of particulate filter produced large quantities of nuclei mode particles, while another did not. For the filter which produced the large quantities, the number increased with higher exhaust temperatures. Biswas et al. (2008) found that two types of particulate filters efficiently suppressed nucleation mode particles. They hypothesized that the young age of a filter could contribute to its ability to store sulfur. Holmén and Ayala (2002) found that a passive DPF yielded both accumulation and nuclei mode number concentrations of particles which were lower than those from the same vehicle using an oxidation catalyst by a factor of 10-100, under most test conditions. Nuclei mode particles are smaller than accumulation mode particles. Nuclei mode particles tend to make up a large percentage of the number of particles in diesel exhaust, while constituting a very low fraction of the total mass (Kittleson et al., 1998).

It can be difficult to develop a single percent effectiveness for a given technology at reducing a given pollutant (even when considering installation on only one potential vehicle). This is largely due to the effect of the duty cycle. Emission rates vary with the engine load. Measuring particle emissions can be particularly challenging, as the manner in which exhaust is diluted after leaving the tailpipe can influence results (Holmén and Ayala, 2002).

For practical purposes, the verified technology lists produced by the EPA and CARB are good sources of product specific emission reduction estimates. EPA numbers are based on engine dynamometer tests (US EPA, 2002a). CARB's verification procedure allows the applicant limited flexibility in choosing between engine and chassis dynamometer testing, depending on what kind of reduction they are seeking to verify (CARB, 2009a).

In order to meet new vehicle emission standards, manufacturers have been incorporating DPFs into their new models, both in the United States and other countries such as Japan. In the United States, the 2007 standard for particulate matter was a driving force behind the change. Selective Catalytic Reduction (SCR) is being used to curtail NO_x emissions. SCR works by reducing NO_x on a selective catalyst using an ammonia reductant (commonly from urea). The systems are often quite complex, partially in order to withstand freezing and thawing of the urea. On the positive side, highly efficient NO_x removal allows engines to operate at maximum fuel efficiency (high engine-out NO_x, low PM) (Johnson, 2009). As a result of these changes to new vehicle designs, replacing a vehicle can lead to lower emissions than retrofits.

Several alternative means of emission reduction, primarily alternative fuels, were described and evaluated in Gao and Stasko (2010). For the sake of brevity, they are not described here. For various reasons, they do not play major roles in the case studies.

CHAPTER 3

COMPATIBILITY AND COST-BENEFIT ANALYSIS

3.1 Compatibility

3.1.1 Introduction

Before the optimal retrofit and replacement strategy can be formulated, the feasible region must be defined. Retrofit technologies must be assessed to determine their compatibility with the vehicles in the fleet, as well as whether they will bring the vehicles into compliance with any relevant regulation. This section will outline how that process was conducted for various vehicles in the NYS DOT heavy duty diesel dump truck fleet. For a more detailed description of the process, see Gao and Stasko (2010).

A relatively small number of types of dump trucks make up a large fraction of the NYS DOT diesel fleet. The top ten vehicle types in the January 1st, 2008 snapshot provided by NYS DOT made up roughly 50% of the fleet, with more than 100 of each present. The top 20 made up more than 75% of the fleet. Common vehicles are listed in Table 2, along with basic characteristics. They made up over 80% of the fleet.

Table 1. Common Vehicles in the NYS DOT Diesel Fleet

Count	Class	Manufacturer	Model	Year	Typical Uses	EGR
29	8	International	7600	2003	Large Dump/Spread	N
83	8	International	7600	2004	Large Dump/Spread	Y
114	8	International	7600	2005	Large Dump/Spread	Y
24	8	International	2574	1993	Large Dump Truck	N
84	8	International	2574	1995	Large Dump Truck	N
12	8	International	2574	1996	Large Dump Truck	N
83	8	International	2574	1997	Large Dump Truck	N
197	8	International	2574	1998	Large Dump Truck	N
102	8	International	2574	1999	Large Dump Truck	N

113	8	International	2574	2000	Large Dump Truck	N
135	8	International	2574	2002	Large Dump Truck	N
79	8	International	2574	2003	Large Dump Truck	N
23	8	Mack	CV713	2006	Large Dump Truck	Y
153	8	Mack	CV713	2007	Large Dump Truck	Y
107	8	Mack	CV712	2006	Large Dump Truck	Y
72	8	Mack	CV712	2007	Large Dump Truck	Y
105	6	Ford	F650	2007	Dump and Stake	Y
39	6	International	4700	1996	Stake and Dump	N
51	6	International	4700	1997	Stake, Dump, Sweeper	N
67	6	International	4700	1998	Stake, Dump, Sweeper	N
145	6	International	4700	2000	Stake and Dump	N
116	6	International	4700	2002	Stake and Dump	N
38	6	International	4600	1990	Stake and Dump	N
43	6	International	4600	1992	Stake and Dump	N
95	6	International	4600	1994	Stake and Dump	N
51	6	International	4200	2004	Stake and Dump	Y
36	6	International	1654	1989	Stake Truck	N
2196						

The dump trucks are on-road vehicles, making them incompatible with retrofit technologies designed solely for non-road vehicles or stationary engines. These technologies were screened out, as were technologies which were no longer available for sale in the United States. The remaining technologies were examined in further depth.

Compatibility requirements are listed both by retrofit device manufacturers and government bodies which verify retrofits. CARB can be particularly detailed in its descriptions of appropriate vehicle types. In order for a retrofit to be verified for use on a particular vehicle, the vehicle may have to meet requirements regarding its model year, horsepower, displacement, use of exhaust gas recirculation (EGR), number of strokes, gross vehicle weight (GVW), exhaust temperature profile, original PM emission certification, whether it is turbocharged or naturally aspirated and whether it has already been retrofitted. Additionally, CARB provides a list of EPA engine family names for which retrofits are verified (CARB, 2009b, 2009c, 2009d). EPA engine family names are 12 character codes. The first nine characters can be determined by the model year, manufacturer, family type (e.g. H for heavy duty) and displacement, while the remainder ensure uniqueness of an engine family.

While numerous efforts were made to collect all relevant data from NYS DOT, vehicle manufacturers, and federal regulators, some data was unavailable while other data was contradictory. As a result, the compatibility analysis in this chapter is optimistic in the sense that it only discusses incompatibilities revealed by the available data. An effort is made to indicate which incompatibilities have particularly high uncertainty.

3.1.2 Non-Duty-Cycle Requirements

This section examines compatibility with requirements that are independent of how the vehicle is used. It does not discuss exhaust temperature profile or regeneration requirements which are related to duty cycles. Duty cycle requirements are discussed in the next section.

Table 2 lists incompatibilities between the pre-2007 vehicles in Table 1 and level 3 active filters, while Table 3 lists incompatibilities with level 3 passive filters and Table 4 lists incompatibilities with level 2 flow through filters as well as level 1 diesel oxidation catalysts. Only verifications for sufficient PM reductions to meet these levels have been included. It is important to note these tables only list incompatibilities found with the available data, and that if complete data were available on all vehicles there would likely be more incompatibilities. As noted, several retrofit technologies are associated with other engine modifications such as crankcase filters and exhaust gas recirculation.

Table 2. Non-Duty-Cycle Incompatibilities Found Between Common Vehicles and Level III Active Filter Technologies

					Level III					
					Active Filters					
Count	Class	Manufacturer	Model	Year	Cleaire Horizon	Cleaire Vista	ESW Canada ThermaCat	Donaldson Semi-Active Electric Filter (SEF)	Engine Control Systems Purifilter Plus	HUSS Umwelttechnik FS-MK
29	8	International	7600	2003			HP			
83	8	International	7600	2004		EGR	EGR	EGR	EGR	EGR
114	8	International	7600	2005		EGR	EGR	EGR	EGR	EGR
24	8	International	2574	1993			HP			
84	8	International	2574	1995			HP			
12	8	International	2574	1996			HP			
83	8	International	2574	1997			HP			
197	8	International	2574	1998		too high PM*	HP		too high PM*	
102	8	International	2574	1999			HP			
113	8	International	2574	2000			HP			
135	8	International	2574	2002			HP			
79	8	International	2574	2003			HP			
23	8	Mack	CV713	2006		EGR	EGR	EGR	EGR	EGR
107	8	Mack	CV712	2006		EGR	EGR	EGR	EGR	EGR
39	6	International	4700	1996						
51	6	International	4700	1997						
67	6	International	4700	1998						
145	6	International	4700	2000						
116	6	International	4700	2002						
38	6	International	4600	1990		year	year	year	year	
43	6	International	4600	1992		year	year		year	
95	6	International	4600	1994						
51	6	International	4200	2004		EGR	EGR	EGR	EGR	EGR
36	6	International	1654	1989		year	year	year	year	

* dependent on engine family name which is uncertain

Table 3. Non-Duty-Cycle Incompatibilities Found Between Common Vehicles and Level III Passive Filter Technologies

					Level III						
					Passive Filters						
Count	Class	Manufacturer	Model	Year	Cleaire Longview	Donaldson Low NO2 Filter (LNF)	Donaldson Low NOx (LXF) Muffler	Engine Control System Purifilter	Johnson Matthey ACCRT	Johnson Matthey CRT reformulated	Johnson Matthey EGRT^
29	8	International	7600	2003							year
83	8	International	7600	2004	EGR					EGR	year
114	8	International	7600	2005	EGR					EGR	year
24	8	International	2574	1993			year		year	year	year
84	8	International	2574	1995			year		year		year
12	8	International	2574	1996			year		year		year
83	8	International	2574	1997			year		year		year
197	8	International	2574	1998	too high PM*	too high PM*	year	too high PM*	year	too high PM*	
102	8	International	2574	1999			year		year		
113	8	International	2574	2000			year		year		
135	8	International	2574	2002							
79	8	International	2574	2003							year
23	8	Mack	CV713	2006	EGR					EGR	year
107	8	Mack	CV712	2006	EGR					EGR	year
39	6	International	4700	1996			year		year		year
51	6	International	4700	1997			year		year		year
67	6	International	4700	1998			year		year		
145	6	International	4700	2000			year		year		
116	6	International	4700	2002							
38	6	International	4600	1990	year	year	year	year	year	year	year
43	6	International	4600	1992	year	year	year	year	year	year	year
95	6	International	4600	1994			year		year		year
51	6	International	4200	2004	EGR					EGR	year
36	6	International	1654	1989	year	year	year	year	year	year	year

* dependent on engine family name which is uncertain

^ includes EGR as well as filter

Table 4. Incompatibilities Found Between Common Vehicles and Level II and I Retrofit Technologies

					Level II	Level I					
					Flow Through Filter	Diesel Oxidation Catalysts					
Count	Class	Manufacturer	Model	Year	DFM diesel multi-stage filter (DMF)	Donaldson DCM 6000	Donaldson 6000 + Spiracle~	Donaldson DCM 6100	Donaldson DCM 6100 + Spiracle~	Engine Control System AZ Purifier & Purifmuffler	Cummins Filtration DOC and CCV System~
29	8	International	7600	2003	year	year	year	year	year		
83	8	International	7600	2004	year	year	year	year	year	year	year
114	8	International	7600	2005	year	year	year	year	year	year	year
24	8	International	2574	1993		year		year			
84	8	International	2574	1995		year					
12	8	International	2574	1996		year					
83	8	International	2574	1997		year					
197	8	International	2574	1998		year					
102	8	International	2574	1999		year					
113	8	International	2574	2000		year					
135	8	International	2574	2002		year					
79	8	International	2574	2003	year	year	year	year	year		
23	8	Mack	CV713	2006	year	year	year	year	year	year	year
107	8	Mack	CV712	2006	year	year	year	year	year	year	year
39	6	International	4700	1996		year					
51	6	International	4700	1997		year					
67	6	International	4700	1998		year					
145	6	International	4700	2000		year					
116	6	International	4700	2002		year					
38	6	International	4600	1990	year			year	year	year	year
43	6	International	4600	1992		year		year			
95	6	International	4600	1994		year					
51	6	International	4200	2004	year	year	year	year	year	year	year
36	6	International	1654	1989	year			year	year	year	year

~ includes crankcase filter as well as DOC

3.1.3 Duty-Cycle Requirements

The duty cycles of NYS DOT vehicles are quite heterogeneous. The region in which a vehicle is used can play a major role. International 2574 trucks in regions 1 and 2 average roughly 14,200 miles/year and 15,700 miles/year respectively, while the same kind of truck in regions 8 and 10 averages 7,100 miles/year and 4,100 miles/year respectively. As Figure 3 shows, regions 8 and 10 contain a large amount of fairly dense NYC suburbs while regions 1 and 2 contain more rural upstate areas. Climate differences might also contribute to the discrepancy. Despite the clear regional trends, there is considerable heterogeneity among vehicles in the same region. Some of this heterogeneity is due to older vehicles experiencing lower usage, but even among vehicles of the same type and model year in the same region considerable heterogeneity remains. To complicate matters, vehicles can change region.

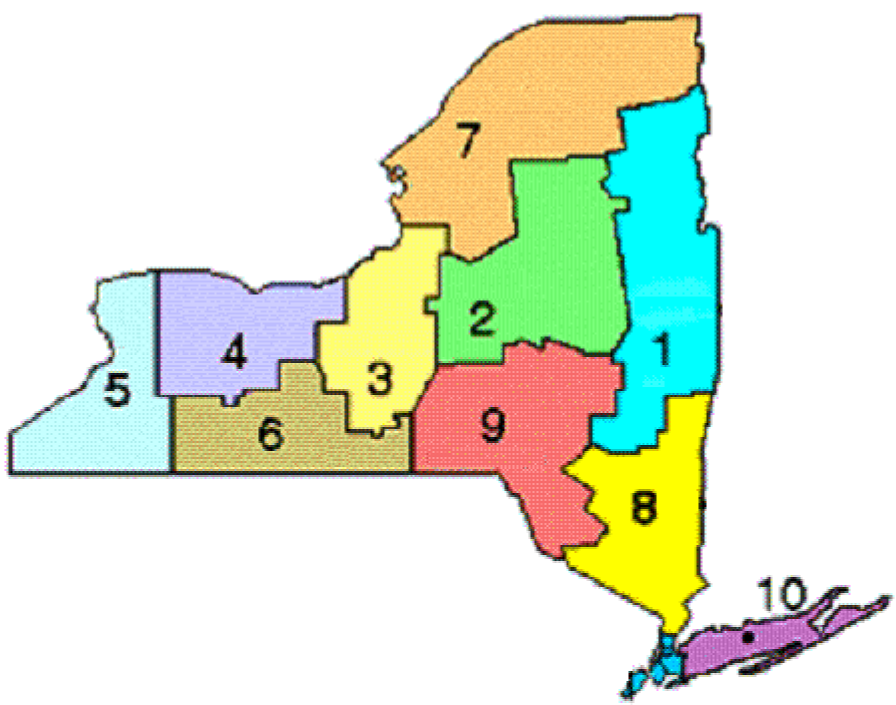

Figure 3. Map of NYS DOT Regions

There is also a seasonal component to the duty cycle of some vehicles. Heavy duty dump trucks are largely used as snow plows in the winter, which involves considerable highway driving at relatively steady speeds. Many of the same trucks are used for work zone protection in the summer, which involves idling for extended periods (sometimes for over 6 hour stretches) in order to keep lights flashing without depleting the batteries.

While mileage data can illustrate some trends in how vehicles are used, it does not necessarily correlate with exhaust temperature profile, which is more important for retrofit compatibility. NYS DOT contracted with Cummins to have exhaust temperature data logging devices installed on numerous large dump trucks in all ten regions. The vehicles included International 2574 model years 2000-2003 and International 7600 model years 2003-2005. NYS DOT wanted to be able to collect data over a much longer period than the standard three day process, which required devices with relatively high storage capacity. Exhaust air temperature was recorded every 5 minutes. The first set of data logs was compiled in winter, and the second was compiled in the late winter and early summer.

The resulting temperature profiles from 76 tested vehicles were checked for compatibility with each of 9 technologies with substantial temperature requirements that were not eliminated for other reasons. Table 5 summarizes the requirements and the results of the 684 checks. Active DPFs without significant temperature requirements were not included in the table. DOCs can have minimum temperature requirements, but they are generally so relaxed (e.g. 100 C) that they cause no concern, so DOCs were left out of the table as well.

Table 5. Temperature Profile Compatibility with On-Road Verified Retrofits*

Technology	Level	Type	Temperature Requirement	% incompatible
ESW Canada ThermaCat	3	Active DPF	>=210 C for >=15% cycle	5.3%
Cleaire Longview	3	Passive DPF	260° C for at least 25% time	55.3%
Donaldson Low NO2 Filter (LNF)	3	Passive DPF	>=235 C for >=40 % cycle, or >=300 C for >=10 % cycle, or average >=237 C	56.6%
Donaldson Low NOx (LXF) Muffler	3	Passive DPF	>=245 C for >=40% cycle, or 310 C for >=10% cycle, or average >= 263 C	68.4%
Engine Control System Purifilter	3	Passive DPF	>=282 C for>=25 % cycle	85.5%
Johnson Matthey ACCRT	3	Passive DPF	>=240 C for >=40% cycle	84.2%
Johnson Matthey CRT reformulated	3	Passive DPF	>=240 C for >=40% cycle	84.2%
Johnson Matthey EGRT	3	Passive DPF/EGR	>=260 C for >=40% cycle	93.4%
Donaldson DFM diesel multi-stage filter (DMF)	2	Flow Through Filter	1991 - 1993 engines: >=230 C for >=40 % cycle and average >=215 C, 1994 - 2002 engines: >=210 C for >=40 % cycle and average >=210 C	79.6%

* Technologies without significant minimum exhaust temperature requirements are not listed.

The final column of Table 5 is the percentage of tested vehicles which had a temperature profile incompatible with the retrofit. The higher this percentage, the less likely this retrofit is to be compatible with vehicles like those tested. This percentage should definitely not be interpreted to mean that all other vehicles are compatible with the retrofit. At least three aspects of the testing approach will tend to make retrofits appear more compatible than they are.

First, the testing used to compile Table 5 was conducted during periods of relatively high activity (large dump trucks tested nearly exclusively in the winter). Compatibility is quite likely to be lower when idling around construction in the summer. A small amount of summer testing was conducted. Temperature profiles from a class 8 truck collected on each of two 6-hour summer days were incompatible with all passive filters, and exhibited a median temperature of only about 150 C. A 1-hour test yielded a slightly higher temperature profile more in line with winter results, but it is unclear whether the truck idled significantly (or at all) during this test.

Second, recall that exhaust temperature profile testing is usually done over shorter intervals than were used here. Even if there are enough high temperatures in the duty cycle, they may not be spaced at regular enough intervals for proper regeneration.

Third, vehicles were only tested in one year of their life. Vehicles may be compatible at one stage of their life, but not later on as their usage profile changes. Vehicles may not be assigned to the same route, location, or even region throughout their life, as they are statewide assets and need to be utilized in response to changing demands, including emergency needs.

Even at first glance, these test results provide fairly strong evidence that four of the seven passive DPF technologies are incompatible with the exhaust temperature profiles of NYS DOT’s class 8 dump trucks. Engine Control System’s Purifilter, Johnson Matthey's ACCRT, Johnson Matthey's CRT reformulated, and Johnson Matthey's EGRT are all incompatible with more than 84% of the vehicles tested. There is also considerable evidence that Donaldson's DFM DMF, the only level 2 FTF, is incompatible with NYS DOT's large dump trucks, given its incompatibility percentage of 79.6.

For the first three passive DPFs, the test results are less immediately conclusive. The Cleaire Longview, Donaldson LNF, and Donaldson LXF were found to be incompatible with 55 to 68 percent of tested vehicles. There was no clear pattern, such as model7600 trucks being compatible while model 2574 are not. There were incompatibilities in every region, though they did appear especially common in region 10. In short, there did not appear to be any straightforward way of predicting which vehicles would be compatible. Furthermore, there is no reason to believe that a vehicle which is compatible one year will remain so the next. For this reason, passive DPFs would be a risky retrofit solution for NYS DOT class 8 dump trucks. Were NYS DOT to install the more tolerant PDPFs on its class 8 dump trucks, it is unlikely that every vehicle would have regeneration problems, but the risk of many vehicles encountering substantial problems would be very high.

The lowest incompatibility was achieved by the ESW Canada Thermacat. Barely over five percent of temperature profiles were incompatible, but it was also the only active DPF listed. Some other active DPFs list no minimum temperature constraints. Temperature profile considerations do not, therefore, prevent active DPFs (or DOCs) from being used on NYS DOT heavy duty dump trucks.

Although not a formal requirement declared in retrofit verifications, vehicle availability is a major concern for much of the NYS DOT fleet. NYS DOT officials are understandably concerned about the possibility of having to take plows off the road at a critical time. Imagine a blizzard moves into New York and plows are immediately sent out before dawn to keep the roads clear. Late that day, it is still snowing heavily and many plows now need to actively regenerate their filters. Plows start to arrive back at their depots and plug in for five hours of regeneration. It's hard to argue that such a situation would not make the roads more dangerous and consequently pose a public safety risk. The regeneration time does not have to be five hours for there to be potential problems. DOT staff have expressed that 20 minutes of downtime can be too much, particularly if it cannot be accurately predicted.

The approximate regeneration times for the Cleaire Horizon, Cleaire Vista, Donaldson Semi-Active Electric Filter and Engine Control Systems Purifilter Plus are 5 hours, 2 hours, 4.5 hours and 2.5-4 hours respectively (Cleaire, 2008, 2009a; Donaldson, 2008; Engine Control Systems, 2008). According to Volvo, the Huss MK DPF (which is available as a retrofit) takes no more than 35 minutes to regenerate (Volvo North America, 2010). No manual could be found on Huss's website to confirm this number, and none was provided upon request. The county of Los Angeles has one Huss DPF which regenerates in roughly 30 minutes. It came installed on a new truck when they purchased it (Nunez, 2009; 2010). This 30 minute regeneration time may still be too long for NYS DOT, and there may be other compatibility issues such as a lack of space for installation. Despite being marketed as an active filter, the ESW ThermaCat regenerates during normal operation (ESW, 2009). However, NYS DOT's class 8 dump trucks exceed the maximum horsepower for the device, and many class 6 trucks violate other requirements such as model year and incompatibility with EGR.

3.2 Cost-Benefit Analysis

3.2.1 Application of Verified Retrofits

Perhaps the most straightforward means of comparing methods of emission reduction is by computing cost effectiveness in terms of emissions prevented per dollar spent. This type of metric is commonly computed when assessing plans for diesel retrofits, and was the subject of an EPA report on diesel retrofits (EPA, 2006). The cost effectiveness was computed for a range of NYS DOT vehicles and retrofit technologies. Similar metrics were computed for repowering vehicles to (partially) use compressed natural gas and for early replacement, though these numbers are more subject to uncertainty. This section will outline the process, as well as many of the data sources used, but it will focus only on class 8 dump trucks for the sake of brevity. For a more complete benefit-cost analysis, please see Gao and Stasko (2010).

The price estimates for retrofits were based on estimates from retrofit manufacturers as well as costs of recent retrofits, as reported by the fleet managers. The resulting equipment price estimates were $1,400 for a level 1 DOC, $8,000 for a level 2 flow through filter, and $15,000 for a level 3 DPF. There is inevitably some uncertainty surrounding these numbers, due to factors such as the volume ordered. The $15,000 figure matches a recent CARB estimate of the capital cost of an ADPF (CARB, 2008b). At first glance, one might wonder why PDPFs aren't given a lower equipment price, as the CARB report estimates their cost at $12,000 per heavy heavy-duty vehicle. Lower priced passive DPFs are indeed available, but the higher-price PDPFs have substantially more relaxed exhaust temperature requirements. As discussed in Chapter 3, NYS DOT vehicles often come nowhere near meeting the duty cycle requirements for the less tolerant PDPFs. Even relatively expensive PDPFs have serious compatibility issues, but for the purpose of the cost-benefit analysis in this section, they will be assumed to be compatible.

Installation cost was added to the equipment prices. The shop labor rate was assumed to be $65/hour. This labor rate was provided by NYS DOT, and it resembles shop labor rates reported in a recent survey of auto body shops (NV DMV, 2010). Installation times were based on NYS DOT experience as well as company estimates (e.g. Donaldson, 2008, 2009a). They were 4 hours for level 1 and 2 retrofits, 6 hours for level 3 PDPFs, and 9 hours for level 3 active ADPFs.

