

ESSAYS IN ENVIRONMENTAL AND DEVELOPMENT ECONOMICS

A Dissertation
submitted to the Faculty of the
Graduate School of Arts and Sciences
of Georgetown University
in partial fulfillment of the requirements for the
degree of
Doctor of Philosophy
in Economics

By

James H. O'Brien, M.A.

Washington, DC
December 19, 2013

Copyright 2014 by James H. O'Brien
All Rights Reserved

ESSAYS IN ENVIRONMENTAL AND DEVELOPMENT ECONOMICS

James H. O'Brien, M.A.

Thesis Advisor: Arik M. Levinson, Ph.D..

ABSTRACT

In this dissertation I examine topics in environmental and development economics. In the first chapter I estimate the value of a statistical life (VSL) for individuals from the age of 18 up to the age of 85 by combining information on vehicle use, household attributes, vehicle prices, crash test results, and yearly fatal accidents for each make, model, and vintage automobile. I calculate a separate willingness to pay for reduced mortality for different age groups and find a significant inverted-U shape to the age-VSL function that ranges from \$1.5 to \$19.2 million (in 2009 dollars). This extends the range of revealed preference estimates of the age-VSL relationship and highlights the importance of considering the specific ages of affected individuals when evaluating public policy.

In the second chapter I present a set of Environmental Engel curves (EECs), which plot the relationship between households' incomes and the amount of pollution embodied in the goods and services they consume. The estimates reveal three clear results. First, EECs are upward sloping; richer households are responsible for more pollution. Second, EECs have income elasticities of less than one; pollution increases with income, but at a decreasing rate. Third, EECs have been shifting over time; at every level of income households are responsible for decreasing amounts of pollution. Using these EECs, I find that 60 percent of the compositional change in household consumption can be attributed to income growth – movement along the

EEC – and 40 percent can be attributed to economy-wide effects such as changing prices due to environmental regulations – shifts in the EEC.

In the third chapter, I examine public works employment programs in India, which offer a guarantee of temporary employment to any household upon request. I estimate the effects of public works participation on the schooling intensity of children in workers' households. To measure this, I rely on a new econometric technique that exploits heteroscedasticity rather than an exclusion restriction. I find that higher public works participation is significantly correlated with lower schooling intensity, and that this effect persists even after controlling for the endogeneity of public works and schooling.

Acknowledgements:

I am grateful to Arik Levinson, Matt Harding, Garance Genicot, Rick Lotspeich, Billy Jack, Mary Evans, Mrittika Shamsuddin, and Liz Schroeder, participants of Camp Resources and SITE Segment 6 workshops, and participants of the Georgetown Economics Department Seminar for helpful suggestions; and to Steve Stafford at the National Automobile Dealers Association for providing batch data from the NADA Used Car Guide.

TABLE OF CONTENTS

Chapter I: Age, Autos, and the Value of a Statistical Life	1
1.1 Introduction.....	1
1.2 Vehicle Choice and the Value of a Statistical Life	4
1.3 Independence of Irrelevant Alternatives	8
1.4 Data	10
1.5 Baseline Results	14
1.6 Correcting for Driver Behavior within Age Groups.....	18
1.7 Caveats	23
1.8 Conclusion	26
Chapter II: Environmental Engel Curves.....	28
2.1 Introduction.....	28
2.2 Movement Along Versus Shifts in the EEC	31
2.3 Data	34
2.4 Estimating Environmental Engel Curves.....	37
2.5 An Application: Decomposing the Effects of Income Growth.....	44
2.6 Conclusion	49
Chapter III: Rural Public Works and Childhood Schooling Intensity	51
3.1 Introduction.....	51
3.2 The National Rural Employment Guarantee Act.....	53

3.3	The Effects of Parent Participation on Child Schooling.....	54
3.4	Econometric Model.....	56
3.5	Estimation.....	60
3.6	Data and Results.....	61
3.7	Conclusion.....	65
Appendix A: Tables and Figures for Chapter I.....		66
Appendix B: Tables and Figures for Chapter II.....		80
Appendix C: Tables for Chapter III.....		93
Appendix D: Data Appendix for Chapter II.....		98
D.1	Introduction.....	98
D.2	Data Consolidation.....	99
D.3	Manually Calculate Total Pollution Intensities.....	106
D.4	Calculate Sub-IO-Level Pollution Intensities and Expenditure.....	107
D.5	Aggregate from Sub-IO to CEX.....	111
D.6	Combine Aggregate Pollution Intensities with NBER CEX Extracts.....	112
Bibliography.....		114

CHAPTER I

AGE, AUTOS, AND THE VALUE OF A STATISTICAL LIFE

1.1 Introduction

The value of a statistical life (VSL) measures people's demonstrated willingness to pay for reduced mortality risk and is used to assign a dollar value to the benefits of regulations that protect health and safety. For many such regulations, those benefits accrue disproportionately to older people. But most current VSL estimates come from hedonic wage studies based on samples that include few older workers and no retirees.¹ Rather than examining the risk-income tradeoff that individuals face in their labor market decisions, I focus on the risk-income tradeoff that individuals face in their automobile purchase decisions. By comparing the relative importance of automobile cost and safety for various age groups, I extend the age range of individuals that can be captured in choice-based estimates of the VSL.