Although a relatively small portion of the overall retrofit cost, installation costs are not necessarily inconsequential. If NYS DOT were to retrofit every vehicle in the January 1st, 2008 snapshot with a DOC (the simplest retrofit installation), at the NYS DOT experience install time of 4 hours/DOC, the total time required would be over 10,000 hours. This equates to more than five full time mechanics working a whole year at 40 hours/week and 50 weeks/year. With level 3 retrofit technologies, the cost would be substantially greater.

In addition to retrofit equipment and installation costs, PDPFs and ADPFs have periodic costs from ash removal. The interval between cleanings was estimated to be 1 year, based on EPA figures (US EPA, 2009d), a CARB report (CARB, 2008e), a Massachusetts Department of Environmental Protection report (M DEP, 2008), equipment manuals (Donaldson, 2009b; Engine Control Systems, 2006; 2008), and DSNY experience (Kim, 2009), all of which either put the interval at 1 year or a range including 1 year. The cost was assumed to be $300 per cleaning, roughly in the middle of estimates of $250-300 by DSNY (Kim, 2009), $200-400 by M DEP (2008), $400 by CARB (2008e), and $500 by the county of Los Angeles (Nunez, 2009). The costs of future cleanings were discounted to the present using a 5% annual interest rate. Newer vehicles with longer expected remaining lifetimes had more expected cleanings, making them more expensive to retrofit.

One more cost was included for ADPFs, and that was the cost of regeneration stations. This equipment is needed much more frequently than de-ashing equipment (sometimes daily regenerations are warranted), and consequently it was assumed that NYS DOT would purchase the equipment. A charging station is not necessarily required for every vehicle, both because filters don't always need to be regenerated daily, and because recharging stations can sometimes be set up to conduct multiple regenerations in sequence overnight (Donaldson, 2008). Cost estimates were based on an Oakland Public Works Agency proposal, which called for 6 stations, costing $6,000 each, to serve 27 filters (OPWA, 2007). This amounts to a cost of $1,333 per ADPF purchased.

The benefits of exhaust treatment retrofits are, of course, the emission reductions. Unretrofitted vehicle emission rates, in grams/mile, were obtained from the EPA MOBILE6.2 model. The emissions impacts of verified retrofit technologies, in terms of percentage reductions, were obtained from EPA and CARB verified technology lists. Retrofits are verified to reduce PM, CO, HC, and (infrequently) NO_x. PM is the basis for CARB's technology levels (CARB, 2009e) and the tiers in NYS DEC regulations (NYS DEC, 2009c), and is therefore the focus of this analysis. Based on fuel usage and mileage reading data, annual mileage is assumed to be 13,750. All vehicles of age 13 or younger are assumed to remain in use until they reach age 14. Vehicles already 14 or older are assumed to remain in use for one more year. The grams of PM emissions prevented by applying different types of retrofits to different age trucks are plotted in Figure 4.

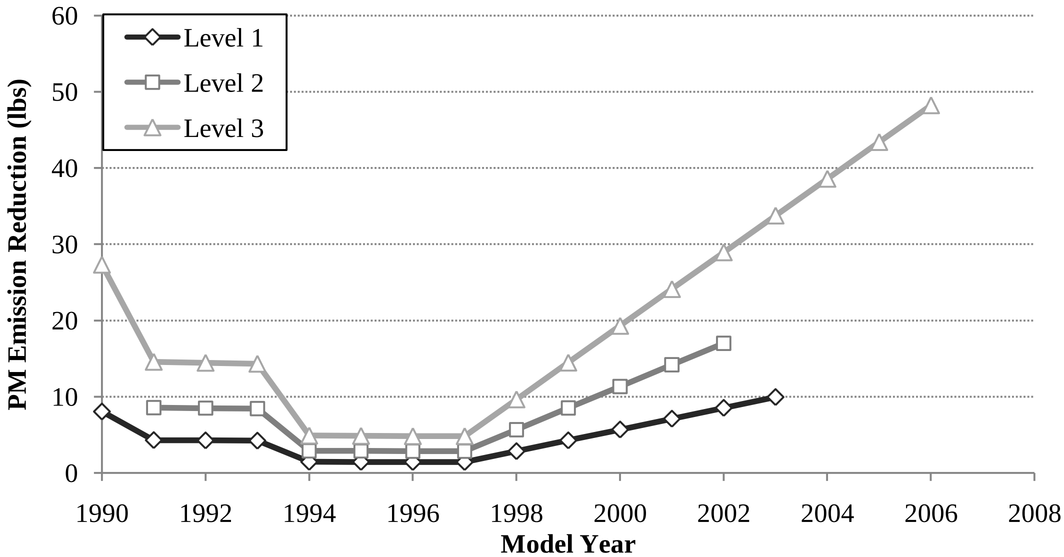

Figure 4. Emission Reductions from Retrofits of Class 8 Trucks

Emissions prevented are plotted for all model years for which it is plausible that one of the EPA or CARB verified retrofits is compatible. Given the incomplete fleet data described earlier, it is likely that several incompatibilities are not represented. Most of the incompatibilities represented are based on model year and EGR.

Regardless of the level of the retrofit, retrofitting the newest (pre-2007) vehicles generates the largest emissions savings. The lower expected remaining mileage of older vehicles is what drives down their emission reductions. Pre-1994 vehicles start to have substantially higher emission rates, causing larger emission reductions.

The grams of PM emission prevented per dollar spent are plotted in Figure 5. This provides a simple “bang per buck” metric, with higher values meaning more emissions reduced per dollar spent. Level 3 is broken down into passively regenerating and actively regenerating technology. Although emissions benefits are generally similar, these different types of filters have different costs and vehicle compatibility.

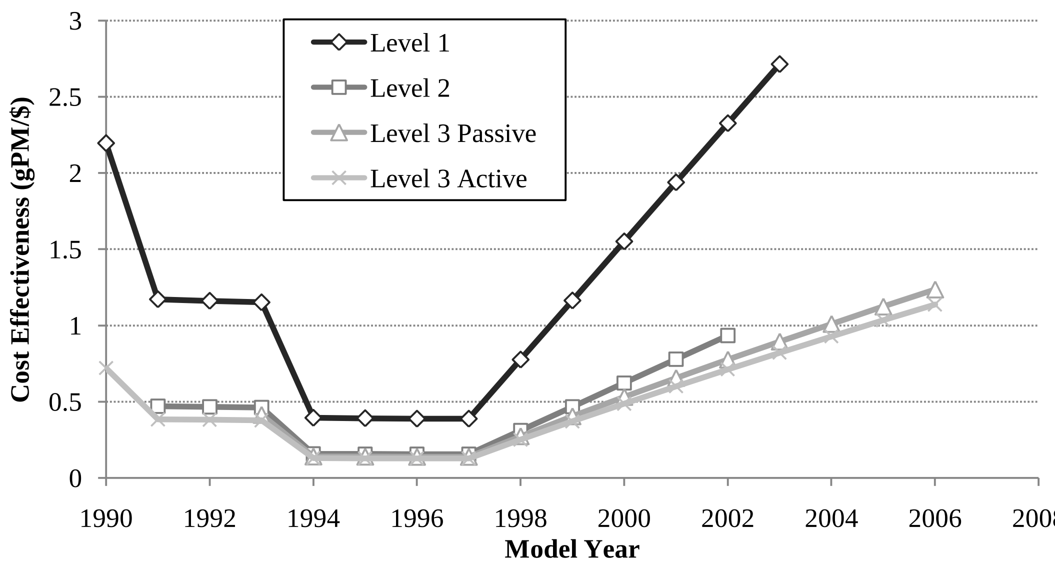

Figure 5. Long-Term Cost Effectiveness of Retrofitting Class 8 Trucks

The general shape of the cost effectiveness curves matches that of the emissions curves, and is driven by the same factors. The cost effectiveness curves are scaled by costs, which makes it apparent that level 1 retrofits (DOCs) offer the highest bang per buck, despite providing the lowest emission benefits.

3.2.2 Early Replacement

The cost structure of early vehicle retirement is inherently quite different from that of retrofits. Early vehicle replacement has very high initial costs, but these costs can be substantially offset by future savings. Imagine, for example, that a vehicle slated for retirement in 3 years in replaced today instead. This requires that the NYS DOT spend the full purchase price of the replacement vehicle. Selling the replaced truck, whether for scrap or for use as a vehicle, does not generate revenue for NYS DOT to offset the purchase price. Three years from now, however, the planned replacement does not have to be made. This difference in when costs and savings are experienced can cause dramatic differences between the initial and long-term cost effectiveness of vehicle replacements, assuming no financing option is utilized.

The new vehicle purchase price of a class 8 dump truck is assumed to be $160,000 (based on NYS DOT estimates). This is the initial cost of replacement. Normally, it could be argued that the long-term cost of losing a vehicle (with no revenue compensation), and consequently having to replace said vehicle, is simply the market value of that vehicle. NYS DOT is in a somewhat unusual situation, however, which makes the picture more complicated. From the NYS DOT perspective, used vehicles are less valuable than they are to most fleet owners. This is because most fleet owners would include future resale revenue in their estimation of a used vehicle's value. For older vehicles, this might even be the majority of the vehicle's value. For the NYS DOT, there is no resale revenue, making older vehicles less valuable from their perspective. For this reason, market values of used vehicles are computed, and then adjusted to create NYS DOT values of used vehicles. For newer vehicles, resale is a relatively small portion of the value, making the market and NYS DOT values very close, but the NYS DOT values drop more sharply with age.

The fact that NYS DOT does not keep vehicle auction revenue has naturally encouraged NYS DOT to wait until vehicles are in very poor condition before selling them. Harsh operating conditions, especially during snow removal, can contribute to vehicle deterioration. It is not unusual for NYS DOT technicians to remove parts which might be needed as spares before selling a vehicle. This means that historic auction prices don't provide a complete picture of used vehicle market values throughout the vehicle lifetime. Historic auction data can be used to estimate vehicle lifetime under current practices, as well as vehicle scrap value, however. Auction data from 2005 to 2009 was used to estimate a class 8 diesel truck lifetime of 14 years, and a scrap sale value of roughly $2,115.

The above information provides the market value of a new truck, and of a very old truck, but not of any truck in between. Unfortunately, few data points for prices of such trucks could be found. These intermediate market values were filled in using the assumption that the market price of a used diesel truck decreases exponentially with age, following a pattern similar to that found for diesel school buses (Gao and Stasko, 2009). The resulting truck values are plotted in Figure 6. Note that these are intended as best guesses, only knowing the vehicle type and age. Individual vehicle values will naturally vary depending on factors such as mileage and condition. Adjusted NYS DOT perspective vehicle values are plotted in Figure 7. These are the long-term replacement costs.

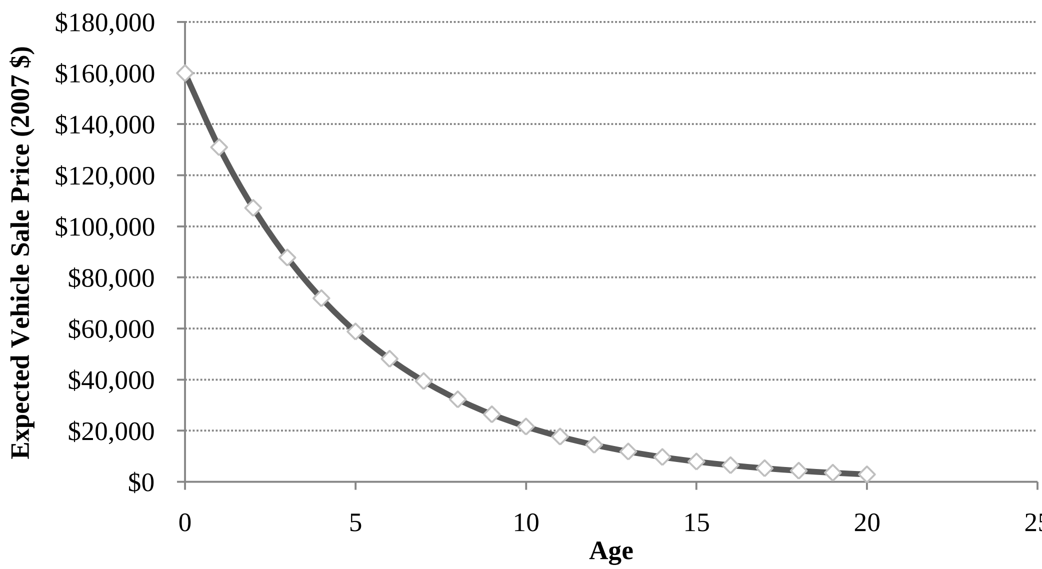

Figure 6. Expected Vehicle Market Price by Age

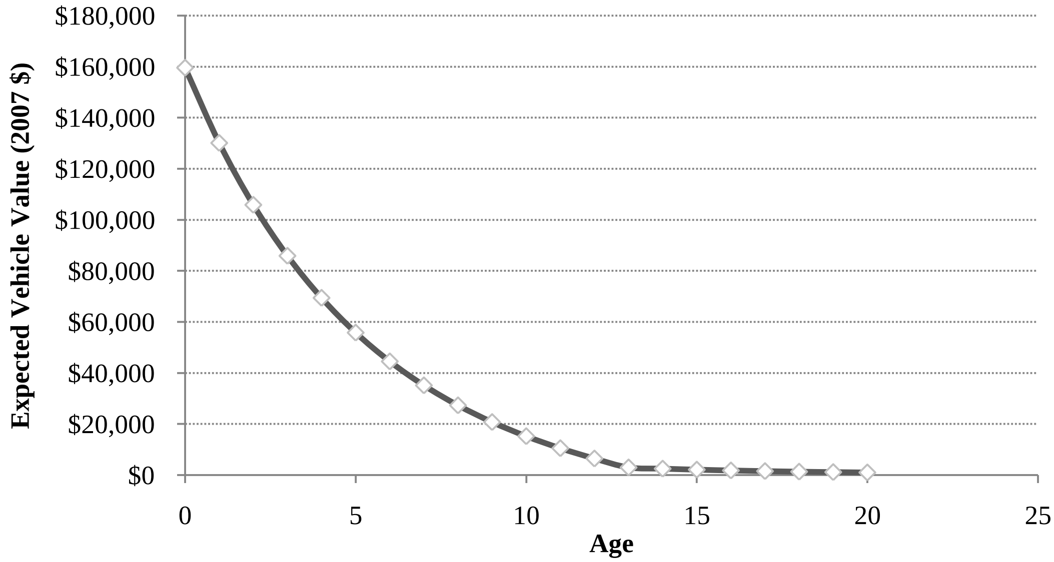

Figure 7. Expected Vehicle NYS DOT Value by Age

The emissions reduction achieved from an early replacement is also somewhat less straightforward than that from a retrofit. The most obvious change occurs when the retired vehicle would have operated, but does not because of the early retirement. For these years, the emissions reduction can be computed based on the difference between the emission rates of the retired vehicle and those of its replacement.

Further along in time, the effects become more complicated. A replacement vehicle purchased today might have higher emission rates than a replacement vehicle purchased in three years. Furthermore, purchasing a replacement three years earlier than previously planned will likely mean that the new vehicle will also have to be replaced earlier than planned. This effect can be carried further and further into the future, with decreasing certainty.

Historically declining emission rates limit the importance of these uncertainties, however. As long as emission rates do not increase, future changes to emission rates will be much smaller than those made in recent history. Recall the simplified histories of PM and NO_x emission standards for new trucks, as presented in Figures 1 and 2 (US EPA, 2003a). Between the late 1980s and 2010, both PM and NO_x emission rate caps have been lowered by over 98%. This means that even if both PM and NO_x emission rates were cut to absolutely zero next year, the change would be dramatically smaller than those seen over the past couple decades. Therefore, future changes to emission rates are ignored. The emission reduction from an early retirement is assumed to be the change in emissions for the additional years the retired vehicle would have been operating if not retired early.

A few additional assumptions are made in order to compute the cost effectiveness of vehicle replacement. As in the previous section on retrofits, annual mileage is assumed to be 13,750. All vehicles of age 13 or younger are assumed to remain in use until they reach age 14. Vehicles already 14 or older are assumed to remain in use for one more year. The resulting long-term cost effectiveness curve (in terms of grams of PM emissions prevented per dollar spent) is plotted in Figures 8 and 9 as a function of model year. The only difference between the two figures is that Figure 9 zooms in on the lower portion of Figure 8 to reveal the variation in recent model years.

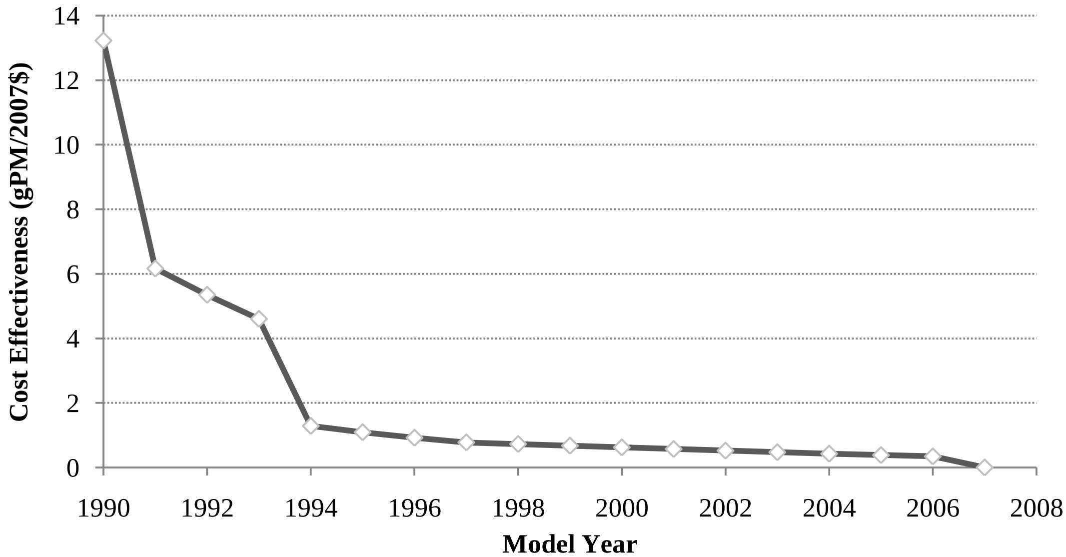

Figure 8. Long-Term Cost Effectiveness of Replacement (Zoomed Out)

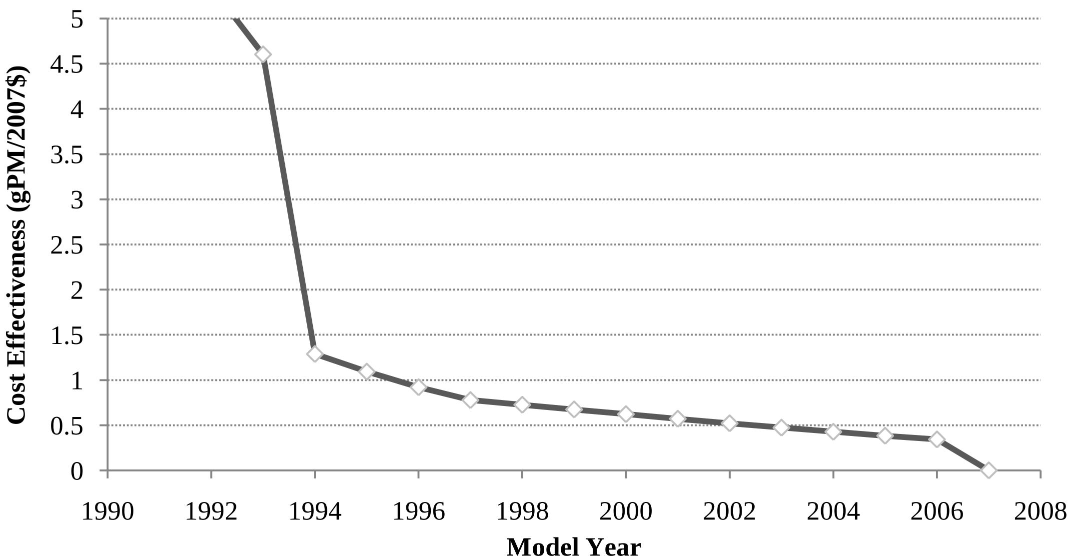

Figure 9. Long-Term Cost Effectiveness of Replacement (Zoomed In)

Numerous factors influence the shape of the long-term cost effectiveness curves. These can be explained by starting with model year 2007 and then gradually looking at older vehicles (moving from right to left across Figures 8 and 9). Replacing a model year 2007 vehicle has no PM emission benefit because the replacement vehicle would have the same emissions rate as the original vehicle. This means the cost effectiveness of replacing a model year 2007 (or newer) vehicle is zero.

Model year 2006 PM emission rates are more than 13 times those for model year 2007. This means there is considerable emissions savings from replacing a model year 2006 truck. Class 8 diesel PM emission rates were unchanged 1996-2006. Within this interval, long-term cost effectiveness gradually increases as vehicles get older. Two competing factors are at work. Older vehicles tend to have lower remaining mileage, meaning lower expected emissions savings. Older vehicles also have lower value, meaning lower long-term replacement cost. The former factor makes replacing older vehicles less cost effective, while the latter (dominant) factor makes replacing older vehicles more cost effective.

For the oldest vehicles on the far left side of Figures 8 and 9, the shape of the curve changes, shooting upward quickly. By the time vehicles are this old, their expected remaining mileage is low but fixed, meaning that this factor no longer influences the shape of the curve. The vehicle values are still decreasing with age, driving down the long-term replacement cost. Changes in emission rates also play an important role. These changes were especially large going from 1993 to 1994 and from 1990 to 1991. Both of these factors make older vehicles more cost effective to replace. The cost effectiveness would continue to grow as vehicles become older if the graphs were extended to the left and earlier model years were added. A quick glance at Figure 8 indicates that replacing the oldest vehicles is likely to be one of the most cost effective emission reduction techniques in the long term, but that replacing new vehicles would be much less cost effective. This difference would be made even more pronounced if multiple unreliable old vehicles could be replaced with a single reliable new vehicle.

The short-term cost effectiveness of replacement (emission reduction divided by new vehicle purchase price) is plotted in Figure 10. The short-term cost effectiveness curve looks quite different from the long-term cost effectiveness curve. The far right of both curves is the same, as the lack of emissions benefit from replacing a model year 2007 vehicle gives this action a cost effectiveness of zero, regardless of timeframe. The era of unchanging emission rates which resulted in a slow increase in long-term cost effectiveness for older vehicles now exhibits lower short-term cost effectiveness for older vehicles. Older vehicles still tend to have lower remaining mileage, meaning lower expected emissions savings. The competing effect of older vehicles having lower market values is gone, however. The short-term cost is independent of model year, as the new vehicle price is independent of the age of the vehicle replaced. Emission rate changes still drive the curve upward on the far left side of Figure 10, but they are no longer compounded by cost changes, making the rise much less dramatic.

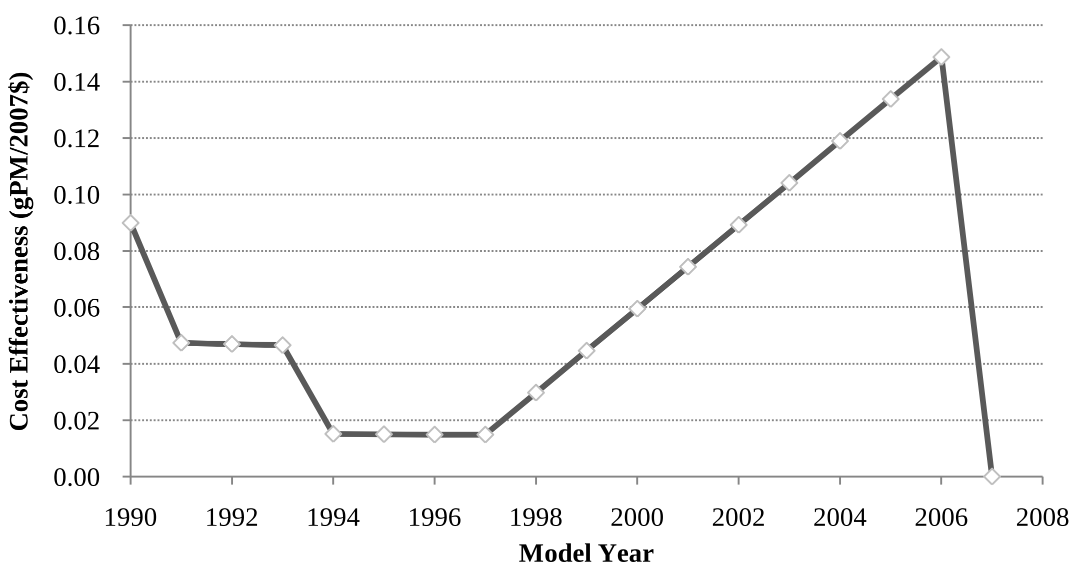

Figure 10. Short-Term Cost Effectiveness of Replacement

3.2.3 Comparison of Costs and Benefits

To the extent possible, retrofits and replacements were analyzed using comparable assumptions, to allow an “apples to apples” comparison. The costs of all options considered are plotted together in Figure 11.

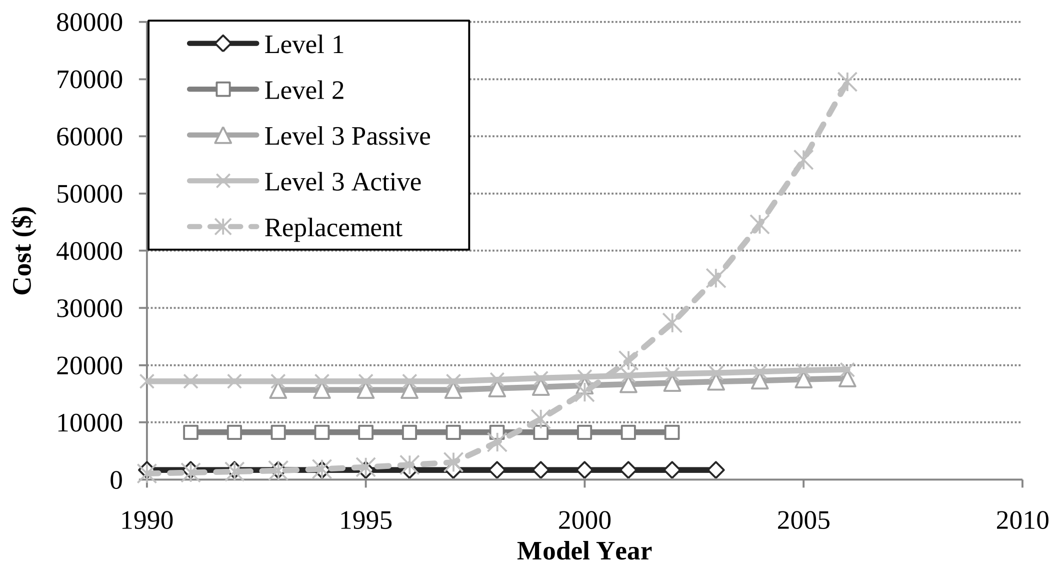

Figure 11. Class 8 Dump Truck Long-Term Costs/Vehicle

Not surprisingly, level 1 retrofits are the cheapest option for all the trucks with more than a couple years of expected life left. For the oldest trucks, it is cheaper in the long term to replace the vehicle than to retrofit with any level technology. The value of the vehicle to NYS DOT is actually less than the cost of retrofitting, even with a DOC. Unsurprisingly, replacing the newest vehicles is the most costly action. Unfortunately, there isn’t a good cheap option for vehicles which are too new for level 1 DOCs.

PM emissions prevented are plotted in Figure 12. The benefits of level 3 technologies follow slightly behind those of replacement, with level 2 and 1 further behind.

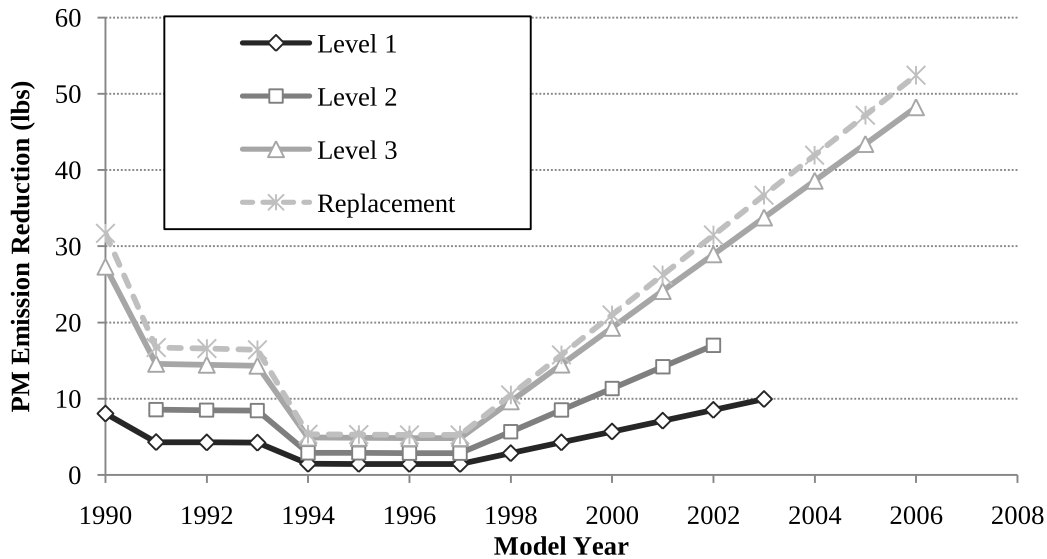

Figure 12. Class 8 Dump Truck Emission Reductions/Vehicle

The long-term cost effectiveness curves for class 8 dump trucks are plotted in Figure 13. In the long term, the most cost effective way to reduce PM emissions from class 8 trucks is to replace the oldest trucks. Another relatively cost effective option is to install level 1 retrofits on relatively new class 8 dump trucks, starting with the newest which are compatible. It is generally more difficult to find cost effective methods of reducing PM emissions for vehicles too new for DOCs, or from the mid 1990’s. CNG conversion may be the most cost effective option (as well as the cheapest) for a subset of the newest class 8 dump trucks, but this would have to be evaluated on a vehicle by vehicle basis because of the complexities of conversion approval discussed in Gao and Stasko (2010).

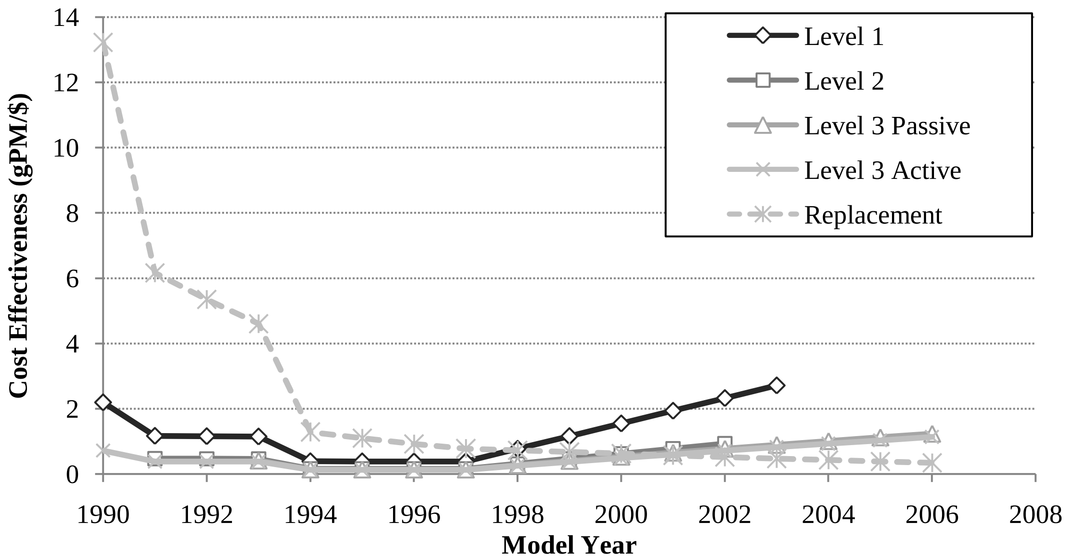

Figure 13. Class 8 Dump Truck Long-Term Cost Effectiveness Comparison

3.2.4 Strengths and Limitations of the Benefit-Cost Analysis

Given a single metric by which to compare emission reduction options, it is fairly straightforward to choose the best approach for each vehicle. One can simply rank the feasible alternatives, and choose the one on top. It is a transparent and intuitive process. No sophisticated algorithms or extensive computing time are required.

The difficulty lies in selecting a single metric. One might argue that emissions reduction per dollar spent is the best metric, but this leaves open the question of which emissions to include. Also, NYS DOT may be most concerned with finding the cheapest way to maintain regulatory compliance, independent of the emission reductions achieved. If this is the case, it makes sense to rank alternatives by cost. Ranking by cost is not as simple as it first appears, however, as there may be short-term budgets as well as long-term cost minimization goals. While the benefit-cost analysis presented has the advantage of being intuitive and easily interpretable, it does not help the fleet manager to balance multiple objectives and constraints that inevitably cloud such decisions.

Cost-benefit analysis is also poorly suited to helping the fleet manager to make multiple related decisions. For example, retrofit and replacement decisions might impact how heavily vehicles should be used, and flexibility in vehicle usage might impact what retrofits and replacements are optimal.

CHAPTER 4

APPLYING INTEGER PROGRAMMING

4.1 Model Formulation

4.1.1 Initial Formulations

Transitioning from a standard benefit-cost analysis to integer programming adds computational complexity, but it also opens up a wide range of possibilities. As a first step, a single-period integer program was developed, with the original case study based on school buses. The full formulation is available in Gao and Stasko (2009a). Because it can largely be regarded as a special case of the multi-period model presented below, it will not be described in detail here. The single-period model minimized long-term costs subject to a short-term budget, as well as constraints on emissions of each pollutant. By changing the right hand sides of these constraints and resolving the integer program, tradeoff curves were produced.

The speed with which the single-period model solved was encouraging, as was the power of the results produced. The model was capable of quickly balancing competing metrics and illustrating the efficient frontier. At the same time, several weaknesses were apparent. Perhaps the most obvious weakness was the model's inability to handle phased-in regulation. This was particularly troubling given that most retrofit mandates are phased in. Also, the original single-period model had no ability to represent flexibility in vehicle usage. Both of these issues were addressed with the development of the multi-period model.

The first formulation of the multi-period model was presented in Stasko and Gao (2010). It combines vehicle purchase and retrofit decisions with aggregated task scheduling. Tasks are assumed to have already been grouped. Vehicles are assigned to a single group of tasks, called a run. The objective was to minimize long-term costs, given short-term budgets. The effects of emissions taxes were examined, and the recommendations of the multi-period model were compared to the recommendations of a sequence of single-period models. In several examples, the single-period models demonstrated a propensity for setting the fleet owner up with an infeasible problem, as they did not fully take into account the future impacts of current decisions. Even where the sequence of single-period models remained feasible, their performance lagged behind that of the multi-period model.

Several changes were made to the multi-period model to produce the revised formulation presented below. First, the value of the fleet at the end of the time horizon was captured, so investments which paid off beyond the time horizon would be properly valued. Second, the model was extended to allow for a more complicated process of retirement in which the fleet manager may not keep all revenue, and can choose to scrap some vehicles before selling them. This better represents the situation facing NYS DOT. Third, additional regulatory alternatives were considered. In addition to emission taxes, emission rate taxes, and clean technology mandates were included.

4.1.2 Building the Objective

The objective remains to minimize long-term costs while performing the fleet's regulated duties, though the exact formulation of this objective differs somewhat from previous models. Costs include vehicle purchases, as well as maintenance, repairs, retrofits, and fuel. Furthermore, from the fleet manager's perspective, vehicle and emissions taxes are included. Finally, there are costs which do not show up as cash flows. If depreciation of the fleet was not considered, the integer program would have a tendency to select policies with low costs in the short term, but which leave the fleet manager with an old and poorly maintained fleet at the end of the time horizon. For this reason, any decrease in the value of the fleet is considered a cost. Because the initial fleet value is a constant, it can be dropped from the objective without changing the optimal policy. This leaves the challenge of defining vehicle values at the end of the time horizon.

One option for estimating asset values is to use market prices. In a perfect market with identical players, this method would work extremely well. The further these assumptions are from the truth, the less helpful market prices will be. As Drinkwater and Hastings (1967) pointed out, market prices will rarely equal residual asset values. There are nearly always significant transaction costs, such as taxes, transportation, and labor. Also, market prices are set by a combination of many players who use and value assets differently. While plugging in market prices for initial asset values could be a reasonable start, it may not be the best approach. Of course, good market price data may not even be available. While a reasonable collection of market prices were found for the school buses analyzed in Gao and Stasko (2009), only very limited market price data was found for vehicles resembling those in the NYS DOT fleet.

Vehicle values can be solved for analytically in the deterministic steady-state case. The integer program is fast enough to run for many years beyond the end of regulatory phase-in and emission rate changes, giving the system time to reach a steady-state. Given that vehicles are replaced with identical new vehicles forever, the value of a used vehicle of any age can be calculated by measuring its ability to delay these costs. First, the optimal lifespan is computed by minimizing the equivalent uniform annual cost as in Newnan et al. (2002). Second, the value of having a vehicle at the end of the current period, which will be retired at that time, is set equal to the resale revenue discounted to the present. Third, all younger vehicles have their values recursively calculated according to expression (2).

$$V\left[a_f\right] = S/(1+\delta) \tag{1}$$

$$V\left[a_f - \mu\right] = \frac{V\left[a_f - \mu + 1\right]}{(1+\delta)} + \frac{\left(L + \frac{L}{(1+\delta)^{a_f} - 1}\right)\left(1 - \frac{1}{(1+\delta)}\right)}{(1+\delta)} - \frac{m[a_f - \mu + 1]}{(1+\delta)} \tag{2}$$

Where:

$V[a]$ is the value of a vehicle of age a

a_f is the last age at which the vehicle is used

S is the resale revenue

δ is the discount rate

L is the lifetime cost of a vehicle held for the optimal lifespan, discounted to the purchase date

$m[a]$ is the maintenance and repair cost paid when using a vehicle of age a

These vehicle values are from the fleet manager's perspective, and may differ from used vehicle market prices if other fleets use the vehicles differently or if there are transaction costs such as taxes, registration fees, or transportation.

The approach of recursively computing asset values as a function of age is not entirely new. Bellman's 1955 paper on equipment replacement outlined a similar technique. Bellman's formulation was designed for equipment which generates "output" possibly in excess of its upkeep, meaning that equipment can be valued according to its potential to generate profit. In the NYS DOT example, and hence in the revised formulation in expression (2), there is no profit. Instead, used vehicles are valued by their ability to prevent costs which would otherwise be incurred to provide mandated service. Current ownership of a used vehicle can impact expected purchase dates infinitely far into the future, which is why the second term in expression 2 contains the solution of a perpetuity.

4.1.3 Constructing the Constraints

Expressions (4)-(12) are constraints which define tracking variables, purely for convenience and readability. Perhaps the most obvious set of constraints enforce the laws of physics as well as human laws against theft. Expressions (13)-(18) essentially require "conservation of vehicles." Vehicles are tracked throughout the time horizon. No vehicles appear without being purchased. No vehicles disappear without being sold. Expression (19) caps the number of vehicles purchased in a given time period, while expression (20) enforces compatibility of retrofits. Expression (21) requires that vehicles are sold on time while expression (22) requires that vehicles are deactivated on time. Expression (24) requires that no inactive vehicles are used.

From the fleet manager's perspective, the duties to be performed are taken as given. The DOT, for example, is not permitted to choose not to plow the roads. There is some degree of flexibility in terms of which vehicles are used more or less heavily, however, as indicated by expression (23).

Any potential clean diesel mandate is imposed through expressions (25)-(27). Expression (27) applies to the first year of the phase-in, while expression (26) applies to the second year, and expression (25) applies to all subsequent years.

The integer programming framework is flexible enough to include a wide range of additional side constraints, including budgets, emission reduction goals, and refueling equipment requirements.

4.1.4 IP Formulation: Sets

$t \epsilon T$ set of time periods (t_f is the final period)

$i \epsilon I$ set of vehicle types

$j \epsilon J$ set of retrofit states (1= active but unretrofitted, 2=active with DOC, 3=active with expensive DPF, 4=active with FTF, 5=active with cheap DPF, 6=inactive)

$a \epsilon A$ set of pollutants

$r \epsilon R[t]$ set of duties (or runs) a vehicle can be assigned to in period t

4.1.5 IP Formulation: Input Parameters

λ_i period by which vehicle type i must become inactive

l_i period by which vehicle type i must retire

f_{ij} number of vehicles of type i with retrofit j in the initial fleet

v_{ijkt} cost to switch vehicle type i from retrofit option j to k at the start of period t

p_{it} cost to purchase vehicle type i at the start of period t

c_{ijtr} cost to cover run r in period t with vehicle type i with retrofit j

e_{ijtra} emissions of pollutant a from vehicle type i with retrofit j, running on r in period t

n_{tr} number of vehicles required for each run r in period t

q_{it} max number of vehicles of type i which can be purchased in period t

u_{ij} 1 if vehicles of type i can use retrofit j

s_{ijt} resale value of vehicle type i at the start of period t

σ_{ij} value of vehicle type i with retrofit j to the fleet at end of the last period of the simulation

ρ discount factor = $1/(1+\delta)$

φ fraction of resale revenue kept by the fleet

t_{MF} period in which mandate takes full effect

θ percent of (active non-exempt) vehicles which must use BART or be post-2006 starting in t_{MF}

h_i highest retrofit tier which is compatible with vehicle type *i*

w_a pollution weight used in taxes for pollutant *a*

α \$/gram tax paid for weighted combo of pollutants

β_{ij} emission rate tax, for vehicle type *i* and retrofit status *j*

4.1.6 IP Formulation: Decision Variables

x_{ijtr} vehicles of type *i* with retrofit *j* assigned to run *r* in period *t*

y_{ijkt} vehicles of type *i* going from retrofit *j* to *k* at the start of period *t*

b_{it} vehicles of type *i* bought at the start of period *t*

g_{ijt} number vehicles of type *i* with retrofit *j* retired at the start of period *t*

$etaxpaid_t$ emissions tax paid in period *t*

$vtaxpaid_t$ emission rate tax paid in period *t*

$resaletostate_t$ resale revenue kept by state in period *t*

$resaletofleet_t$ resale revenue kept by the fleet in period *t*

fleetvalatend value of the fleet at the end of the last period

$noneBART_t$ number of active pre-2007 vehicles not compliant with any retrofit in period *t*

$DOCBART_t$ number of active vehicles using a DOC as BART in period *t*

$FTFBART_t$ number of active vehicles using a FTF as BART in period *t*

$DPFBART_t$ number of active vehicles using a DPF as BART in period *t*

$post2006_t$ number of active post-2006 model year vehicles in period *t*

4.1.7 IP Formulation: Objective

$$\min \sum_{t \in T} \rho^{t-1} \left(etaxpaid_t + vtaxpaid_t + \sum_{i \in I} b_{it} p_{it} + \sum_{i \in I} \sum_{j \in J} \sum_{k \in J} y_{ijkt} v_{ijkt} + \sum_{i \in I} \sum_{j \in J} \sum_{r \in R[t]} x_{ijtr} c_{ijtr} - resaletofleet_t \right) - \rho^{t_f} fleetvalatend \tag{3}$$

4.1.8 IP Formulation: Constraints

$etaxpaid_t = \sum_{i \in I} \sum_{j \in J} \sum_{r \in R[t]} \sum_{a \in A} \alpha e_{ijtra} x_{ijtr} w_a$	for all {t in T}	(4)
$vtaxpaid_t = \sum_{i \in I} \sum_{j \in J} \sum_{r \in R[t]} \beta_{ij} x_{ijtr}$	for all {t in T}	(5)
$resaletostate_t = \sum_{i \in I} \sum_{j \in J} s_{ijt} g_{ijt} (1 - \varphi)$	for all {t in T}	(6)
$resaletofleet_t = \sum_{i \in I} \sum_{j \in J} s_{ijt} g_{ijt} (\varphi)$	for all {t in T}	(7)
$noneBART_t = \sum_{i \in I} \sum_{j \in J} y_{ij1t}$: h_i=0 and *modelyear*[*i*]<2007	for all {t in T}	(8)
$DOCBART_t = \sum_{i \in I} \sum_{j \in J} y_{ij2t}$: h_i=1	for all {t in T}	(9)
$FTFBART_t = \sum_{i \in I} \sum_{j \in J} y_{ij4t}$: h_i=2	for all {t in T}	(10)
$DPFBART_t = \sum_{i \in I} \sum_{j \in J} (y_{ij3t} + y_{ij5t})$: h_i=3	for all {t in T}	(11)
$post2006_t = \sum_{i \in I} \sum_{j \in J} y_{ij1t}$: *modelyear*[*i*]>2006	for all {t in T}	(12)

$b_{i1} + f_{i1} = g_{i11} + \sum_{k \in J} y_{i1k1}$	for all {i in I}	(13)
$f_{ij} = g_{ij1} + \sum_{k \in J} y_{ijk1}$	for all {i in I, j in J:j>1}	(14)
$\sum_{j \in J} y_{ijk1} = \sum_{r \in R[1]} x_{ik1r}$	for all {i in I, k in J}	(15)
$b_{it} + \sum_{r \in R[t-1]} x_{i1(t-1)r} = g_{i1t} + \sum_{k \in J} y_{i1kt}$	for all {i in I, t in T:t>1}	(16)
$\sum_{r \in R[t-1]} x_{ij(t-1)r} = g_{ijt} + \sum_{k \in J} y_{ijkt}$	for all {i in I, j in J, t in T:t>1, j>1}	(17)
$\sum_{j \in J} y_{ijkt} = \sum_{r \in R[t]} x_{iktr}$	for all {i in I, k in J, t in T:t>1}	(18)

$b_{it} \leq g_{it}$	for all {i in I, t in T}	(19)

$\sum_{j\in J} y_{ijkt} \leq u_{ik} * 9999999$ for all {i in I, k in J, t in T} (20)

$\sum_{j\in J} \sum_{t\in 1..l_i} g_{ijt} = \sum_{j\in J} f_{ij} + \sum_{t\in 1..l_i} b_{it}$ for all {i in I} (21)

$\sum_{j\in J} \sum_{k\epsilon 1..5} y_{ijkt} = 0$ for all {i in I, t in T: λ_i<=t} (22)

$\sum_{i\in I} \sum_{j\epsilon J} x_{ijtr} \geq n_{tr}$ for all {t in T, r in R[t]} (23)

$x_{i6tr} = 0$ for all {t in T, i in I, r in R[t]:r<4} (24)

$$DOCBART_t + FTFBART_t + DPFBART_t + post2006_t \geq \frac{\theta(\sum_{i\in I} \sum_{j\in J} \sum_{k\in 1..5} y_{ijkt} - noneBART_t)}{100}$$
for all {t in T: t>=t_{MF}} (25)

$$DOCBART_t + FTFBART_t + DPFBART_t + post2006_t \geq \frac{2}{3} * \frac{\theta(\sum_{i\in I} \sum_{j\in J} \sum_{k\in 1..5} y_{ijkt} - noneBART_t)}{100}$$
for all {t=t_{MF}-1} (26)

$$DOCBART_t + FTFBART_t + DPFBART_t + post2006_t \geq \frac{1}{3} * \frac{\theta(\sum_{i\in I} \sum_{j\in J} \sum_{k\in 1..5} y_{ijkt} - noneBART_t)}{100}$$
for all {t=t_{MF}-2} (27)

4.2 Case Study and Results

4.2.1 Example Fleets

The fleet manager's integer program was run for four representative fleets. Each fleet resembled the NYS DOT class 8 dump truck fleet in terms of size, vehicle type, and usage. The fleets differed in terms of the compatibility of the vehicles with various retrofit technologies. While the regulators at the NYS DEC would likely encounter only minimal difficulty in estimating fleet size and composition, they would have a much harder time predicting retrofit technology incompatibilities. As discussed in Chapter 3, retrofit technologies are verified for use on specific vehicle types by the EPA and CARB. These verifications are highly specific, requiring details ranging from horsepower and displacement to the EPA engine family name. NYS DOT does not maintain records of this data for many of its vehicles. Verifications also frequently include requirements for the exhaust temperature profile, which is a function of the way the vehicle is used. For these reasons, compatibility assumptions would have been a major source of uncertainty when constructing regulations. The four compatibility scenarios are laid out in Table 6. DOC compatibility is held constant in all four scenarios; all model year 2003 and older vehicles are considered DOC compatible.