Policymakers recognize that health and safety benefits often accrue disproportionately to older groups, but this isn't currently incorporated into regulatory cost-benefit analyses in the US. Other countries, such as Canada and the European Commission, have at times recommended the use of an age-adjusted VSL for seniors (Aldy and Viscusi, 2007). In 2003 the US Environmental Protection Agency (EPA) released a cost-benefit analysis applying a "senior discount" of 37 percent to the VSL of people over the age of 70 (EPA, 2002). The political backlash that followed led Congress to prohibit the use of age adjustments in Federal cost-benefit analyses.² Currently, Federal agencies are instructed not to use age-adjustment factors due to "continuing

¹ For example, the EPA uses a VSL estimate derived from a meta-analysis of 26 studies, including 21 hedonic wage studies (US EPA, 2010b).

² See, for example, Viscusi (2011). See also Evans and Smith (2006) and Scotton and Taylor (2011) for a discussion of several House and Senate bills limiting adjustments to the VSL.

questions over the effect of age on VSL estimates” (US OMB, 2003). Similarly, the OECD also recommends “no adjustment for adults due to inconclusive evidence” (OECD, 2012).

Policymakers’ mistrust of an age-adjusted VSL is well placed. In theory, the age-VSL relationship can take any shape (Johansson, 2002; Aldy and Viscusi, 2003). Empirical work on this topic has also produced mixed results. Table A.1 and Figure A.1 summarize several prominent empirical age-VSL studies and their findings.³ These can be separated into two types: those that rely on revealed preferences and those that rely on contingent valuation. Among the revealed preference studies, the vast majority employ a hedonic wage framework to estimate a wage-risk tradeoff of workers across different industries, occupations, and demographics.⁴ Although there is considerable variation, hedonic wage studies generally support a flat or inverted-U-shaped age-VSL relationship, such as in Aldy and Viscusi (2003 and 2008).⁵ The survey-based contingent valuation approach is not limited by labor market participation and captures a more diverse set of individuals. Stated-preference studies typically show flat, slightly decreasing, or weakly inverted-U shaped age-VSL estimates, such as in Krupnick et al. (2002), Alberini et al. (2004) and DeShazo and Cameron (2004).⁶

³ See Viscusi and Aldy (2003), Aldy and Viscusi (2007), Evans and Smith (2006), and Krupnick (2007) for additional review of the literature surrounding age and VSL.

⁴ A handful of non-labor market studies have also obtained mixed results, such as in Portney (1981) (housing prices and local air pollution), who finds a substantially negative age-VSL relationship, Andersson (2008) (hedonic vehicle prices), who finds no age effect, Mount et al. (2004) (hedonic vehicle prices) who find different VSLs for families with seniors, adults, and children, and Rohlfs et al. (2012) (hedonic vehicle prices and air bags) who find no effect.

⁵ Other studies have focused on revealed preferences without relying on labor market data, but do not address the age-VSL relationship. Examples include Ashenfelter and Greenstone (2002) (mandated speed limits); Atkinson and Halvorsen (1990) and Dreyfus and Viscusi (1995) (hedonic vehicle prices); Blomquist et al. (1996) (seatbelt, helmet, and child seat use); Schnier et al. (2009) (commercial fishing decisions), and León and Miguel (2013) (estuary crossings in Sierra Leone).

⁶ See for example Arrow, et al. (1993), Portney (1994), Diamond and Hausman (1994), and Hausman (2012) for a discussion of the advantages and disadvantages of contingent valuation.

In this paper I use a revealed-preference approach to estimate the VSL over an age range previously captured only by stated preferences. A key challenge I face in estimating the VSL based on automobile choice is that the number of fatalities in each vehicle depends on the individuals who drive those vehicles. For example, if Prius drivers and Camaro drivers differ from one another in unobservable ways, estimates based on the relative riskiness of those vehicles could be biased. I control for the effects of drivers' behavior on vehicles' measured safety in two ways. First I use age-specific risk factors for evaluating each individual's purchase decision. Second, I directly estimate driver behavior based on differences between laboratory crash test results and real-world single-vehicle fatalities.