Table 6. Compatibility Scenarios

Number	Scenario Name	Fraction of vehicles which are compatible with the cheap DPF	Fraction of the rest which are compatible with the expensive DPF
1	Diverse	0.2	0.5
2	Complete	1	NA
3	Minimal	0	0
4	Expensive	0	1

Because the example fleets are based on NYS DOT, the fleet manager does not receive revenue generated by used vehicle auctions. That revenue, which is tracked, would go to the state government as a whole. The fleet manager can salvage a few hundred dollars worth of parts from vehicles before selling them, independent of the vehicles' age.

4.2.2 Cost Curves

Although the optimizations were run from the perspective of the fleet manager, these cost curves are presented from the point of view of the state as a whole. This means that emission taxes and emission rate taxes, which are transfers from one state department to another, do not show up. Also, the auction revenue which NYS DOT does not get to keep is subtracted from the costs because it would become state revenue.

Three instruments were tested independently, at varying levels of intensity. They were emission taxes, emission rate taxes (based on emission rate), and a technology mandate similar to that which was actually implemented. Emission taxes were always proportional to the weighted combination of pollutants, but were scaled up by a common factor. These taxes require that the regulator know how heavily each vehicle is used. Emission rate taxes were always proportional to the weighted composite emission rate, with a common scaling factor. The weights were applied to different pollutants based on the way CARB weights pollutants in the Carl Moyer Program. All active vehicles were weighted equally, regardless of their mileage, so the regulator need not have detailed mileage data. Emission rate taxes were not applied to inactive vehicles being used solely for scrap parts. Mandates were scaled by changing the percentage of vehicles which were required to be compliant at the end of the phase-in. In the previous year, two thirds of this percentage needed to be compliant, while two years before one third of this percentage needed to be compliant.

Environmental economic theory indicates that in the absence of market distortions (apart from the emissions), the optimal policy would be to tax emissions with a “Pigouvian tax” set equal to the marginal societal cost of emissions. The remainder of this chapter will compare the impacts of the instruments and discuss how and why the findings can differ from what might be expected from conventional theoretical models.

Figure 14 shows the complete cost curves for all three instruments in compatibility scenario 1 (diverse). It is immediately apparent that the cost curves for the emission tax and emission rate tax are very similar. In fact, many of the places where they appear to be different could simply be the result of the linear interpolations made between points. There are, however, a few places where it is clear that the emissions tax is slightly more efficient, as theory would predict. This stems from the ability the emissions tax has to encourage heavier usage of low emitting vehicles. As it turns out, this ability is often unnecessary, as low emitting vehicles are frequently newer vehicles with lower maintenance costs. With or without the encouragement of the emissions tax, the fleet manager will tend to use these newer vehicles more (mileage readings on DOT vehicles confirm this trend). This can be taken as good news, as it means that mileage data, which could be difficult for the regulator to reliably obtain for large numbers of vehicles, is largely unnecessary.

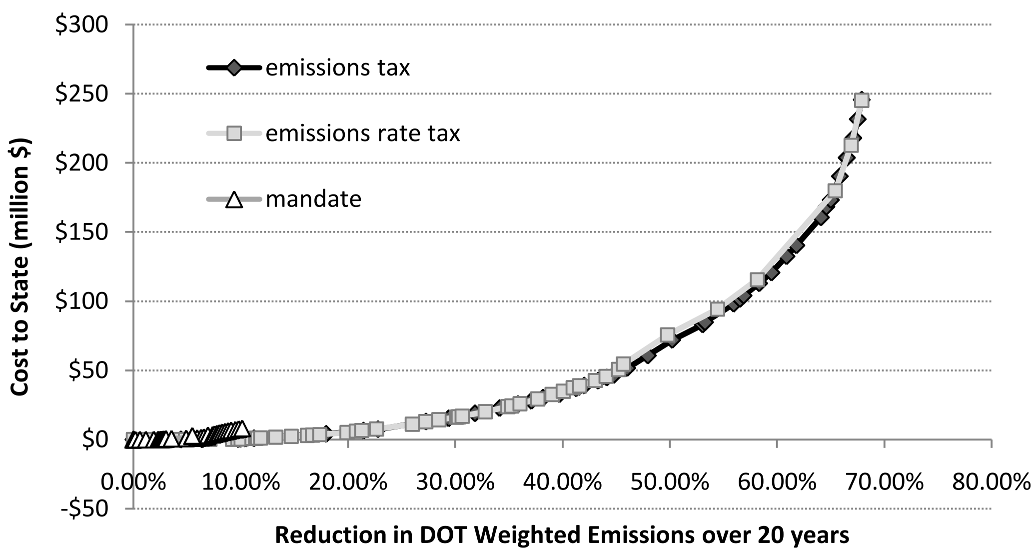

Figure 14. Cost to State vs. Emission Reduction in Compliance Scenario 1

It is also immediately obvious from Figure 14 that the mandate is not capable of producing the large emission reductions which can be caused by emission and emission rate taxes. This is partly due to the fact that the mandate has a three-year phase-in, while the taxes do not. A larger factor is the fact that the mandate requires that each vehicle use a compliant retrofit or be replaced. There is never any incentive for the fleet manager to go beyond the least expensive means of achieving compliance, even if doing so could achieve relatively high emission reductions per dollar spent.

When comparing the different compliance scenarios, it makes sense to focus on the lower left portion of the cost curves for two reasons. First, the actions taken in the upper portions are largely (very) early replacements, which are relatively independent of retrofit compatibility. Second, as the next section will discuss further, the upper portion of the curve is very costly and likely far from optimal policy.

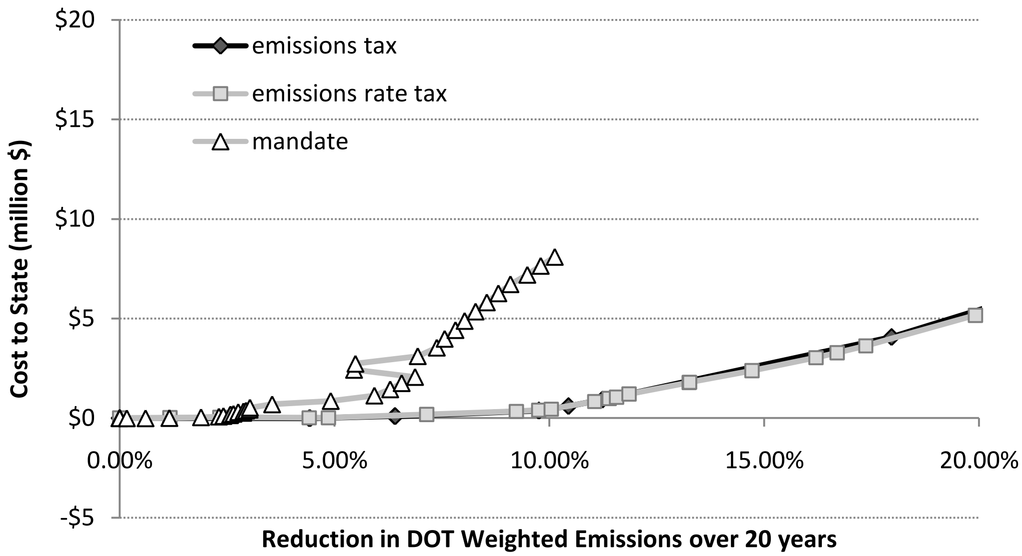

Figure 15. Cost to State vs. Emission Reduction in Compliance Scenario 1 (Zoomed-in)

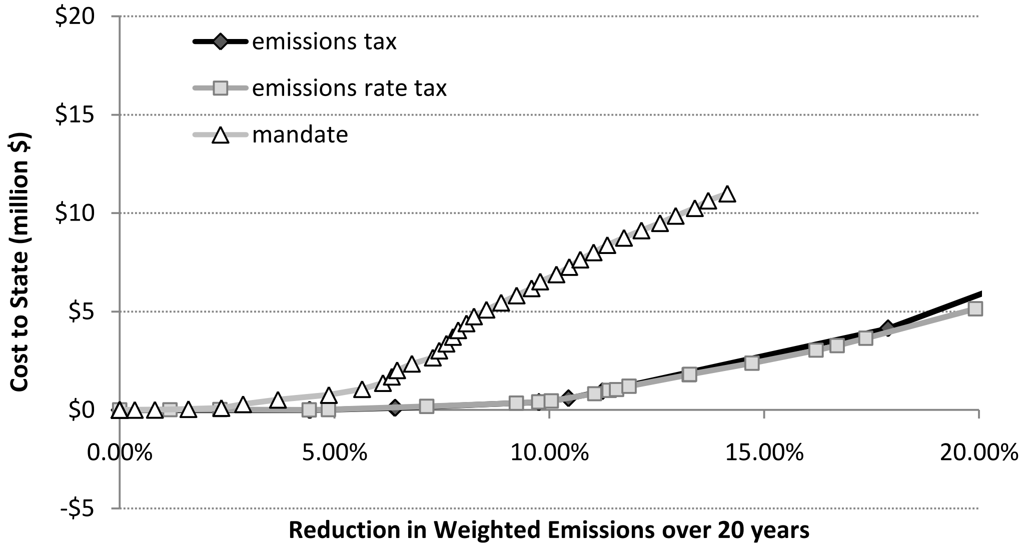

Figure 16. Cost to State vs. Emission Reduction in Compliance Scenario 2 (Zoomed-in)

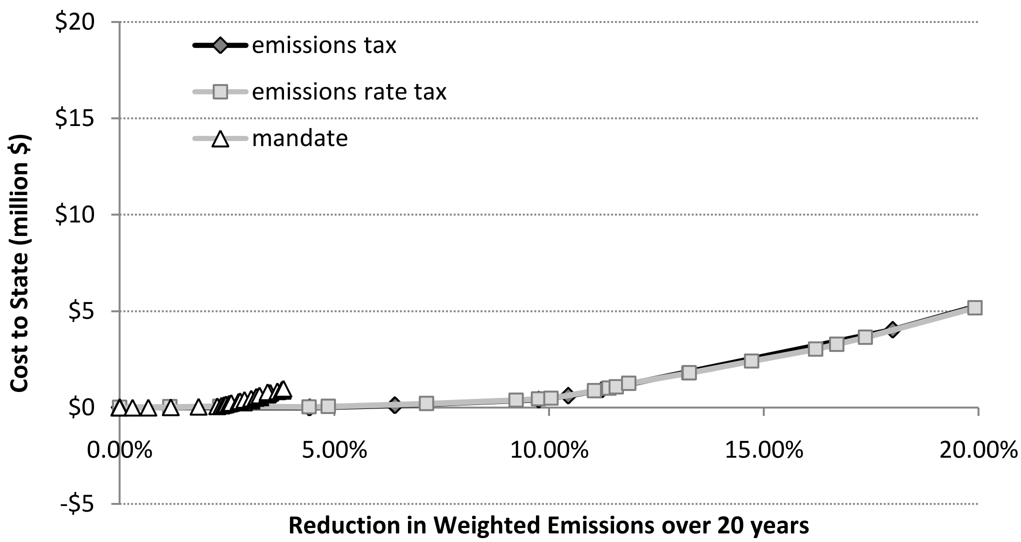

Figure 17. Cost to State vs. Emission Reduction in Compliance Scenario 3 (Zoomed-in)

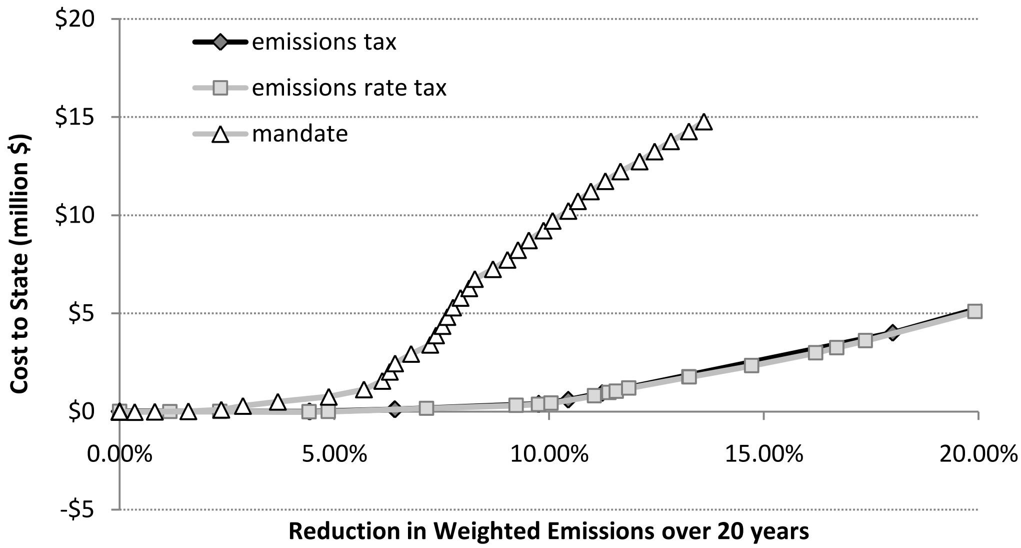

Figure 18. Cost to State vs. Emission Reduction in Compliance Scenario 4 (Zoomed-in)

Regardless of the compatibility scenario, it is apparent that for very small reductions, the mandate has costs which are relatively close to those of the taxes. This is because modest levels of all three instruments cause the fleet manager to replace old vehicles slightly faster. The mandate becomes noticeably more expensive when the fleet manager starts to apply retrofits earlier and more broadly than she would under emissions or emission rate taxes. This is not to say that retrofits are universally inefficient. The efficiency depends heavily on the application. The degree to which these additional retrofits cause the mandate to grow in cost depends heavily on the compatibility scenario. More compatibility is not always better. The sharpest increase comes when vehicles are compatible with the DOCs and expensive DPFs, but not the cheap DPFs. This is because the fleet owner must use the expensive DPFs, which offer the lowest emission reduction per dollar spent. The cheapest scenario is the one with the lowest compatibility because the fleet owner can use DOCs. Naturally, the scenario in which vehicles can use cheap DPFs and the mixed scenario are in the middle. NYS DOT responded to the regulation largely by arguing that their vehicles could use only DOCs, and by installing DOCs.

An element of the mandate cost curve in the mixed scenario (Figure 15) jumps out. At one point, increased mandate intensity causes an increase in cost while increasing emissions. This is not typical of such curves, and seems somewhat counterintuitive. One is prompted to ask what the extra cost paying for. The complicating issue is the metric used to measure emission reductions. While the mandate had clauses relating to NO_x, it was heavily focused on reducing PM. PM reductions actually do continue to decrease at this point, but NO_x emissions jump upward enough that the weighted combination increases. This occurs as the fleet manager transitions into heavier use of retrofits, which reduce PM but not NO_x. Early replacements reduce both. There is no consensus on the best way to measure emission reductions. As described in the next section, the weights used are based on those used by CARB's Carl Moyer Program.

4.2.3 Selecting the Optimal Level of Regulation

Given the cost curves in the previous section, it is natural to ask what level of regulation is optimal. In order to answer this question, it is necessary to estimate the value of emissions reductions. This is a hotly debated topic, with no clear consensus. Because the value of emission reductions is not the focus of this dissertation, values will be inferred from current government policy. In particular, the values will be drawn from the cost effectiveness threshold used by CARB's Carl Moyer Program, which funds diesel retrofits and early replacements. CARB won't fund a project through the program unless it costs less than \$16,000 per weighted ton of NO_x, reactive organic gases (ROG), and PM_{10} reduced (CARB, 2008b). The weights for

NO_x and ROG are 1, while the weight for PM_{10} is 20 because it has been identified as a toxic air contaminant.

The Carl Moyer Program criterion implies that CARB believes that reducing weighted emissions by one gram is worth at least $0.0176. This can be used to value emission reductions and estimate optimal levels of regulation. It is possible that limited funds force CARB to increase its cost effectiveness cutoff beyond what its staff believes to be socially optimal. To the extent that this is true, estimates of optimal regulation intensity will be lower bounds on the truly socially optimal levels.

Tables 7-9 list the optimal regulation levels if weighted emissions reductions are valued at $0.0176 per gram. None of the instruments have purely convex cost curves (though emissions and emission rate taxes come close), meaning that simply setting marginal cost equal to marginal benefits can yield multiple solutions. In such situations, the solution with the highest net social benefit was selected. Net social benefit was defined as the value of emissions reductions less the extra financial cost to the state. All optimal levels are approximate due to the fact that only a finite number of potential levels were evaluated.

Note that fleet manager behavior is not a continuous function of the tax or mandate intensity. Behavior changes discontinuously when policy intensity thresholds are passed. As a result, ranges of policy intensity can yield the same behavior and the same results. This means that there may be optimal tax or mandate ranges, as opposed to a single optimal level as is commonly expected.

Table 7. Optimal Emission Tax Levels

Compatibility Scenario	Tax Level ($/g)	% Emission Reduction	Net Social Benefit
1	$0.0075-$0.0125	9.76%	$2,718,181
2	$0.0150-$0.01525	10.45%	$2,714,576
3	$0.0150-$0.01525	10.45%	$2,711,225

4	$0.0150-$0.01525	10.45%	$2,720,989

Table 8. Optimal Emission Rate Tax Levels

Compatibility Scenario	Tax Level ($/(g/mi))	% Emission Reduction	Net Social Benefit
1	$120-$160	10.05%	$2,741,868
2	$120-$160	10.05%	$2,733,596
3	$120-$160	10.05%	$2,703,362
4	$120-$160	10.05%	$2,756,366

Table 9. Optimal Mandate Levels

Compatibility Scenario	Mandate Level (Final %)	% Emission Reduction	Net Social Benefit
1	66%	5.93%	$745,412
2	42%	4.87%	$786,045
3	46%	2.27%	$652,293
4	42%	4.87%	$785,441

The most prevalent trend in the optimal regulation levels is that optimal taxes cause substantially larger emission reductions and higher net social benefits. The differences between the net benefits of the two types of taxes are relatively small. Contrary to what one might expect from theory, the optimal net benefit from emission rate taxes can be slightly higher than that from an emissions tax.

Theory generally tells us that the optimal policy is a Pigouvian emission tax equal to the marginal value of emission reductions. The integer programming results indicate that the optimal emission tax is not Pigouvian (it is lower than the Pigouvian tax), and that the optimal tax structure is not necessarily even an emissions tax. This type of finding is not unprecedented, as several papers have found that Pigouvian taxes can be suboptimal if there are other market distortions (Parry, 1994; Parry and Bento, 1999). In this example, the second market distortion is that the fleet owner does not keep resale revenue, meaning that the marginal cost of vehicle replacement from the fleet owner's point of view is not the true marginal cost. This distortion impacts behavior with and without environmental regulation. To make matters more complicated, the distortion only applies to resale and not to retrofits, meaning that the fleet manager might not only abate the "wrong" amount, she might abate in the "wrong" way as well. "Wrong" here means socially suboptimal, not suboptimal from the fleet owner's point of view. Because this distortion can cause suboptimal behavior even in the absence of environmental regulation, the financial costs of small emissions reductions can even be negative, as in Parry and Bento (1999) with the "second dividend" coming from resale revenue received by the state. In other words, very modest environmental regulation which encourages earlier vehicle replacement may actually save the state money, because the current handling of auction revenue causes suboptimal replacement decisions.

4.4 Strengths and Limitations of Integer Programming

The primary strengths of the integer programming approach are the breadth of objectives and constraints it can consider, and the speed with which it converges. Even for relatively large fleets with dozens of vehicle types and a half dozen retrofit options, it rarely took CPLEX11.2.1 more than a few seconds to solve the integer program. This makes it easy to rerun the program hundreds of times, generating the graphs in the previous section. The integer program appropriately represents the discrete nature of the decisions involved and it can be applied to steady-state situations as well as periods of transition.

Unfortunately, the integer programming approach is not well suited to dealing with uncertainty. Many forms of uncertainty are present when making retrofits and replacements. Future prices are unknown. From the regulator's perspective, compatibility is unknown. From the fleet manager's perspective, future regulation may be unknown. Perhaps most importantly, vehicle lifetimes are unknown, as are the timing and severity of future breakdowns. To some degree, some of these uncertainties can be addressed through sensitivity analysis, as was conducted using multiple compatibility scenarios in the previous section. Sensitivity analysis alone will not reveal how to best mitigate risk, however.

The issue of future price uncertainty was addressed with a nonlinear program presented in Gao and Stasko (2009b). The objective was to minimize the expected lifetime cost of vehicles purchased over a period of time. The number of vehicles to be purchased at each point in time was considered fixed, as was the lifetime of each vehicle, but the optimizer could choose vehicle technologies. A constraint limited "well-to-wheel" emissions, and another constraint limited the variance of the cost. The quadratic program was able to quickly balance expected cost, emissions, and financial risk to produce tradeoff curves. Much the cost uncertainty came from fuel price fluctuations, making fuel efficient vehicles a financially conservative choice. While the model was capable of taking advantage of low or negative price correlations to reduce risk, there was not a lot of opportunity to do so among the technologies considered (conventional gasoline, diesel, E85, and gasoline-electric grid-independent hybrid). The fuel prices were all strongly correlated. The integrality of decisions was dropped in order to allow for an interior point solver to quickly produce solutions, limiting the applicability small fleets. While this model has potential for dealing with technology selection, it provides no insights into when vehicles should be replaced or retrofitted.

Uncertainty in vehicle breakdowns and usable lifetimes is both important to replacement and retrofit decisions and extremely difficult to represent within an integer programming framework. The first section of Chapter 5 will discuss this in more detail. In general, the integer programming approach can be useful for quickly revealing aggregate trends, which would likely be the focus of the regulator. Because the repair costs for a given vehicle will often deviate substantially from the average, integer programming may not always be terribly helpful for prescribing what to do with each vehicle. From a fleet manager's perspective, an approach which acknowledges the uncertainty surrounding vehicle breakdowns and repair costs could prove more fruitful.

CHAPTER 5

APPLYING APPROXIMATE DYNAMIC PROGRAMMING

5.1 Choosing an Approach

Dynamic programming shares many strengths with integer programming. It is well equipped to handle the discrete nature of retrofit and replacement decisions. It can also be quite flexible in terms of including multiple kinds of constraints, depending on how the program is solved. The primary advantage of dynamic programming is its ability to handle stochastic breakdowns and repair costs.

The importance of stochastic breakdowns and repair costs is illustrated by Figure 19. It plots the probability that a vehicle will still be in the fleet at a range of ages. It assumes a steady-state situation in which there is no retrofit regulation. As discussed in Chapter 1, when breakdowns and repair costs are deterministic, there is a single optimal replacement age. Deterministic integer programs replace all vehicles at the same age, as do deterministic dynamic programs. The red curve, which is derived from the actual auction dates of 331 class 8 International 2574 dump trucks, indicates that vehicles are not in fact retired at a consistent age. In this case, they are phased out over a period of roughly five years. Most vehicles are replaced when they would need a major repair to remain operational.

The stochastic approximate dynamic program (ADP) result strongly resembles the actual replacement pattern. The stochastic ADP result does not match current behavior perfectly, nor would one necessarily expect it to. Current practice is not necessarily optimal, but the fact that current practice is so dramatically different from the integer program result calls the assumption of deterministic breakdowns and repair costs into question. Breakdowns and repair costs are, in fact, stochastic. The replacement theory literature discussed in Chapter 1 confirms that fixed retirement ages are generally not optimal where repair costs are stochastic. It is therefore highly desirable to have a model which can handle this stochasticity and develop optimal policies which take it into account.

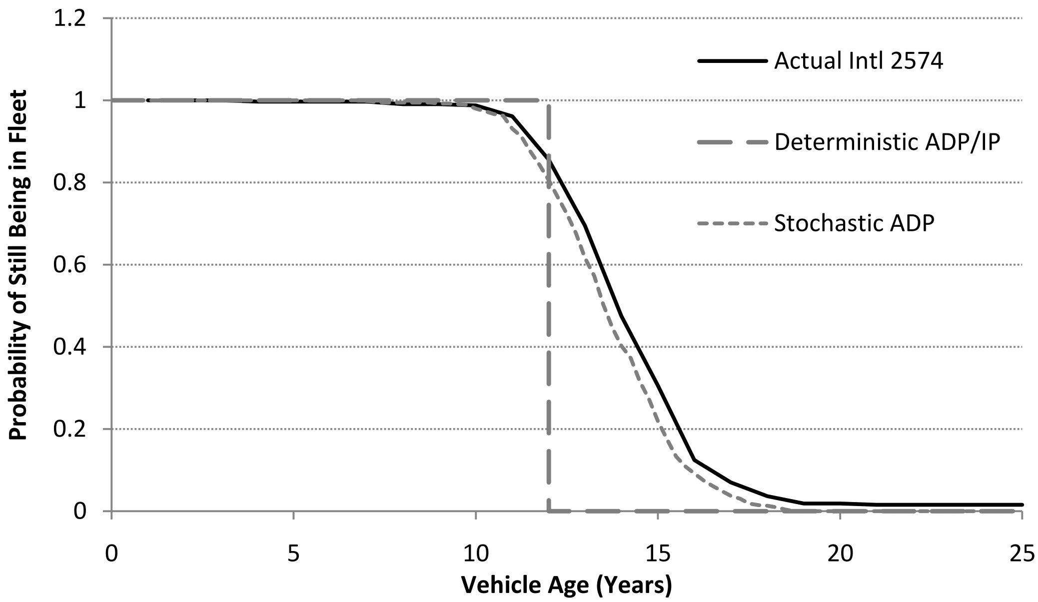

Figure 19. Probability of Remaining in Fleet until a Given Age

The stochasticity of repair costs is challenging to represent because it cannot easily be captured by a set of predetermined scenarios. Stochastic breakdowns prompt the fleet manager to react (perhaps by repairing or replacing the vehicle). These reactions then influence the probabilities of future breakdowns. The feedback effect means that the distributions from which breakdown events are drawn are not known until the actions taken in previous periods are known. This type of uncertainty can be well represented by stochastic dynamic programming. Furthermore, stochastic dynamic programming is capable of representing the fleet manager's changing access to information over time. When making decisions at a given point in time, the fleet manager knows the current state of the fleet, as well as something about the likelihood of future maintenance events. This continuously changing knowledge is easily captured by a stochastic dynamic programming framework, but it is extremely difficult to handle with integer programming.

Once the fleet upkeep problem is framed as a stochastic dynamic program, the next step is to select a method for solving the program. It is well known that dynamic programs grow quickly with the dimension of the state space. The state space is the set of possible conditions the fleet can be in. It is defined not just by the size of the fleet, but by the age of each vehicle, as well as any other relevant characteristics of each vehicle, such as the retrofits installed and the repair status. Naturally, the state space can quickly become enormous. The resulting computational difficulty is often referred to as dynamic programming's "curse of dimensionality." Powell (2007) argues that there are in fact two additional curses of dimensionality, one for the action space (representing different feasible combinations of decisions for a given period) and one for the outcome space (representing different potential combinations of random variable realizations for a given period).

A relatively simple example fleet might have 25 possible ages, three repair statuses, and two retrofit statuses (compliant and non-compliant). The state space could be represented using 150 variables, one for each vehicle category. If no category ever has more than 19 vehicles, there are a whopping 20^{150} possible states of the fleet. In the test problems, the average time to compute the value of a state was at least 0.01 seconds. At this pace, it would take roughly $4.5*10^{185}$ years to evaluate the value of every state for a single period, which is significantly longer than the estimated of the age of the universe (NASA, 2009).

Traditional backwards dynamic programming requires computing the value of every such state in every time period, an obviously infeasible task. Alternatively, dynamic programs can be reformulated as linear programs with variables for each state and constraints for each state-action combination (Bertsekas, 1987). Given the vast size the of the state and action spaces, however, this approach also has limited applicability. Both memory and computation time requirements would quickly become astronomical as the problem was scaled up from toy examples.

Approximate dynamic programming techniques are often able to produce high quality solutions, despite examining only a small fraction of possible states. Among ADP approaches, value iteration is particularly well suited for the fleet upkeep problem, because retrofit constraints will generally change over time. Policy iteration, an alternative, is popular for steady-state infinite horizon problems (Powell, 2007). The next section will present a customized value iteration algorithm for parallel asset replacement and retrofit problems, while the following sections will test and apply it in the context of the NYS DOT's fleet management problem.

5.2 The ADP Algorithm

5.2.1 Overview and Value Function Definition

The core of any dynamic program is its recursion equation. A stochastic version of Bellman's equation is given by expression (28) where $V_t(S_t)$ is the value of being in state S_t at the start of period t (assuming optimal behavior), $N_t(S_t,x_t)$ is the net benefit experienced at the start of t when taking action x_t in state S_t, ρ is the discount factor, and E[] designates the expected value. Notation for the ADP is different from notation for the integer program presented in Chapter 4.

$$V_t(S_t) = \max_{x_t} \{N_t(S_t, x_t) + \rho \mathrm{E}[V_{t+1}(S_{t+1}(S_t, x_t)) | S_t]\} \tag{28}$$

In the fleet upkeep problem, the state space describes the set of possible conditions the fleet could be in at any given point in time. The state of the fleet is defined by a set of integer variables f_{ajk}, each indicating how many vehicles exist in a relevant category. A category is defined by a vehicle age a, maintenance status j, and retrofit status k. Actions described by x_t are vehicle purchases, sales, repairs, and retrofits. N_t is the vehicle sales revenue (if kept by the fleet owner) minus the costs due to other actions. Uncertainty stems from the fact that future maintenance statuses and vehicle failures are not known. Thus, the expectation is taken over possible vehicle maintenance statuses and vehicle failure combinations.

An outline of the value iteration approach employed is provided in Table 10. Forward passes through time act as sequences of simulation steps and optimization steps, capturing random effects and acting in response to them. With each step, estimates of vehicle values are improved. This approach allows the optimizer to focus on understanding regions of the state space which are of greatest importance, and to extrapolate based on the findings.

Table 10. Outline of Value Iteration Approach

1	Initialize. Input data on current fleet status, future demands, and future retrofit regulation. Set period = 1.
2	Solve single-period IP using network flow LP formulation defined by expressions (29)-(33).
3	Update the value function approximation by using expressions (34) and (35) to adjust the value of each vehicle category in the previous period.

4	If not the last period: a) Using the transition function described in Section 5.2.9, update fleet status based on manager actions (repairs, sales, purchases, retrofits), and then based on random breakdown events. b) Move to the next time period. Update demand and retrofit requirements. c) Go to step 2. If the last period: a) Update the final fleet values to equal the average over the last several periods assumed steady-state). b) If not final iteration: i. Reset to initial fleet status and period 1 demand/retrofit requirements. ii. Go to step 2. a) If final iteration: i. Go to step 5.
5	Compute performance metrics. Output results.

Developing an appropriate value function form is one of the key challenges when formulating an ADP. There is a tremendous amount of flexibility. On one extreme, a value can be defined individually for each state, without any functional form. While this method is free of functional form restrictions, it is slow to improve. The value of a given state won't be updated until it is visited, the odds of which are extremely low even in a mid-sized ADP. There are, of course, serious memory capacity issues as well. On the other extreme, the value function could be a simple function of a very small number of parameters. This approach causes new knowledge to be broadly applied quickly, but a simple form may not be able to accurately capture the true values of diverse states.

This model employs a linear value function which assigns a value to each vehicle category, and allows these values to change over time (e.g. when regulatory mandates take effect). The value of the fleet is simply the sum of the values of the vehicles it contains. This functional form allows for subproblems to be solved efficiently, and the parameters that define it have intuitive meaning. This makes results easily interpretable, and facilitates identification of errors in implementation.

5.2.2 Network Flow Description

Given expression (28), the next step is to determine how to solve for the optimal set of actions, x_t^*. Largely because of the form selected for the value function approximation, this problem can be modeled as an integer program. The objective and constraints are all linear in the decision variables. This single-period integer program effectively forms the policy used to make decisions at a given point in time.

In order for the ADP to converge, a considerable number of single-period subproblems need to be solved. IPs are NP-hard, and no polynomial time algorithm for solving them is known. There are well known polynomial time algorithms for solving linear programs (LPs) without integrality constraints, as well as a worst-case exponential algorithm (known as the simplex method) which works very well in practice (Kleinberg and Tardos, 2006). For this reason, it is natural to seek a linear program formulation which will yield integer solutions.

There are several classes of network flow problems for which the simplex algorithm produces integer solutions. The minimum cost flow problem is one such problem class (Sierksma, 1996), and it can be used to model the single-period asset upkeep problem. This is possible because the discrete nature of vehicles will yield only integer supplies, demands, and upper bounds on link flows. A network illustration of a simple single-period fleet upkeep problem is presented in Figure 20.

Vehicles in the initial fleet flow from source S1, while potential new purchases flow from source S2. All flows terminate at sink T. Costs on the edges produce the proper objective, and the capacity constraints combine with conservation of flow to construct the proper feasible region. For example, the capacity of the link from ``Not

Available" to T ensures there are enough vehicles available to meet demand. Extensions such as allowing purchases of non-compliant vehicles or multiple maintenance states complicate the picture, but they can still be represented as a network flow problem. A linear program which allows for such extensions is outlined below. This is the core LP formulation used for single-period subproblems of the ADP, but several variants were developed to allow for other extensions. When extending the LP, it is important to be aware of implications on how value functions are updated (see Section 5.2.8).

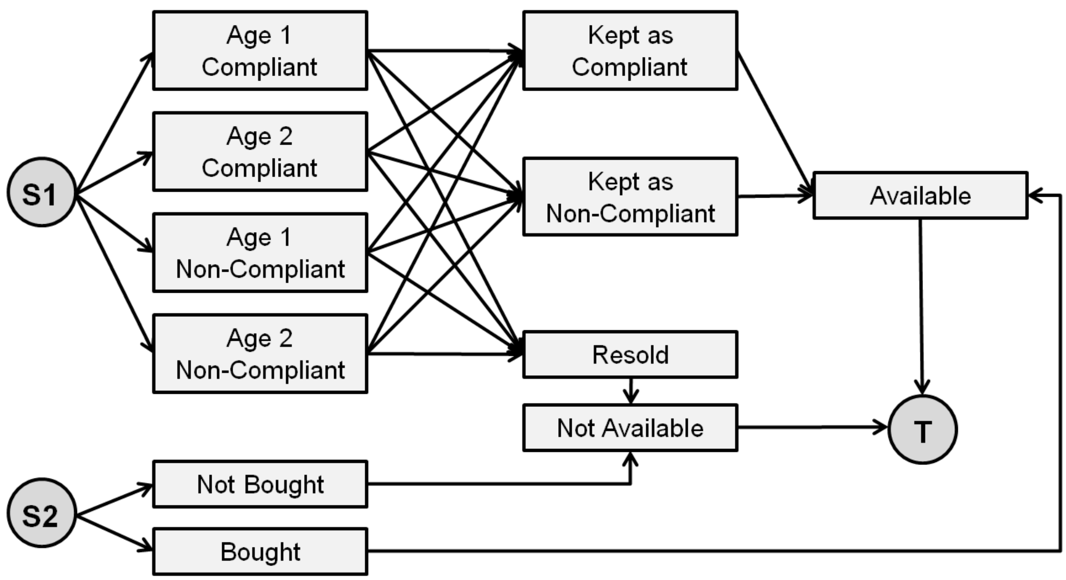

Figure 20. Network Flow Representation of a Simple Single-Period Fleet Upkeep Problem

The objective, given by expression (29), is to maximize the discounted future value of the fleet, plus vehicle sales revenue from the current period, minus costs from the current period (e.g. repairs, retrofits, and purchases). Expression (30) is a constraint requiring conservation of flow for vehicles in the existing fleet. Expression (31) requires that there are enough vehicles to meet demands. Expression (32) caps the number of vehicles bought or kept in each retrofit status. Expression (33) caps vehicle purchases. This cap is used to create a network flow formulation, and is assumed to be high enough that the problem remains feasible.

5.2.3 LP Formulation: Sets

A set of vehicle ages

J set of maintenance statuses

K set of retrofit statuses

T set of time periods (assumed to have same resolution as vehicle ages)

5.2.4 LP Formulation: Input Parameters

f_{ajk} number of age a vehicles in maintenance status j and retrofit status k at the start of the period

c_{ajkd} cost of keeping an age a vehicle currently in maintenance status j and retrofit status k, to be put in new retrofit status d

p_k price of a new vehicle in retrofit status k

u_k maximum number of new vehicles in retrofit status k which can be purchased

r_{aj} net resale revenue for a vehicle of age a in maintenance status j

v_{ak} discounted future value of a kept vehicle of age a in retrofit status k

w_k discounted future value of a bought vehicle in retrofit status k

ϕ demand for vehicles in current period which must be met

ψ_k maximum number of vehicles in retrofit state k to be held or bought

5.2.5 LP Formulation: Decision Variables

g_k number of new vehicles in retrofit status k to be bought

h_{ajkd} number of vehicles of age a vehicle currently in maintenance state j and retrofit status k, to be kept and put in new retrofit status d

q_{ajk} number of vehicles of age a vehicle currently in maintenance state j and retrofit status k to be resold

5.2.6 LP Formulation: Objective

$$\max \sum_{k\in K}\left\{(w_k - p_k)g_k + \sum_{a\in A}\sum_{j\in J}\sum_{d\in K}(v_{ak} - c_{ajkd})h_{ajkd}\right\} + \sum_{a\in A}\sum_{j\in J}\sum_{d\in K} r_{aj}q_{ajk} \tag{29}$$

5.2.7 LP Formulation: Constraints

$q_{ajk} + \sum_{d\in K} h_{ajkd} = f_{ajk}$ for all {a in A, j in J, k in K) (30)

$\left\{\sum_{a\in A}\sum_{j\in J}\sum_{k\in K}\sum_{d\in K} h_{ajkd}\right\} + \sum_{k\in K} g_k \geq \phi$ (31)

$g_d + \sum_{a\in A}\sum_{j\in J}\sum_{k\in K} h_{ajkd} \leq \psi_d$ for all {d in K} (32)

$g_k \leq u_k$ for all {k in K} (33)

5.2.8 Updating the Value Function

Once optimal actions are determined for the current period, the next step is to update the value function estimate. Vehicle shadow prices from the current period are used to update the vehicle value estimates used in the previous period's LP. Naturally, these improved value estimates won't be used until the next forward pass.

The value function in iteration *n* is defined by a set of parameters, $v_{\mathrm{akt}}^{\mathrm{n}}$, indicating the expected discounted future value of keeping a vehicle of age *a*, and retrofit status *k*, at the start of period *t*, as well as a set of parameters, $w_{\mathrm{kt}}^{\mathrm{n}}$, indicating the expected discounted future value of a new vehicle in retrofit status *k*, bought at the start of period *t*. Because $v_{\mathrm{akt}}^{\mathrm{n}}$ and $w_{\mathrm{kt}}^{\mathrm{n}}$ are the means of a random variables, it does not make sense to ignore the previous estimates whenever new observations are found.

Instead, the new estimates are weighted combinations of the old estimates and the discounted average of shadow prices from period t+1, as given by expressions (34) and (35). The shadow prices are averaged over the different maintenance states (including complete failure).

$$v_{akt}^{n} = (1-\alpha_{n-1})v_{akt}^{n-1} + \alpha_{n-1}\rho\left[\tau_{(a+1)}\varphi + \sum_{j\in J}\pi_{(a+1)j}\,\lambda_{(a+1)jk(t+1)}^{n-1}\right] \tag{34}$$

$$w_{kt}^{n} = (1-\alpha_{n-1})w_{kt}^{n-1} + \alpha_{n-1}\rho\left[\tau_{1}\varphi + \sum_{j\in J}\pi_{1j}\,\lambda_{1jk(t+1)}^{n-1}\right] \tag{35}$$

In expressions (34) and (35), φ is the resale revenue of a vehicle which has failed beyond repair, while τ_a is the probability of complete failure for a vehicle of age a, π_{aj} is the probability of being in maintenance status j for a vehicle of age a, and λ_{ajkt}^{n} is the shadow price for a vehicle of age a in maintenance status j and retrofit status k during period t of the nth iteration.

Selecting appropriate alpha values for step sizes is critically important. Both theory and experience can guide step size selection. Theory comes from conditions for convergence proofs of stochastic gradient algorithms. Step sizes must be nonnegative, and their infinite sum must be infinite while the infinite sum of their squares must be finite. These rules essentially require that step sizes decline according to a something resembling harmonic sequence. Experience, on the other hand, indicates that a simple $\alpha_{n-1} = 1/n$ step size rule drops too quickly (Powell, 2007). As a result, the current

ADP implementation uses a well known step size rule which is based on the generalized harmonic sequence given in expression (36).

$$\alpha_{n-1} = \frac{\delta}{\delta+n-1} \tag{36}$$

In order to implement expressions (34) and (35), it is necessary to more precisely define shadow prices, and develop a method for estimating them. In general, a shadow price is the rate of change in the optimal objective function with respect to change in the amount of one resource. A simple means of obtaining shadow prices is to add or subtract a unit of the resource in question, resolve the LP, and compare objective values. While reliable, this method can be very time consuming when many shadow prices are required.

In linear programming, dual variables are commonly used to determine shadow prices. Unfortunately, obtaining shadow prices isn't always as simple as outputting the dual variables corresponding to the optimal solution. Classical linear programming texts have been criticized for misleading readers about the equivalence of shadow prices and dual variables (Akgül, 1984). The equation of dual variables and shadow prices is based on the assumption of non-degeneracy. If the optimal primal solution is degenerate, however, there may be alternative dual values, meaning that the shadow prices are no longer necessarily equal to the set of dual variables output by the solver (Lin, 2010). Even a simple fleet upkeep problem with only a few vehicles can exhibit primal degeneracy. This can cause non-compliant vehicles to be erroneously assigned the same shadow price as compliant vehicles, significantly impacting results.

The operations research community has been struggling to deal with shadow prices of degenerate LPs for some time. Various approaches require solving different, albeit smaller LPs for each shadow price sought (Akgül, 1984; Lin, 2010).

Additionally, it has been proven that if the set of optimal dual solutions is y in D^*, then:

$$\lambda_z^+ = \min\{y_z : y \in D^*\} \tag{37}$$

$$\lambda_z^- = \max\{y_z : y \in D^*\} \tag{38}$$

where λ_z^+ is the shadow price for an additional unit of resource *z*, while λ_z^- is the shadow price of the last unit of resource *z*, and the primal problem is a maximization (Lin, 2010). The min and max operators are switched for a primal minimization problem (Roos et al. 1997). Essentially, this means that dual variable for a particular vehicle is an upper bound on the value of another such vehicle.

Performing degenerate pivots around the optimal solution can provide multiple dual solutions, which could provide progressively tighter bounds on shadow prices. Various techniques for doing so were explored, and were successful at tightening bounds, but proved to be too time consuming to be worth the effort, given the large number of possible pivots and the large number of dual variables of interest.

The ADP presented uses dual variables as upper bounds on shadow prices, and it constructs lower bounds for comparison. Lower bounds are constructed by considering what could be done with the additional vehicle, and constructing the corresponding paths through the network. If kept, the vehicle might be retrofitted, and it might eliminate the need for a new vehicle purchase, depending on which constraints are binding. Alternatively, the vehicle could be sold. If upper and lower bounds are sufficiently close (within $100 for sample problems), then the average of the bounds is used for the shadow price. Otherwise, the actual shadow price is determined by perturbing the right-hand-side vector and resolving the LP.

This hybrid approach to shadow price estimation proved far more accurate than depending on dual variables and far faster than perturbing the right-hand-side vector in every case. On relatively small sample problems, pure perturbations took more than 30 times longer than the hybrid approach, which only needed to perform perturbations roughly 1-10% of the time. Larger sample problems could not be solved using pure perturbations in a reasonable time frame, but were solvable using the hybrid approach. Despite being relatively small, and having ``warm starts," the perturbed LPs associated with each shadow price are time consuming to solve, and are better used as a last resort than as a standard approach.

5.2.9 Transitioning between States

Recall that each iteration of the ADP is a simulated forward pass through time. Once actions are determined for a given period, and the value function for the past period has been updated, the next step is to move forward to the next period. This is accomplished with the transition function. Using several nested loops, the transition function generates random variables describing maintenance events, and produces the pre-decision state of the fleet for the next time period. All vehicles are aged by one period. Some vehicles fail completely and are sold for scrap. Those remaining are randomly assigned a maintenance status according to the appropriate probabilities.

The transition function also estimates usage levels for all of the vehicles which were kept, and calculates an emissions inventory based on these usage levels. The best manner in which to do so is problem specific, but in the case of NYS DOT there is no reason to believe usage levels won't continue to follow a pattern similar to recent history. Average annual mileage for class 8 dump trucks is plotted as a function of age in Figure 21. There is clearly a relatively linear trend of decreasing use with age. Lower repair and maintenance costs encourage higher usage of newer vehicles to the extent possible, and a technology mandate won't change that. It is possible that the size or age distribution of the fleet could change, so it is not reasonable to assume that the average mileage of each vehicle age will remain the same. Such an assumption could cause the model to predict that a younger fleet would automatically be driven much more, even if the demand remained the same.

Instead, it is assumed that the relative levels of usage will remain the same. For example, a six year old truck averages roughly 80% of the mileage of a one year old truck, while a ten year old truck averages roughly 63% of the mileage of a one year old truck. These ratios are held constant. Based on these ratios, the usage of each vehicle is dynamically determined for each period, as a function of the makeup of the entire fleet at that point in time.

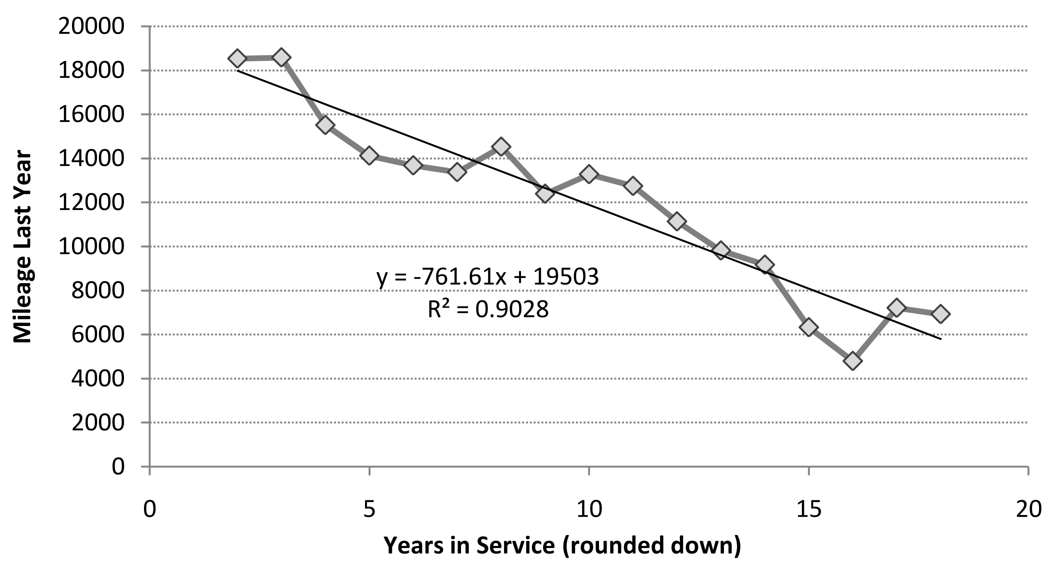

Figure 21. Average Annual Mileage by Age for Class 8 Dump Trucks

5.3 Testing and Convergence

The convergence properties of approximate dynamic programs have been studied by numerous authors. It has been demonstrated that under a fairly broad set of conditions these algorithms do converge eventually (Powell, 2007). That is to say, ADPs will converge as the number of iterations approaches infinity. This is somewhat reassuring, and the proofs can be useful in guiding the structure of the formulation, as they did with the step size formula described in Section 5.2. At the same time, these proofs are of practically no value when it comes to predicting exactly how long convergence will take for a given application. If convergence will take an infinite amount of time, one might as well use backward induction, which would take an extremely long (yet finite) amount of time. Because the convergence pattern can vary depending on the algorithm and application, this section tests convergence on a series of problems based on the NYS DOT fleet. All example problems include well over 1,000 vehicles and use time steps of three months over a period of twenty five years. All test IPs and ADPs were run on the same desktop computer, which used a dual core 3.0GHz Intel Xeon processor and roughly 3GB of RAM.