Although estimating a VSL using automobile choice is challenging because measured risk depends on driver behavior, this approach also has several distinct advantages. The first is that automobiles are purchased by adults of all ages. For example, in the 2009 National Household Transportation Survey (NHTS) roughly 11 percent of automobile acquisitions involved someone over the age of 65. Second, this study provides VSL estimates in a context that is more appropriate for transportation-related regulations than the currently-used estimates based on occupational risk. Third, by measuring the relative importance of fatality risk and cost in the consumer's underlying utility function, I separately observe the effects of each as they flow into the final VSL estimate. This will be the first study to identify whether age differences in the VSL are the result of different preferences for risk, cost, or both.

I find the age-VSL function over the lifecycle follows an inverted-U shape. Moreover, this inverted-U is the result of two distinct patterns in the marginal disutility of cost and risk. Seniors and middle-aged drivers have similar disutility from cost, but seniors place less

importance on safety, which results in a lower VSL for seniors relative to middle-aged drivers. Young drivers place less importance on safety and more importance on cost relative to middle-aged drivers, which implies a lower VSL for young drivers relative to middle-aged drivers. Overall, the VSL ranges from \$2.4 million for the youngest group (18 to 24 years) to \$19.2 million for the middle-aged group (55 to 64 years), falling again to \$8.2 million for the oldest group (75 to 85 years). These results are comparable in scale to existing wage-based estimates, but represent the first examination of revealed-preference VSL for the 18-85 age range.

1.2 Vehicle Choice and the Value of a Statistical Life

To calculate average willingness to pay for reduced fatality risk I start with a basic discrete choice framework. When facing the decision to purchase an automobile, each consumer gains utility from various vehicle attributes and chooses the car that gives her the highest total utility. These attributes can include safety, cost, and other features such as spaciousness, performance, origin, and luxury. By estimating the relative importance of safety and cost for different age groups, I calculate the average VSL for each group as the trade-off between fatality risk and cost that keeps that group at a constant utility.

To begin, suppose individual n obtains utility from a given vehicle choice j in the set of alternatives J . I model individual n 's utility as a linear function of various vehicle attributes and a mean-zero error term:

$$U_j^n = \beta_1 risk_j + \beta_2 cost_j + X_j \gamma + \varepsilon_j^n \quad (1)$$

where $risk_j$ measures the perceived annual probability of a fatal accident in vehicle j , $cost_j$ is a combination of a discounted purchase price and an annual operating costs, including fuel, insurance, and taxes, and X_j captures other vehicle attributes that contribute to utility.⁷

If the errors are distributed according to an extreme value distribution then McFadden (1974) shows that the probability of a consumer choosing a particular automobile takes the form:

$$P_j^n = \frac{e^{V_j^n}}{\sum_{i=1}^J e^{V_i^n}} \quad (2)$$

where V_j^n includes the non-error terms from equation 1, i.e. $V_j^n = \beta_1 risk_j + \beta_2 cost_j + X_j \gamma$.

Given this probability, the coefficients in equation 1 can be estimated using standard maximum likelihood techniques.

The VSL is likely to vary across several individual characteristics; income and age being the two most salient. But multinomial logit estimation is based on differences in utility across various alternatives within a given individual's choice set, so I can only observe the effects of individual-specific characteristics insofar as those characteristics interact with vehicle attributes.⁸ To account for the possibility that higher income households are less sensitive to vehicle costs, I incorporate income by expressing cost as a fraction of income for each alternative. In addition, to see how different groups may value riskiness and cost, I interact the risk and cost variables with an age-group indicator. In this case, the utility function takes the form:

$$U_j^n = \sum_c \beta_{1c} risk_j * I_c(age_n) + \sum_c \beta_{2c} \frac{cost_j}{income_n} * I_c(age_n) + X_j^n \gamma + \varepsilon_j^n \quad (3)$$

⁷ A linear utility function implies that consumers are risk-neutral over all vehicle attributes, including safety. In this context, I assume that consumers face a constant marginal disutility from fatality risk.

⁸ See, for example, Train (2009) for further discussion.

where $I_c(age_n)$ is an indicator function that equals one when person n is a member of a given age cohort, c . With the utility function specified in this manner, I estimate a separate risk and cost coefficient for each age group.

The VSL represents individuals' willingness to pay for reduced fatality risk. This can be calculated as the reduction in annual cost that leaves a consumer indifferent to an increase of one fatality per year. By totally differentiating the utility function in equation 3 and setting $dU=0$, I solve for the marginal rate of substitution between cost and fatality risk. The VSL in this context is the negative average marginal rate of substitution for each group:

$$VSL^c = -\frac{\partial cost}{\partial risk^c} = \frac{\beta_{1c}}{\beta_{2c}} * \overline{income}^c \quad (4)$$

where \overline{income}^c is the average income for households in cohort c and corrects for the fact that vehicle cost enters the utility function as a share of income. Since I estimate separate risk and cost coefficients for each age group in equation 3, I also estimate equation 4 separately for each group.