Perhaps the most obvious metric by which to measure convergence is the speed by which the objective approaches the optimum value. In particular, one could measure the time it takes to find a solution within 1% of the optimal objective. For the stochastic case, this is challenging to test because the optimum objective is generally not known. For the deterministic case, on the other hand, a branch-and-bound procedure can produce an optimal solution, or at least one which is provably extremely close to optimality. The tolerance for CPLEX was set so that IP solutions were always within 0.01% of optimality.

A deterministic version of the problem was constructed for the purpose of comparing the policies found by the ADP with those found by solving a single large

IP. The problem included the introduction of regulation followed by a period of steady-state. In order to solve this problem with an IP, the objective was changed slightly so that the fleet remaining at the end of the simulation was sold at exogenous market prices, instead of valued at endogenously determined prices. This preserved linearity in the IP, which made it possible to use standard linear IP solvers like CPLEX.

Convergence rates naturally depend on the quality of the initial asset value guesses. While there are a number of reasonable ways to generate initial guesses, this trial used a deliberately naïve initial guess that all vehicles are worth $80,000, independent of age and condition. Nonetheless, the ADP produced a solution within 1% of the CPLEX11.2.1 IP optimum by the 23rd iteration, which took approximately one and a half hours. The discounted net cost is plotted as a function of the iteration in Figure 22, with the IP optimum designated by a dashed line. In the absence of an IP optimum for comparison, several factors can offer clues that the ADP has reached an optimum. Perhaps the most obvious sign is a decrease in the rate of improvement of the objective function, as is clearly the case in Figure 22. Slowed improvement can be deceiving however, and may not indicate optimality. It is possible that the step sizes have simply declined to the point where the value function is changing too slowly to noticeably improve the objective.

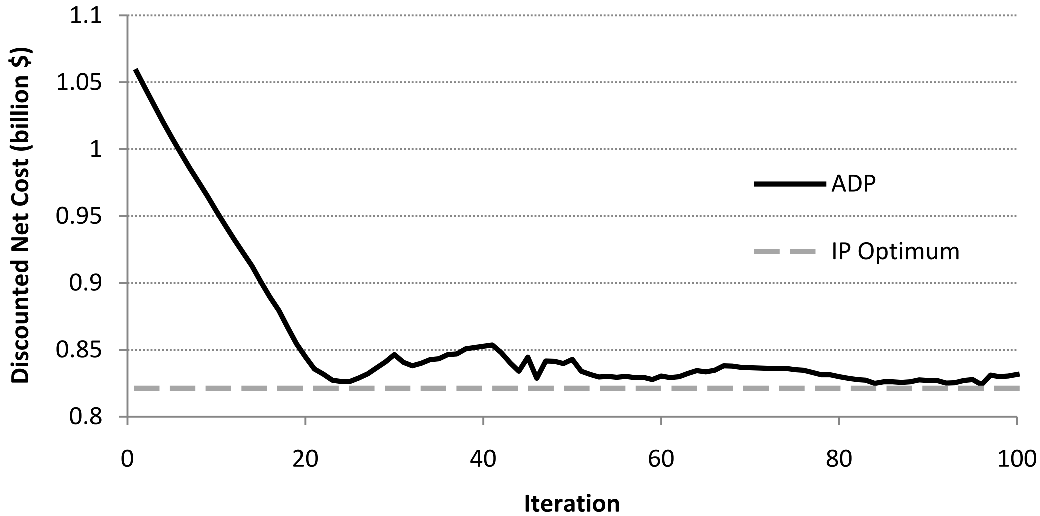

Figure 22. Convergence of the Objective in a Deterministic Example

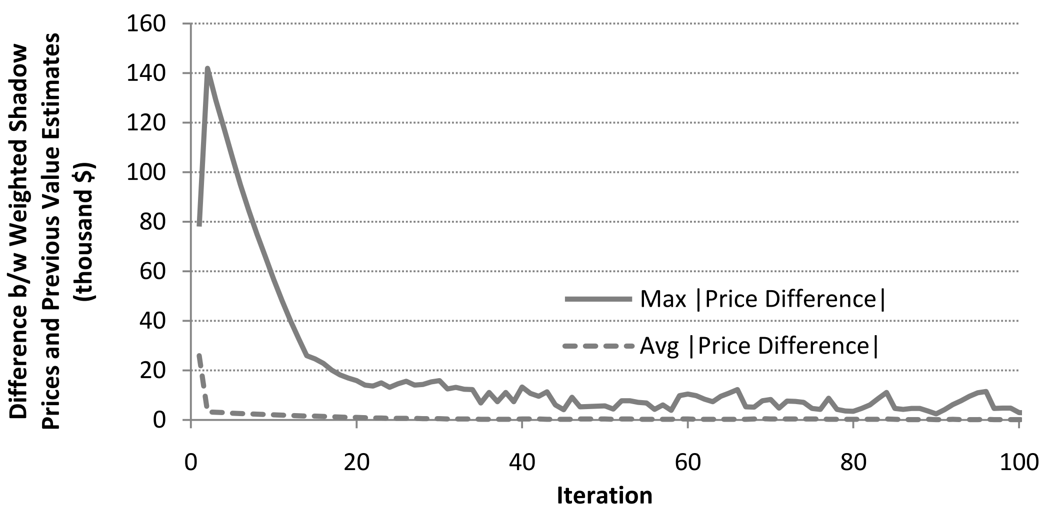

Figure 23. Convergence of Gap in Value Estimates

In order to avoid step size issues, one can directly compare shadow prices to the previous iteration's estimates of vehicle values. The difference between the estimate of a vehicle's value based solely on current shadow prices and the previous iteration's estimate of that vehicle's value is an upper bound on how much the corresponding value function parameter can change in the next iteration. Figure 23 plots the maximum and average absolute differences between value estimates based on current shadow prices and previous value estimates. At the start of the ADP, the average absolute difference is in the tens of thousands dollars, with the maximum absolute difference topping $100,000. By the 100th iteration, the average absolute difference is a little over $100, and the maximum is a few thousand dollars. The fact that the value function is not going to change dramatically, regardless of the step size, provides a helpful hint that the ADP has converged, but it does not equate to a guarantee.

If the ADP is initialized with more reasonable vehicle value guesses, the convergence can be much faster. For example, the analytical deterministic steady-state values were used as initial guesses, based on expressions (1) and (2). The solution was not terribly close to optimal on the first iteration because several assumptions of analytical solution were violated (e.g. unchanging technology, no regulation). The first iteration solution had a cost which was roughly 18% above optimal, but by the sixth iteration (22 minutes) the cost was within 1% of optimal and by the seventh iteration (25 minutes) the cost was only 0.35% higher than the IP optimum.

In order to better represent reality and more fully illustrate some of the ADP's capabilities, a stochastic version of the problem was developed, including uncertain vehicle lifetimes and maintenance costs. The expected vehicle lifetime and maintenance costs match those used in the deterministic case. As shown in Figure 24, convergence follows a similar pattern to the deterministic case, though there is no IP optimum available for comparison. The ADP was further run to 350 iterations, yielding insignificant change in the objective function. It was initialized with the naïve guess of all vehicles being worth $80,000.

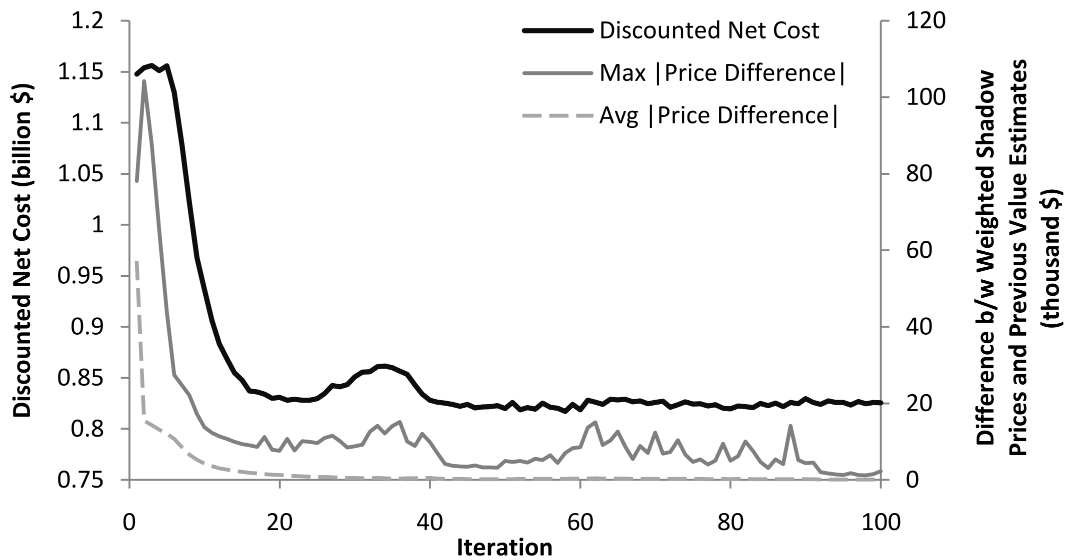

Figure 24. Convergence in a Stochastic Example

The speed at which the objective approaches optimality is an important convergence metric, but it is by no means the only relevant metric. One of the advantages of ADP is that it provides the fleet manager with an estimate of how much each vehicle is worth, from that fleet's perspective. Drinkwater and Hastings (1967) pointed out early on that vehicle values are inherently related to repair limits. In the case where vehicles cannot be resold or scrapped for parts, the repair limit equals the vehicle value. Vehicle values are at the core of the policies recommended by the ADP, but they do not necessarily converge at the same rate as the objective. It is possible, for example, for relatively inaccurate vehicle value estimates to yield nearly optimal behavior, if the fleet manager is lucky.

As with the objective function, it can be difficult to measure the convergence of vehicle values for the full stochastic problem due to a lack of a true value for comparison. For the deterministic steady-state case without retrofits or regulation, however, there is an analytical solution for vehicle values described by expressions (1) and (2). The ADP was set up to analyze this problem, and it was given the deliberately naïve initial guess that all vehicles are worth $80,000, independent of age and condition. The evolution of vehicle value estimates was recorded and the mean absolute percentage error (MAPE) was computed for each iteration. The mean is taken over all vehicles age 15 or younger because older vehicles are correctly valued as scrap from the second iteration onward. As a result, the older vehicle values are all identical and there is no error in their estimates. The results are plotted in Figure 25. The initial guess, which is close to the average vehicle value, yielded a MAPE of roughly 37%. At first, estimates worsened, with the MAPE peaking at just over 64% in the second iteration. The ADP quickly recovered, however, improving to a MAPE just under 1% by the 20th iteration. By the end of the simulation, the MAPE was hovering between 0.1% and 0.3%.

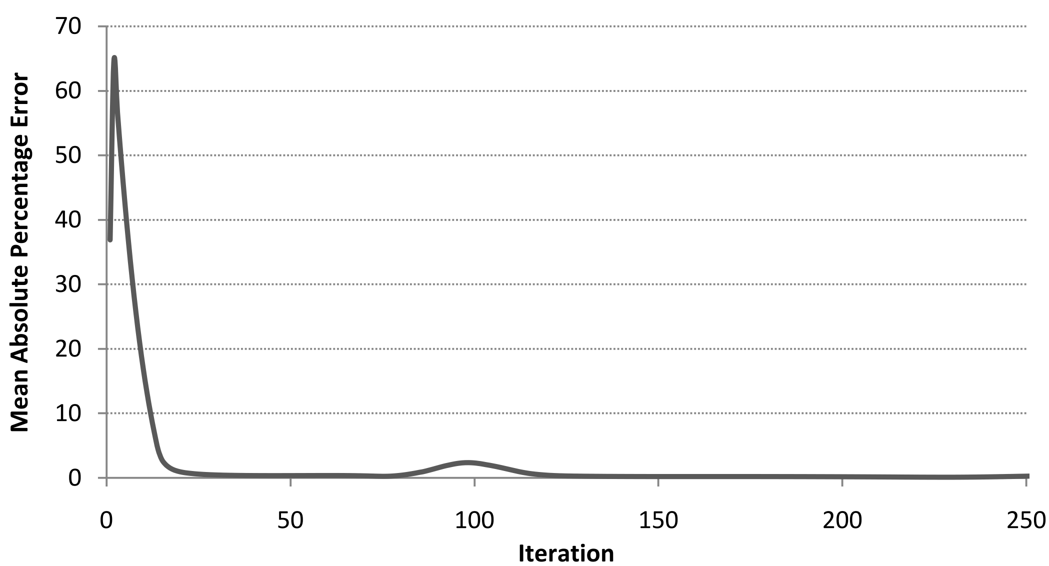

Figure 25. Convergence of Vehicle Value Estimates toward Analytical Solution

The convergence of vehicle values remained relatively consistent as various parameters, such as scrap values, were changed. In one test, the initial value guess was set an order of magnitude too high, at $800,000 per vehicle, well above the $160,000 new vehicle purchase price. Convergence was noticeably slower, but the ADP had clearly managed to head in the right direction despite the very cold start. By the end of 250 iterations (taking 4.36 hours), the R^2 was a respectable 0.839.

Unlike the analytical formula, the ADP can estimate vehicle values in situations which are non-steady-state, or which include the option of retrofitting, or which involve stochastic maintenance costs, or when there is relevant regulation, or all of the above. These vehicle values inform the ADP's dynamic policy recommendations, which come in the form of linear programs to run for each decision period. While these linear programs run very quickly (typically well under a second), the ADP as a whole is slower than the IP described in Chapter 4.

The ADP is likely fast enough to provide recommendations regarding an individual fleet within a few hours of runtime. This should often be sufficient, as it allows for the program to be run several times, using different sets of assumptions, before decisions are made. When a regulator is looking to model a large number of heterogeneous fleets responding to a wide range of regulatory options, it might make sense to use IPs, but that does not mean the ADP does not have a role to play. The IPs require vehicle values at the end of the time horizon, and the ADP can produce better estimates than the analytical solution if the analytical solution's assumptions are violated. The ADP could therefore be used to help set up the IPs, using a limited number of ADP runs.

5.4 Case Study Results

As was discussed in Chapter 3, some aspects of compatibility are quite rigid and easily defined, while others are less so. In the case of NYS DOT, it is relatively clear that FTFs and inexpensive DPFs will not work for most of the fleet. It is debatable whether the most expensive DPFs are compatible, depending on whether the regeneration time is considered an unacceptable burden. This case study will generate optimal policies for the class 8 dump truck fleet under three scenarios. The first is a base case without regulation. The second scenario is based on the assumption that DOCs are considered the best available retrofit technology (BART) at an installation cost of $1,660, while the third scenario assumes that relatively expensive DPFs are considered BART at an installation cost of $16,918.

Vehicle breakdowns and costs are stochastic, with each vehicle being placed in one of three maintenance statuses in each period. They correspond to needing major work ($70,000), moderate work ($7,500) or minimal work ($500). The probability of ending up in a given status depends on vehicle age. Apart from the additional costs to clean filters ($300/year), retrofits are assumed to not impact maintenance costs. Policies, costs, and emissions reductions are computed for both regulation scenarios, as well as the base case without regulation.

In all three scenarios, convergence was relatively fast, after initialization with analytical steady-state vehicle values. The ADP was run to 250 iterations for each scenario, but objective function improvement leveled off within the first dozen iterations. The first 100 iterations are shown in Figure 26.

Beyond convergence, cost fluctuations are largely caused by how “lucky” NYS DOT is with stochastic vehicle failures during a given iteration. Close inspection of Figure 26 reveals that the fluctuations are very similar for the DOC BART, DPF BART, and no regulation scenarios. All were started with the same random seed, so all had the same random numbers determining failures.

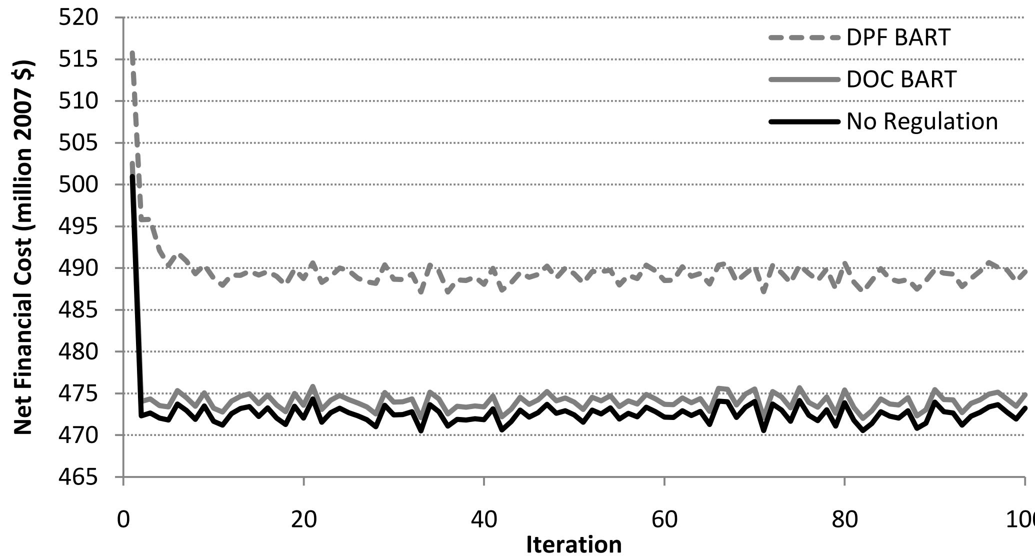

Figure 26. Net Financial Cost Convergence with and without Regulation

Given optimal behavior, it is clear that the regulation is noticeably more costly when DPFs are considered BART, as opposed to DOCs. In the DPF BART scenario, the regulation adds roughly $16.3 million to the cost of operating the fleet. In the DOC BART scenario, the additional cost is roughly $1.6 million. DOCs cost approximately one tenth as much as DPFs, so the ratio makes sense.

Whether DOCs or DPFs are required as BART, the cost imposed by the regulation is significantly less than the cost of retrofitting the entire initial fleet. This cost would be $21.7 million or $2.1 million in the DPF BART and DOC BART scenarios respectively, assuming retrofits are conducted at the last minute before deadlines. The savings is due to both natural and accelerated retirement during the mandate phase-in. As would be expected, retirements are fastest in the DPF BART scenario. In the last period of the regulation phase-in during the final iteration, 21 vehicles are replaced in the DOC BART scenario while 83 are replaced in the DPF

BART scenario. For comparison, 20 vehicles are replaced in this period when there is no regulation.

The impact of the regulation on emissions is presented in Figures 27-30. These figures plot the fleet-wide percent reduction of various emissions in each time period, for both the DPF BART and DOC BART scenarios. PM_{10}, VOC, and CO follow similar patterns. There are no reductions initially, but the reductions spike up as compliance deadlines are reached. Afterward, the reductions slowly decrease. Regardless of the scenario, all vehicles are eventually replaced with vehicles which meet the 2010 standards, causing equivalent emissions in the last periods. The DPF technology is better at reducing PM_{10}, VOC, and CO emissions than DOC technology, but the gap varies by pollutant.

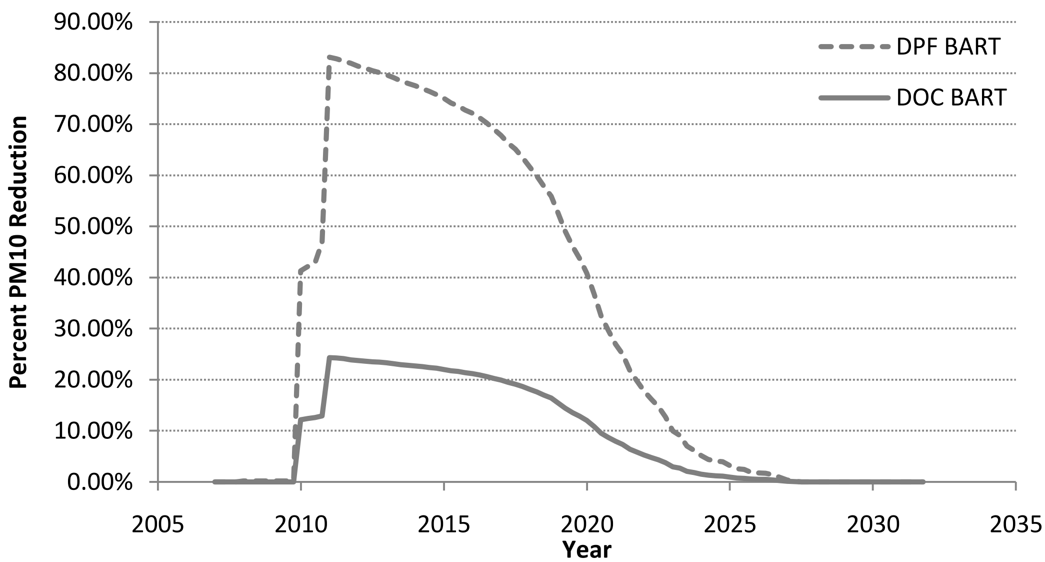

Figure 27. Percent PM_{10} Reduction, Compared to No Regulation Scenario

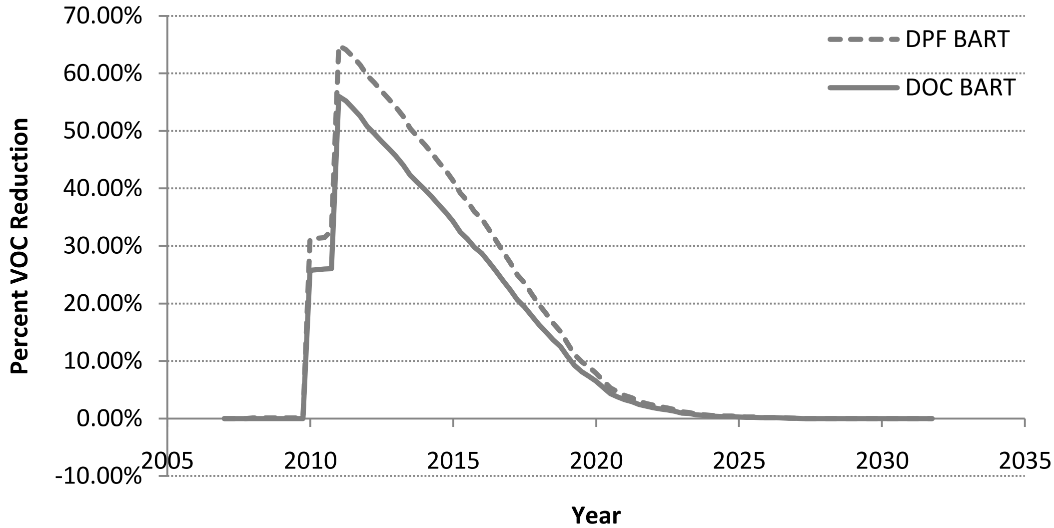

Figure 28. Percent VOC Reduction, Compared to No Regulation Scenario

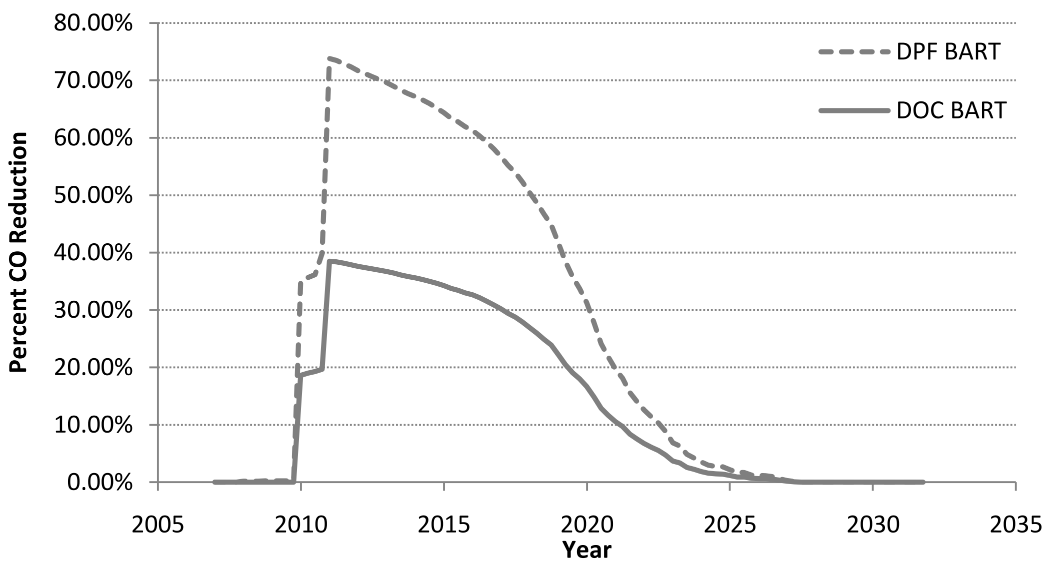

Figure 29. Percent CO Reduction, Compared to No Regulation Scenario

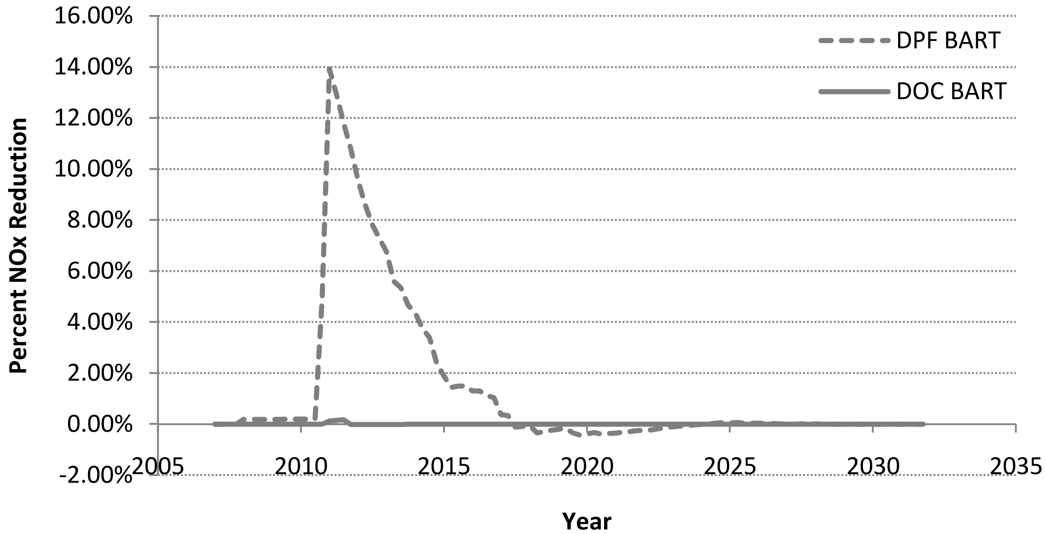

Figure 30. Percent NO_x Reduction, Compared to No Regulation Scenario

Neither DOCs nor DPFs reduce NO_x emission rates, but that does not mean the regulatory mandates do not impact NO_x emissions. In order to avoid paying for retrofits, some vehicles are replaced early, and this causes a drop in NO_x emission rates. The effect is very small for DOCs, as few vehicles are replaced early. For DPFs, a significant number of vehicles are replaced early, resulting in noticeable NO_x emission reductions. The NO_x reductions do not last as long as the reductions of other pollutants. In later years, total NO_x emissions are even slightly higher than without regulation. This is because the vehicles which were bought as early replacements would have had lower NO_x emission rates if they were bought when originally planned.

Given emission reductions and financial costs, the obvious next step is to estimate the net social benefit of the regulation as it was implemented. Emissions were valued in each time period using the same prices as in Chapter 4. Emission values were discounted back to the first period (Q1 2007) using the same discount rate as was applied to financial costs, 1% per quarter. The result was a net social benefit of roughly \$1.33 million in the DPF BART scenario and roughly \$3.7 million in the

DOC BART scenario. NYS DOT ended up applying DOCs nearly exclusively, making the DOC BART scenario a reasonable approximation of reality. It is worth noting that the net social benefit is higher than the IP would have predicted. The IP's reliance on deterministic maintenance costs cause it to underestimate how much older (dirtier) vehicles would be used. As vehicles age, a few outliers noticeably impact the average repair costs, and cause the IP to retire vehicles earlier than the ADP, and earlier than is done in practice (see Figure 19).

The net benefits of regulation are sensitive to changes in the way emissions are valued. If the value of a gram of emissions is deemed to be 7.57% lower than currently assumed, the DPF BART regulation will have zero net benefit. A larger drop will cause a negative net benefit. The positive net benefit associated with the DOC BART regulation is somewhat more robust. A decrease of 69.78% in emissions value would be required to yield a net benefit of zero. On the other hand, if emissions are currently undervalued the DPF BART regulation may actually yield a higher net benefit than the DOC BART regulation. An increase of 19.18% in the value of emissions will cause the DOC BART and DPF BART scenarios to have equivalent net benefits.

The ADP output can be used to measure how the value of the initial fleet changes as a result of regulation. The impact of the regulation on noncompliant vehicle values at the start of the simulation is displayed in Figure 31. For older vehicles which are unlikely to be in use when the regulation takes effect, the impact is minimal. For younger noncompliant vehicles, regulation causes a drop in value under both DOC BART and DPF BART scenarios. Naturally, the drop in value is larger in the DPF BART scenario.

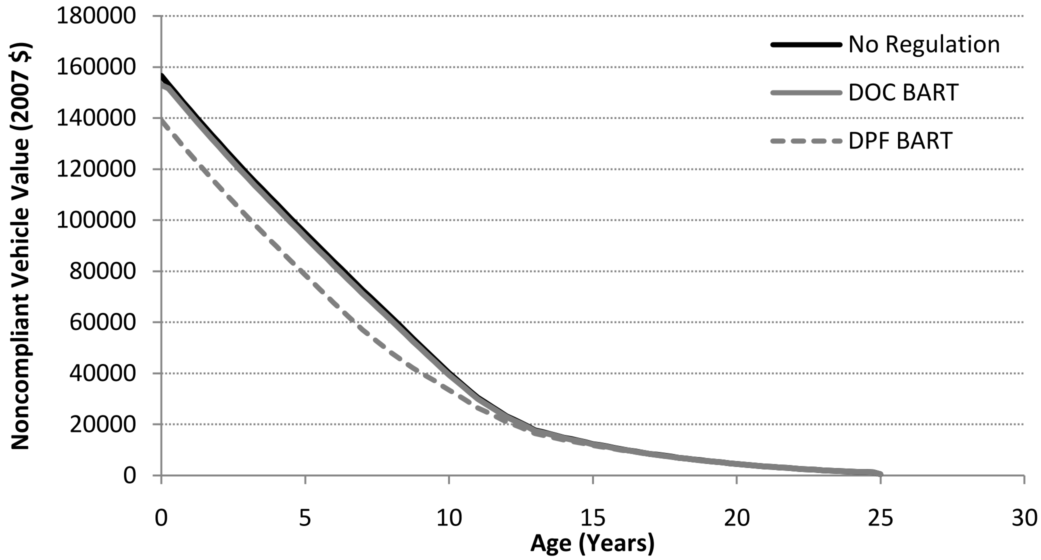

Figure 31. Value of Noncompliant Vehicles in the First Time Period

The vehicle values output by the ADP can reveal how much more the fleet manager would be willing to pay for compliant vehicles, throughout the simulation. Figure 32 plots how much more a vehicle is worth if it has DOC BART installed, as a function of that vehicle's age. At the time when the mandate takes full effect, TMF, the manager is willing to pay the full installation cost of a DOC for young vehicles, but the manager is unwilling to pay this cost for older vehicles which are unlikely to be around long enough to warrant the investment. In earlier time periods, the same trend applies, but the marginal values are lower for two reasons. First, the DOC is not needed yet, so it could be installed later and the cost could be discounted back to the present. Second, the vehicle with the DOC might fail before regulations require the DOC.

The same general pattern applies to the marginal value of a DPF when DPFs are considered BART, as shown in Figure 33. There is one way in which the marginal value of a DPF is different from that of a DOC. Unlike a DOC, a DPF increases the operational cost of the vehicle because of the required filter cleanings ($300 per year).

This impacts the shape of the curves in Figure 33 and means that old vehicles which are unlikely to be around when regulation kicks in are actually less valuable with a DPF installed.

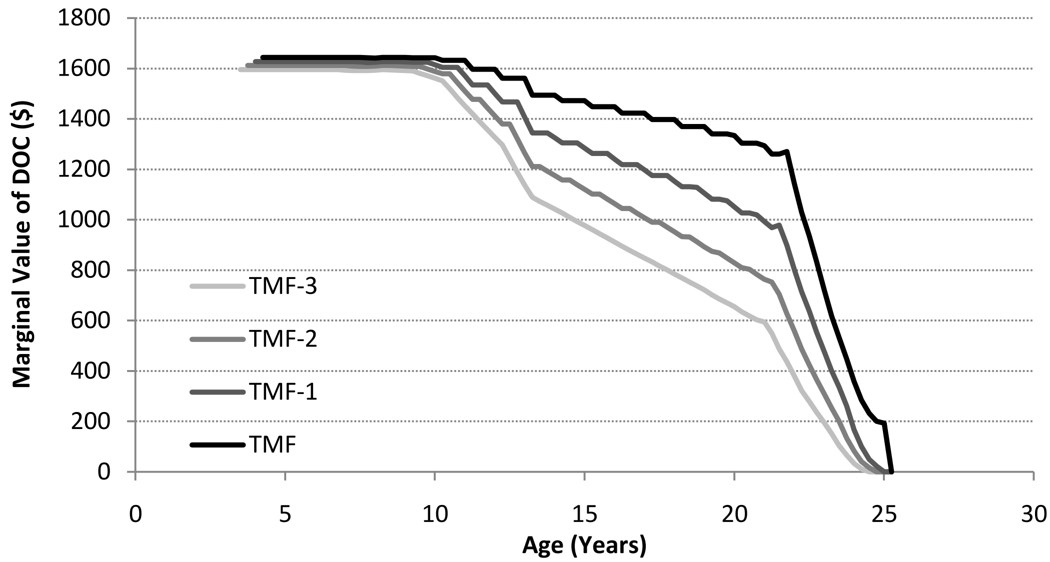

Figure 32. Marginal Value of a DOC on an Additional Vehicle by Vehicle Age, in a Range of Time Periods

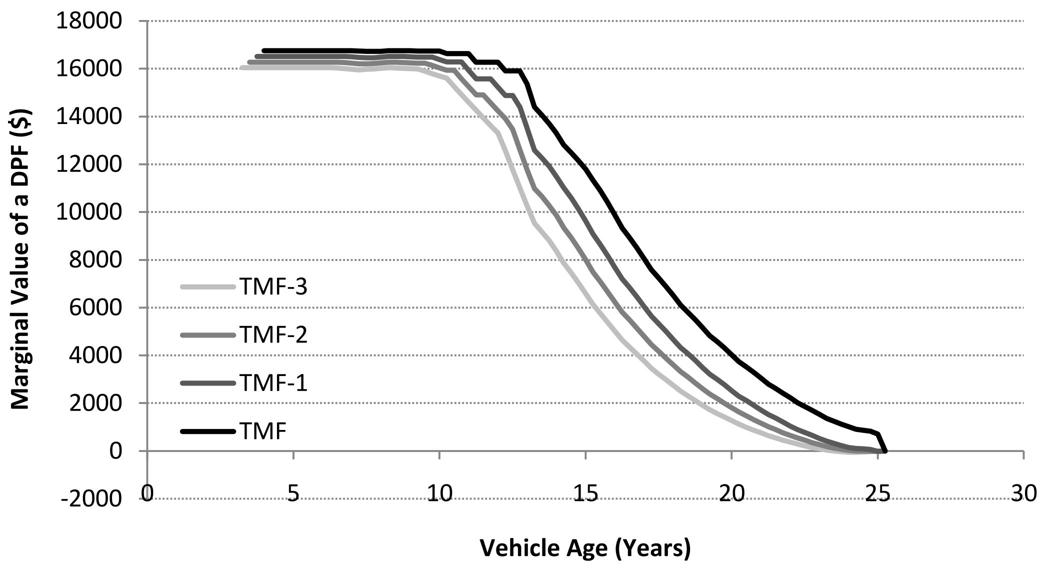

Figure 33. Marginal Value of a DPF on an Additional Vehicle by Vehicle Age, in a Range of Time Periods

In the steady-state, the ADP effectively recommends a policy resembling "repair limit" theory. Vehicles are mostly replaced when they require major repairs, but occasionally older vehicles are replaced when needing more moderate repairs. A single age-dependent repair limit rule clearly cannot be applied while regulations are being phased in, however, as the optimal policy changes in each time period. In these periods, vehicles requiring only minimal repairs may be replaced.

The policies recommended by the ADP are built around the vehicle value estimates, but they are not constrained to a simple formula. The ADP generates a sequence of linear programs, one for each time period, to serve as policies. These linear programs are formulated like that described in Section 5.2 and include the vehicle value estimates. In each time period, the fleet manager updates the appropriate LP by inputting the current state of the fleet, and the LP produces recommendations regarding vehicle purchases, sales, retrofits, and repairs.

CHAPTER 6

CONCLUSIONS

This paper presents a pair of models designed to assist in the management of multiple deteriorating real assets, given financial and environmental concerns. Whether the assets are buildings or vehicles or machines, their purchase and upkeep can be costly, making optimal management policies valuable. The existing literature on this subject is substantial, and the models presented build upon previous work. They incorporate numerous factors which have been modeled previously, though generally not together. These include technological change, linked decisions for multiple assets, and non-steady-state demand. They stand out from previous literature due to their ability to model asset retrofits, as well as repairs and replacements. These retrofits can have initial as well as ongoing costs, and can impact externalities, making them relatively general. For the diesel vehicles in the case study, retrofits cut emissions that cause air pollution, reducing negative externalities.

The first model is an integer program. Its speed allows it to be run many times sequentially, even for fairly large diverse collections of assets. This can be useful from the perspective of a regulator who wants to evaluate the impacts of a wide range of regulatory scenarios on a variety of regulated entities. The primary limitation of this model is that asset failures and repair costs are assumed to be deterministic. In reality, asset failures and repair costs are typically stochastic, and this can influence optimal policies.

The second model, an approximate stochastic dynamic program (ADP), includes both stochastic failures and repair costs. It outputs optimal policies in the form of adjustable LPs, which include parameters that are updated whenever asset failures occur. The ADP also outputs estimates of the value of every asset included in the model. In realistic deterministic example problems, the ADP objective converges toward the IP optimum, and the asset value estimates converge toward analytical solutions. Convergence patterns appear similar for stochastic examples, though true values are not available for comparison.

The downside of the ADP is that it typically converges in a matter of minutes or hours, as opposed to seconds for the IP. This is still fast enough for a single asset manager to evaluate a small group of scenarios, but it may prove too time consuming for a regulator seeking to evaluate hundreds or thousands of potential regulations on a large number of regulated entities. In such situations, the ADP can still be of use in that it can be used to compute asset values which are input parameters for the IP.

Both the IP and the ADP were applied to case studies based on clean diesel regulation in New York State. The IP case study compared a wide range of emissions taxes, emissions rate taxes, and technology mandates, from the perspective of the regulator. As traditional economic theory would predict, emission taxes provided greater net social benefits than regulatory mandates. Emission rate taxes performed surprisingly well, essentially matching the benefits of emissions taxes, despite their inability to impact usage levels. This is a symptom of the facts that 1) the total usage was fixed (assumed to be mandated responsibilities of a government fleet) and 2) the distribution of tasks among vehicles was already nearly optimal because older vehicles tend to be both dirtier and more expensive to operate. Finally, the optimal level of the tax did not necessarily equal the Pigouvian level. This can be explained by the presence of a market distortion resulting from the fleet owner not keeping vehicle resale revenue.

The ADP case study focused on the regulation as it was actually implemented, and the optimal response of NYS DOT. Multiple runs measured the impact of the regulation on NYS DOT's costs, as well as the value of its existing fleet, depending on what retrofits were required to be in compliance. When DPFs were required, the extra costs experienced by NYS DOT were roughly ten times as high as when DOCs were required, though in both cases the cost was far lower than the cost of retrofitting the entire fleet. The ADP was used to measure a net social benefit for the regulation, as applied to the example fleet. The policies produced by the ADP resemble traditional "repair limit" theory in the stochastic steady-state, but when regulation is being phased in the additional flexibility of the ADP structure becomes apparent. It is able to produce dynamic policies which change with the regulation. The ADP is capable of doing the same for changes in demand.

Future research could apply the ADP to other asset types, such as machines or buildings. Future research could also further improve computation speed, allowing the powerful ADP approach to be applied in situations when many model runs are required.

WORKS CITED

Akgül, M. 1984. A note on shadow prices in linear programming. The Journal of the Operational Research Society 35(5) 425-431.

Bellman, R., 1955. Equipment replacement policy. Journal of the Society for Industrial and Applied Mathematics 3(3) 133-136.

Berg, M., Epstein, B., 1978. Comparison of age, block, and failure replacement policies. IEEE Transactions on Reliability R-27(1) 25-29.

Bertsekas, D. P., 1987. *Dynamic Programming: Deterministic and Stochastic Models*. London: Prentice Hall International.

Biswas, S., Hu, S., Verma, V., Herner, J. D., Robertson, W. H., Ayala, A., Sioutas, C., 2008. Physical properties of particulate matter (PM) from late model heavy-duty diesel vehicles operating with advanced PM and NOx emission control technologies. Atmospheric Environment 42(22) 5622-5634.

California Air Resources Board [CARB] (2006). Fact Sheet: Fleet Rule for Public Agencies and Utilities. Available from: <http://www.arb.ca.gov/msprog/publicfleets/publicfleetsfactsheet.pdf>. [Accessed 15 January 2011].

California Air Resources Board [CARB] (2008a). Summary of Adverse Impacts of Diesel Particulate Exposure. Available from: <http://www.arb.ca.gov/research/diesel/diesel_health_effects_summary_7-5-05-1.pdf>. [Accessed January 6 2011].

California Air Resources Board [CARB] (2008b). The Carl Moyer Program Guidelines: Appendices. Available from: <http://www.arb.ca.gov/msprog/moyer/guidelines/cmp_guidelines_part4.pdf>. [Accessed January 6 2011].

California Air Resources Board [CARB] (2008c). The Carl Moyer Program Guidelines: Approved Revision 2008. Available from: <http://www.arb.ca.gov/msprog/moyer/guidelines/cmp_guidelines_part1_2.pdf>. [Accessed January 6 2011].

California Air Resources Board [CARB] (2008d). Staff Report: Initial Statement of Reasons for Proposed Rulemaking: Proposed Regulation for in-use on-road Diesel Vehicles. Available from: <http://www.arb.ca.gov/regact/2008/truckbus08/tbisor.pdf>. [Accessed January 15 2011].

California Air Resources Board [CARB] (2008e). Technical Support Document: Proposed Regulation for in-use on-road Diesel Vehicles. Available from: <http://www.arb.ca.gov/regact/2008/truckbus08/tsd.pdf>. [Accessed January 15 2011].

California Air Resources Board [CARB] (2009a). California Code of Regulations, Title 13, Division 3, Chapter 14. Available from: <http://www.arb.ca.gov/diesel/verdev/reg/procedure_jan1009.pdf>. [Accessed 10 May 2010].

California Air Resources Board [CARB] (2009b). Diesel Certifications: Verification Procedure – Level 3. Available from: <http://www.arb.ca.gov/diesel/verdev/level3/level3.htm>. [Accessed 20 October 2009].

California Air Resources Board [CARB] (2009c). Diesel Certifications: Verification Procedure – Level 2. Available from: <http://www.arb.ca.gov/diesel/verdev/level2/level2.htm>. [Accessed 20 October 2009].

California Air Resources Board [CARB] (2009d). Diesel Certifications: Verification Procedure – Level 1. Available from: <http://www.arb.ca.gov/diesel/verdev/level1/level1.htm>. [Accessed 20 October 2009].

California Air Resources Board [CARB] (2009e). Verification Procedure – Currently Verified. Available from: <http://www.arb.ca.gov/diesel/verdev/vt/cvt.htm>. [Accessed 20 October 2009].

California Air Resources Board [CARB] (2011). Truck and Bus Regulation: On-Road Heavy-Duty Diesel Vehicles (In-Use) Regulation. Available from: <http://www.arb.ca.gov/msprog/onrdiesel/onrdiesel.htm>. [Accessed January 15 2011].

Carr, E.L., Ernst, D.A., Rosenbaum, A., Glass, G., Hartley, S., 2007. Analyzing, Documenting, and Communicating the Impacts of Mobile Source Air Toxic Emissions in the NEPA Process. Available from: <http://onlinepubs.trb.org/onlinepubs/archive/NotesDocs/25-25(18)_FR.pdf>. [1 September 2009].

Chen, Z. 1998. Solution algorithms for the parallel replacement problem under economy of scale. Naval Research Logistics 45(3) 279-295.

Childress, S., Durango-Cohen, P., 2005. On parallel machine replacement problems with general replacement cost functions and stochastic deterioration. Naval Research Logistics 52(5) 409-419.

Cho, D. I., Parlar, M., 1991. A survey of maintenance models for multi-unity systems. European Journal of Operational Research 51(1) 1-23.

Cistemas, M., Blanc, P., Yen, I., Katz, P., E., G., Eisner, M., Shiboski, S., Yelin, E. (2003). A comprehensive study of the direct and indirect costs of adult asthma. Journal of Allergy and Clinical Immunology. 111(6) 1212-1218.

Claggett, Michael and Jeff Houk (2006), Introduction and Background on Mobile Source Air Toxics, Proceedings of the Mobile Source Air Toxics Peer Exchange Meeting. October 5-6, 2006. Monticello, Illinois.

-94.

Clark, N. N., Kern, J. M., Atkinson, C. M., Nine, R. D., 2002. Factors affecting heavy-duty diesel vehicle emissions. Journal of the Air & Waste Management Association 52(1) 84-94.

Clark, N. N., Lyons, D. W., 1999. Class 8 truck emissions testing: Effects of test cycles and data on biodiesel operation. Transactions of the ASAE 42(5) 1211-1219.

Clean Air Task Force, 2005. *Diesel health in America: the lingering threat*. Boston: Clean Air Task Force.

Cleaire (2008). Horizon Owner's Manual. Cleaire Advanced Emissions Controls, San Leandro, CA.

Cleaire (2009a). Vista Owner's Manual. Cleaire Advanced Emissions Controls, San Leandro, CA.

Cummins, (2009a). Clearing the Fog with Crankcase Ventilation. http://www.cumminsfiltration.com/html/en/products/crankcase/crankcase_ventilation.html.

Currie, J., Hanushek, E., Kahn, E. M., Neidel, M., Rivkin, S., 2007. Does Pollution Increase School Absences? NBER Working Paper.

Dekker, R., Wildeman, R. E., van der Duyn Schouten, F. A., 1997. A review of multi-component maintenance models with economic dependence. Mathematical Methods of Operations Research 45(3) 411-435.

Donaldson (2008). Level 3+ Retrofit Emission Solution... The Donaldson SEF Muffler System Delivers the Highest PM Reduction: Using Semi-active DPF Technology. Available from: <http://www.donaldson.com/en/exhaust/support/datalibrary/063725.pdf>. [Accessed 10 May 2010].

Donaldson (2009a). Retrofit Emissions Solutions for Diesel-Powered Vehicles. Available from: <http://www.donaldson.com/en/exhaust/support/datalibrary/002421.pdf>. [Accessed 10 May 2010].

Donaldson (2009b). Installation and Operational Manual: LNF (Low NO2 Filter) Muffler and LXF (Low NOx Filter) Muffler: Manual No. P484796 Rev 2. Donaldson Company, Minneapolis, MN.

Drinkwater, R. W., Hastings, N. A. J., 1967. An economic replacement model. Operational Research Quarterly 18(2) 121-138.

Engine Control Systems (2006). Purifilter: Passively Regenerated Diesel Particulate Filter System: Operation and Maintenance Manual. Engine Control Systems, Thornhill, Ontario.

Engine Control Systems (2008). Purifilter Plus: Hybrid Diesel Particulate Filter Operation and Maintenance Manual. Engine Control Systems, Thornhill, Ontario.

Environmental Solutions Worldwide [ESW] (2009). ThermaCat Active Level III+ Diesel Particulate Filter System. Available from: <http://www.eswcanada.com/ThermaCat_Front_ON.pdf>. [Accessed 10 May 2010].

Fowler, M., Davenport, M., Garg, R., 1992. School functioning of US children with asthma. Pediatrics 90(6) 939-944.

Gao, H. O., Stasko, T. H., 2009a. Diversification in the driveway: mean-variance optimization for greenhouse gas emissions reduction from the next generation of vehicles. Energy Policy 37(12) 5019-5027.

Gao, H. O., Stasko, T. H., 2009a. Cost-minimizing retrofit/replacement strategies for diesel emissions reduction. Transportation Research D 14(2) 111-119.

Gao, H. O., Stasko, T. H., 2010. Diesel Retrofit Assessment for NYS DOT to Retrofit its Existing Engine Fleet: Research Project C-07-12. Available from: <https://www.nysdot.gov/divisions/engineering/technical-services/trans-r-and-d-repository/C-07-12%20Final%20report.pdf>. [Accessed 22 May 2011].

Hastings, N. A. J., 1968. Some notes on dynamic programming and replacement. Operational Research Quarterly 19(4) 453-464.

Herner, J. D., Hu S., Robertson, W. H., Huai, T., Collins, J. F., Dwyer, H., Ayala, A., (2009). Effect of Advanced Aftertreatment for PM and NOx Control on Heavy-Duty Diesel Truck Emissions. Environmental Science and Technology 43(15) 5928-5933.

Hesterberg, T., Bunn, W., McClellan, R., Hart, G., Lapin, C., 2005. Carcinogenicity studies of diesel engine exhausts in laboratory animals: a review of past studies and a discussion of future research needs." Critical Reviews in Toxicology. 35(5) 379-411.

Hill L. Bruce, Zimmerman Neil, Gooch James (2005). A Multi-City Investigation of the Effectiveness of Retrofit Emissions Controls in Reducing Exposures to Particulate Matter in School Buses. Clean Air Task Force.

Holmén, B. A., Ayala, A., 2002. Ultrafine PM Emissions from Natural Gas, Oxidation-Catalyst Diesel, and Particle-Trap Diesel Heavy Duty Transit Buses. Environmental Science and Technology 36(23) 5041-5050.

Hong Y-C., Lee J-T., Kim H., Ha E-H., Schwartz J., Christiani D., 2002. Effects of Air Pollutants on Acute Stroke Mortality. Environmental Health Perspectives 110(2) 187-191.

Howard, R. A., 1960. *Dynamic Programming and Markov Processes*. Cambridge, MA: The MIT Press.

Hsu, C-I., Li, H-C., Liu S-M., Chao, C-C., 2011. Aircraft replacement scheduling: a dynamic programming approach. Transportation Research Part E 47(1) 41-60.

International Agency for Research on Cancer [IARC] (1998). Volume 46: Diesel and Gasoline Engine Exhausts and Some Nitroarenes, Summary of Data Reported and Evaluation. World Health Organization, International Agency for Research on Cancer. Available from: <http://monographs.iarc.fr/ENG/Monographs/vol46/volume46.pdf>. [27 September 2009].

Jin, D., Kite-Powell, H. L., 2000. Optimal fleet utilization and replacement. Transportation Research Part E 36(1) 3-20.

Johnson Matthey (2009). EPA Verifies Johnson Matthey's Level 3 ACCRT™ System – The First Passively Regenerated Particulate Filter for 2.5g NOx EGR and non EGR Engines. Available from: <http://ect.jmcatalysts.com/news2.asp?id=75>. [Accessed 10 May 2010].

Johnson T.V., (2009). Review of diesel emissions and control. International Journal of Engine Research 10(5) 275-284.

Jones, P. C., Zydiak, J. L., Hopp, W. J., 1991. Parallel machine replacement. Naval Research Logistics 38(3) 351-365.

Kansas Green Schools (2009). Kansas Green Schools Grant Program. Available from: <http://www.kansasgreenschools.org/kansas-green-schools-grant-program>. [Accessed 21 March 2009].

Karabakal, N., Lohmann, J. R., Bean, J. C., 1994. Parallel replacement under capital rationing constraints. Management Science 40(3) 305-319.

Kentucky Division for Air Quality [KDAQ] (2008). Request for Proposals for the Kentucky Clean School Bus Grant Program. Available from: <http://www.air.ky.gov/NR/rdonlyres/1AD8AA1A-F170-4343-955B-0BB19C0474B3/0/RequestforProposalsUpdated_RFP2008.pdf>. [Accessed 21 March 2009].

Kittleson, D. B., Abdul-Khalek, I. S., Graskow, B. R., Brear, F., Wie, Q., 1998. Diesel Exhaust Particle Size: Issues and Trends. SAE International Document Number 980525.

Kittleson, D.B., Watts, W.F, Johnson, J.P., Rowntree, C., Payne, M., Goodier, S., Warrens, C., Preston, H., Zink, U., Ortiz, M., Georsmann, C., Twigg, M.V., Walker, A.P., Caldow, R., 2006. On-road evaluation of two Diesel exhaust aftertreatment devices. Journal of Aerosol Science 37(9) 1140-1151.

Kim, Kangwook (2009). The City of New York Department of Sanitation. Email conversation.

Kim, M. J., Makis, V., 2009. Optimal maintenance policy for a multi-state deteriorating system with two types of failures under general repair. Computers & Industrial Engineering 57(1) 298-303.

Kleinberg, J., Tardos, É ., 2006. *Algorithm Design*. Boston: Pearson Addison Wesley.

Legislative Analyst's Office [LAO] (2009). 2009-10 Budget Analysis Series: Transportation: Caltrans—Other Issues. Available from: <http://www.lao.ca.gov/analysis_2009/transportation/trans_anl09004007.aspx>. [Accessed 1 March 2010].

Li, N., Sioutas, C., Cho, A., Schmitz, D., Misra, C., Sempf, J., Wang, M., Oberley, T., Froines, J., Nel, A., 2003. Ultrafine particulate pollutants induce oxidative stress and mitochondrial damage." Environmental Health Perspectives 111(4) 455-460.

Lin, C. 2010. Computing shadow prices/costs of degenerate LP problems with reduced simplex tables. Expert Systems with Applications 37(8) 5848-5855.

Love, C. E., Rodger, A., Blazenko, G., 1982. Repair limit policies for vehicle replacement. INFOR 20(3) 226-236.

Massachusetts Department of Environmental Protection [M DEP] (2008). Diesel Engine Retrofits in the Construction Industry: A How to Guide. Available from: <http://www.mass.gov/dep/air/diesel/conretro.pdf>. [Accessed 1 March 2010].

Moghaddam, K. S., Usher, J. S., 2011. Preventive maintenance and replacement scheduling for repairable and maintainable systems using dynamic programming. Computers & Industrial Engineering 60(4) 654-665.

McClurg, T., Chand, S., 2002. A parallel machine replacement model. Naval Research Logistics 49(3) 275-287.

McConnell, R., Berhane, K., Gilliland, F., London, S., Islam, T., Gauderman J., Avol E., Margolis, H., Peters J., 2002. Asthma in exercising children exposed to ozone: a cohort study. The Lancet, 359(9304) 386-391.

M.J. Bradley & Associates, Inc. & Gruzen Samton LLP (2004). Local Law 77: DDC Ultra-Low Sulfur Diesel Manual. Available from: <http://www.nyc.gov/html/ddc/downloads/pdf/lowsulfur.pdf>. [Accessed 1 March 2010].

Nadeau, K., McDonald-Hyman, C, Noth, E. M.., Pratt B., Hammond S. K., Balmes, J., Tager, I., 2010. Ambient air pollution impairs regulatory T-cell function in asthma. Journal of Allergy and Clinical Immunology 126(4) 845-852.

National Aeronautics and Space Administration [NASA], 2009. How old is the universe? Available from: <http:// map.gsfc.nasa.gov/universe/uni_age.html>. [Accessed 3 April 2010].

Nassau County Government, 2009. Healthy Nassau – Air Quality and Energy Conservation. http://www.nassaucountyny.gov/healthynassau/air/airquality.html

Nebraska Department of Environmental Quality [NDEQ] (2009). Nebraska Clean Diesel Grant Program. Available from: <http://www.deq.state.ne.us/RFP.nsf/c491a3a62d5c1f1686257265005dc889/6ce01deef3310945862575510075b1ed?OpenDocument>. [Accessed 21 March 2009].

Nemmar, A., Hoet, P.H.M., Vanquickenborne, B., Dinsdale, D., Thomeer, M., Hoylaerts, M.F., Vanbilloen, H., Mortelmans, L., Nemery, B., 2002. Passage of inhaled particles into the blood circulation in humans. Circulation: Journal of the American Heart Association, 105(4) 411-414.

Nevada Department of Motor Vehicles [NV DMV] (2010). State of Nevada Department of Motor Vehicles 2010 Auto Body Shop Survey. Available from: <https://dmvapp.nv.gov/DMV/OBL/BodyShop_Survey/Pages/BD_Survey_Results.aspx>. [Accessed 1 March 2010].

New Jersey Department of Environmental Protection [NJ DEP] (2006). Reducing Diesel Emissions from School Buses. New Jersey Department of Environmental Protection, Camptown Bus Lines, Public Service Electric & Gas, and PSEG Power, LLC.

New York State Department of Environmental Conservation [NYS DEC], (2009a). 6 NYCRR Part 248 Regulatory Impact Statement Summary. Available from: <http://www.dec.ny.gov/regulations/56182.html>. [Accessed 10 October 2009].

New York State Department of Environmental Conservation [NYS DEC], (2009b). Adopted Part 248 - Use of Ultra Low Sulfur Diesel Fuel and Best Available Retrofit Technology for Heavy Duty Vehicles, and Part 200 - General Provisions. Available from: <http://www.dec.ny.gov/regulations/56126.html>. [Accessed 10 October 2009].

New York State Department of Environmental Conservation [NYS DEC], (2009c). 6 NYCRR Part 248 Express Terms. Available from: <http://www.dec.ny.gov/regulations/56222.html>. [Accessed 10 October 2009].

New York State Department of Environmental Conservation [NYS DEC], (2011). Notice on the Diesel Emissions Reduction Act of 2006 (DERA) ECL 19-0323/6NYCRR Part 248. Available from: <http://www.dec.ny.gov/chemical/74288.html>. [Accessed 15 January 2010].

New York City Council [NYCC] (2005a). Local Law 2005/039: City's use of ultra low sulfur diesel fuel and the best available technology. Available from: <http://legistar.council.nyc.gov/Legislation.aspx>. [Accessed 10 October 2009].

New York City Council [NYCC] (2005b). Local Law 2005/040: Reducing emission of pollutants from vehicles that handle solid waste and recyclable materials. Available from: <http://legistar.council.nyc.gov/Legislation.aspx>. [Accessed 10 October 2009].

New York City Council [NYCC] (2005c). Local Law 2005/042: Reduce the emissions of pollutants from vehicles that transport children to and from school. Available from: <http://legistar.council.nyc.gov/Legislation.aspx>. [Accessed 10 October 2009].

New York State Energy Research and Development Authority [NYSERDA] (2004a). New York

State Clean Air School Bus Program. Available from: <http://www.nyserda.org/Programs/transportation/AFV/case_studies/CleanAirSchoolBus.pdf>. [Accessed 7 January 2009].

New York State Energy Research and Development Authority [NYSERDA] (2004b). Clean Air School Bus Program. http://www.nyserda.org/Programs/transportation/AFV/cleanairschoolbus.asp>. [Accessed 7 January 2009].

New York State Energy Research and Development Authority [NYSERDA] (2007a). Clean Air School Bus Program: Webster Central School District. Available from: <http://www.nyserda.org/Programs/transportation/AFV/case_studies/webster.pdf>. [Accessed 7 January 2009].

New York State Energy Research and Development Authority [NYSERDA] (2007b). Clean Air School Bus Program: Long Beach City School District. Available from: <http://www.nyserda.org/Programs/transportation/AFV/case_studies/LongBeach2.pdf >. [Accessed 7 January 2009].

Newnan, D. G., Lavelle, J. P., Eschenbach, T. G., 2002. *Essentials of Engineering Economic Analysis*. 2nd ed. New York: Oxford University Press.

Nguyen, D. G., Murthy, D. N. P., 1981. Optimal repair limit replacement policies with imperfect repair. The Journal of the Operational Research Society 32(5) 409-416.

Nunez, Marie (2009). County of Los Angeles. Email conversation.

Nunez, Marie (2010). County of Los Angeles. Email conversation.

Oregon Department of Environmental Quality [ODEQ] (2008). Clean Diesel Retrofit Incentives: Grants and Tax Credits. Available from: <http://www.deq.state.or.us/aq/factsheets/03-AQ-007CleanDiesel.pdf

Oakland Public Works Agency [OPWA], (2007). Public Works Committee Agenda Report: June 26, 2007. Available from: <http://clerkwebsvr1.oaklandnet.com/attachments/16678.pdf>. [Accessed 2 March 2010].

Parry, I. W. H., 1995. Pollution taxes and revenue recycling. Journal of Environmental Economics and Management 29(3) S-64-S-67.

Parry, I. W. H., Bento, A. M., 1999. Tax deductions, environmental policy, and the 'double dividend' hypothesis. Journal of Environmental Economics and Management 39(1) 67-96.

Pham, H., Wang, H., 1996. Imperfect maintenance. European Journal of Operational Research 94(3) 425-438.

Powell, W.B., 2007. *Approximate Dynamic Programming: Solving the Curses of Dimensionality*. Hoboken, NJ: John Wiley & Sons.

Rim, D., Siegel, J., Sprinhirne, J., Webb, A., McDonald-Buller, E., 2008. Characteristics of Cabin Air Quality in School Buses in Central Texas. Atmospheric Environment. 42(26) 6453-6464.

Rockland County Government, 2006. Local Law No. 3 of 2006 – County of Rockland. Available from: <http://www.co.rockland.ny.us/Legislature/Local/past/06/law_03_2006.htm>. [Accessed 8 February 2011].

Roos, C., Terlaky, T. Vial, J.-Ph., 1997. *Theory and Algorithms for Linear Optimization*. Chichester, England: John Wiley & Sons.

Rudell, B., Ledin, M., Hammarström, U., Stjernberg, N., Lundbäck, B., Sandström T., 1996. Effects on Symptoms and Lung Function in Humans Experimentally Exposed to Diesel Exhaust. Occupational and Environmental Medicine 53(10) 658-662.

Rust, J. 1987. Optimal replacement of GMC bus engines: an empirical model of Harold Zurcher. Econometrica 55(5) 999-1033.

Sierksma, G. 1996. *Linear and Integer Programming: Theory and Practice*. New York: Marcel Dekker.

Slaughter, J., Lumley, T., Sheppard, L., Koenig, J., Shapiro, G., 2003. Effects of ambient air pollution on symptom severity and medication use in children with asthma. Annals of Allergy, Asthma, and Immunology 91(4) 346-353.

Stasko, T. H. and Gao, H. O., 2010. Reducing transit fleet emissions through vehicle retrofits, replacements, and usage changes over multiple time periods. Transportation Research D 15(5) 254-262.

Texas Commission on Environmental Quality [TCEQ] (2009a). Emissions Reduction Incentive Grants (ERIG) Application Forms and Information. Online at Available from: <http://www.tceq.state.tx.us/implementation/air/terp/erig_apps.html>. [Accessed 7 March 2010].

Tolbert, P., Mulholland, J., MacIntosh, D., Xu, F., Daniels, D., Devin, O., Carlin, B., Klein, M., Dorley, J., Butler, A., Nordenberg, D., Frumkin, H., Ryan, B., White, M., 2000. Air quality and pediatric emergency room visits for asthma in Atlanta, Georgia. American Journal of Epidemiology 151(8) 798-810.

Tsai, S-S, Goggins, W., Chiu, H-F., Yang, C-Y., 2003. Evidence for an association between air pollution and daily stroke admissions in Kaohsiung, Taiwan. Stroke 34(11) 2612-2616.

U.S. Centers for Disease Control and Prevention [US CDC], (2007). National Surveillance for Asthma --- United States, 1980—2004. Available from: <http://www.cdc.gov/mmwr/preview/mmwrhtml/ss5608a1.htm>. [Accessed 2 October 2009].

U.S. Centers for Disease Control and Prevention [US CDC], (2008). Asthma Fast Facts. Available from: <http://www.cdc.gov/asthma/pdfs/asthma_fast_facts_statistics.pdf >. [Accessed 2 October 2009].

U.S. Centers for Disease Control and Prevention [US CDC], (2009). Asthma: Basic Information. Available from: <http://www.cdc.gov/asthma/faqs.htm>. [Accessed 2 October 2009].

U.S. Department of Health and Human Services [US DHHS] (2005). Diesel Exhaust Particulates: Reasonably Anticipated to be a Human Carcinogen. U.S. Department of Health and Human Services Report on Carcinogens, 11th edition. Available from: <http://ntp.niehs.nih.gov/ntp/roc/eleventh/profiles/s069dies.pdf>. [Accessed 2 October 2009].

U.S. Department of Justice [US DOJ] (1998). "Notice of Filing of Consent Decree Under the Clean Air Act." Federal Register Vol. 63, No. 212, November 3, 1998, pp.59330-59334.

U.S. Environmental Protection Agency [US EPA] (2002a). Generic Verification Protocol for Diesel Exhaust Catalysts, Particulate Filters, and Engine Modification Control Technologies for Highway and Nonroad Use Diesel Engines. Available from: <http://www.epa.gov/etv/pubs/05_vp_devrev.pdf>. [Accessed 20 October 2009].

U.S. Environmental Protection Agency [US EPA] (2003a). Diesel Exhaust in the United States. Available from: <http://www.epa.gov/OMS/retrofit/documents/420f03022.pdf>. [Accessed 4 February 2009].

U.S. Environmental Protection Agency [US EPA] (2006). Diesel Retrofit Technology: An Analysis of the Cost-Effectiveness of Reducing Particulate Matter Emissions from Heavy-Duty Diesel Engines through Retrofits. Available from: <http://www.epa.gov/diesel/documents/420s06002.pdf>. [Accessed 4 February 2009].

U.S. Environmental Protection Agency [US EPA] (2007a). Clean School Bus USA: Retrofit. Available from: <http://www.epa.gov/otaq/schoolbus/retrofit.htm

U.S. Environmental Protection Agency [US EPA] (2007b). Verified Retrofit Technologies from Donaldson. Available from: <http://www.epa.gov/otaq/retrofit/techlist-donaldson.htm>. [Accessed 4 February 2009].

U.S. Environmental Protection Agency [US EPA] (2007c). Verified Retrofit Technologies from Johnson Matthey. Available from: <http://www.epa.gov/otaq/retrofit/techlist-johnmatt.htm>. [Accessed 4 February 2009].

U.S. Environmental Protection Agency [US EPA] (2007d). Diesel Oxidation Catalysts: Informational Update. Available from: <http://www.epa.gov/otaq/diesel/documents/420f07068.htm>. [Accessed 4 February 2009].

U.S. Environmental Protection Agency [US EPA] (2007e). Nitrogen Dioxide Limits from Retrofit Technologies. Available from:

<http://www.epa.gov/otaq/retrofit/documents/420b08005.pdf>. [Accessed 4 February 2009].

U.S. Environmental Protection Agency [US EPA], (2009a). Health Effects Notebook for Hazardous Air Pollutants. Available from: <http://www.epa.gov/ttnatw01/hlthef/hapindex.html>. [Accessed 5 October 2009].

U.S. Environmental Protection Agency [US EPA], (2009b). Nitrogen Dioxide: Health. Available from: <http://www.epa.gov/air/nitrogenoxides/health.html>. [Accessed 5 October 2009].

U.S. Environmental Protection Agency [US EPA], (2009c). IRIS Glossary. Available from: <http://www.epa.gov/ncea/iris/help_gloss.htm>. [Accessed 5 October 2009].

U.S. Environmental Protection Agency [US EPA], (2009d). Tips for a Successful Diesel Retrofit Project. Available from: <http://www.epa.gov/diesel/retrofit-tips.htm>. [Accessed 18 October 2009].

U.S. Environmental Protection Agency [US EPA], (2011a). Carbon Monoxide: Health Available from: <http://www.epa.gov/airquality/carbonmonoxide/health.html>. [Accessed 3 May 2011].

U.S. Environmental Protection Agency [US EPA], (2011b). Health Effects of Ozone in the General Population. Available from: <http://www.epa.gov/o3healthtraining/population.html>. [Accessed 3 May 2011].

U.S. Federal Highway Administration [US FHWA], (2011a). Air Quality: Congestion Mitigation and Air Quality Improvement (CMAQ) Program. Available from: <http://www.fhwa.dot.gov/environment/air_quality/cmaq/>. [Accessed 15 January 2011].

U.S. Federal Highway Administration [US FHWA], (2011b). Air Quality: Congestion Mitigation and Air Quality Improvement (CMAQ) Program: Weighted Populations. Available from: <http://www.fhwa.dot.gov/environment/air_quality/cmaq/populations/>. [Accessed 15 January 2011].

U.S. Environmental Protection Agency [US EPA], (2011c). National Clean Diesel Campaign: Awarded Projects Overview. Available from: <http://epa.gov/cleandiesel/projects/>. [Accessed 15 January 2011].

U.S. Environmental Protection Agency [US EPA], (2011d). National Clean Diesel Campaign: Grants & Funding Overview. Available from: <http://www.epa.gov/otaq/diesel/grantfund.htm >. [Accessed 15 January 2011].

U.S. Environmental Protection Agency [US EPA], (2011e). National Clean Diesel Campaign: Grants & Funding National Funding Assistance Program. Available from: <http://www.epa.gov/otaq/diesel/prgnational.htm>. [Accessed 15 January 2011].

U.S. Environmental Protection Agency [US EPA], (2011f). National Clean Diesel Campaign: Grants & Funding Emerging Technologies Program. Available from: <http://www.epa.gov/otaq/diesel/prgemerg.htm>. [Accessed 15 January 2011].

U.S. Environmental Protection Agency [US EPA], (2011g). National Clean Diesel Campaign: Verification: Emerging Technologies List. Available from: <http://www.epa.gov/cleandiesel/verification/emerg-list.htm>. [Accessed 15 January 2011].

U.S. Environmental Protection Agency [US EPA], (2011h). National Clean Diesel Campaign: Grants & Funding SmartWay Finance Program. Available from: <http://www.epa.gov/otaq/diesel/prgfinance.htm>. [Accessed 15 January 2011].

U.S. Occupational Safety & Health Administration [US OSHA], (2009). Partial List of Chemicals Associated with Diesel Exhaust. Available from: <http://www.osha.gov/SLTC/dieselexhaust/chemical.html>. [Accessed 5 October 2009].

Utah Department of Environmental Quality [UDEQ] (2009). Clean Fuels & Vehicle Technology Grant and Loan Program. http://www.airquality.utah.gov/Planning/Mobile/

Clean_fuels_doc/CFVT_Programoverview.pdf>. [Accessed 2 March 2009].

Van Setten, B. A. A. L., Makkee, M., Moulijn, J. A., 2001. Science and Technology of Catalytic Diesel Particulate Filters. Catalysis Reviews 43(4) 489-564.

Volvo North America (2010). HUSS Engine Retrofit System. Available from: <http://www.volvo.com/constructionequipment/na/en-us/products/innovations/environmental/retrofit/introduction.htm>. [Accessed 3 February 2010].

Wang, H., 2002. A survey of maintenance policies of deteriorating systems. European Journal of Operational Research 139(3) 469-489.

Wargo J., Brown D., 2002. *Children's Exposure to Diesel Exhaust on School Buses*. North Haven, CT: Environment and Human Health, Inc.

Westchester County Government, 2010. Transportation – About Us – Environmental Issues. Available from: <http://transportation.westchestergov.com/index.php?option=com_content&view=article&id=669&Itemid=77>. [Accessed 2 February 2011].

Yatsenko, Y., Hritonenko, N., 2008. Properties of optimal service life under technological change. International Journal of Production Economics 114(1) 230-238.