Automobiles are durable products and a driver may use the same automobile for many years. I assume that consumers have the continuous option to re-purchase a car that optimally reflects their preferences for safety and cost. In this way, the coefficients in equations 3 reflect consumers' preferences at the time of purchase and are unaffected by past or future preferences. For example, when a 55-year-old driver purchases an automobile, I assume she doesn't consider her expected preferences at age 65 because she will be able to change her car if her preferences change. This is necessary to interpret each group's VSL estimates from equation 4 as a cohort-specific value that is unaffected by the estimates for other age groups.

Further, in order for the marginal rate of substitution in equation 4 to have an intuitive interpretation as a VSL, the units on the right-hand side must combine to yield a value of dollars per statistical life. To accomplish this, I measure the riskiness of an automobile as expected deaths per year and the cost of an automobile as dollars per year, so that the ratio of the two coefficients on those variables can be measured in dollars per death.

Measuring automobile safety in terms of expected deaths per year is necessary in order to obtain a meaningful value in equation 4, but it also introduces an empirical challenge. Using the number of fatalities to proxy for individuals' perceptions of safety ignores the fact that fatality rates in a given automobile depend on the behavior of the people who choose to drive that car.⁹ The number of deaths may not be a good approximation of individuals' perceptions of safety if safe and unsafe drivers systematically opt for different vehicles. For example, Figure A.2 shows the number of deaths per mile for the ten safest and riskiest vehicles in my sample. The car with the fewest deaths per mile is the Toyota Prius whereas the car with the most deaths per mile is the Chevrolet Camaro. Part of this difference may be due to the fact that Prius drivers may be safer than Camaro drivers, having nothing to do with the physical cars themselves. In this case, the number of deaths per mile in a Camaro may not accurately represent an individual's perceptions about her own safety if she believes that she is a safer driver than the *typical* Camaro owner.

I account for the effects of drivers' behavior on vehicles' measured riskiness in two ways. First I use an age-specific measure of fatality risk and later I add an additional correction derived using vehicle crash test results. Because age is a particularly salient dimension for individual

⁹ Statistical deaths per mile is commonly used to proxy for vehicle fatality risk (see for example Dreyfus and Viscusi (1995) and Atkinson and Halvorsen (1990)).

sorting, I start by calculating fatalities per mile for each make, model, and vintage automobile separately for each age cohort. Measuring risk in this manner accounts for differing vehicles and driving patterns across age groups. For example, I assume the risk to seniors from driving a Camaro is captured by the number of deaths per mile suffered only by other seniors who also drove a Camaro; the number of teenager fatalities in a Camaro does not enter the calculation.¹⁰

In this context the model for underlying utility becomes:

$$U_j^n = \sum_c \beta_{1c} risk_j^c * I_c(age_n) + \sum_c \beta_{2c} \frac{cost_j}{income_n} * I_c(age_n) + X_j^n \gamma + \varepsilon_j^n \quad (5)$$

where the only difference between equations 4 and 3 is that equation 4 includes a cohort-specific measure of fatality risk for each alternative (indexed by c).

In addition to age-specific fatality risk, I also account for drivers' effects on measured safety by using vehicle crash test results to estimate relative driver behavior across vehicle types. By comparing vehicles with similar crash test performance, but different numbers of single-vehicle fatalities, I estimate the relative propensity to experience a fatality based on driver characteristics for different types of vehicles. Then I use these estimates of relative driver behavior to reweight the risk measure for each vehicle within each individual's choice set to more accurately reflect the riskiness of each vehicle for a given driver.

1.3 Independence of Irrelevant Alternatives

¹⁰ I could add additional variables such as gender and education to reduce unobserved heterogeneity when calculating fatalities per mile. The underlying assumption would be that individuals perceive the number of deaths per mile in each make, model, and vintage automobile that are suffered exclusively by people of their same age-gender-education cohort. The number of fatal accidents per group decreases substantially with additional covariates, making it impractical to group individuals into narrowly defined cohorts.

As is well known, assuming independent extreme value errors leads to the property of independence of irrelevant alternatives (IIA). This means that the probability of selecting a given vehicle relative to a specific alternative depends only on the characteristics of those two vehicles and not on any characteristics of any other vehicle. A first step towards mitigating this limitation is to add an alternative-specific constant to the utility function.¹¹ In this context, the ideal would be to include a make/model/vintage-specific dummy in the utility specification of equations 4 and 6. But since I assume that consumers could have purchased any other make, model, and vintage combination that was purchased in the same calendar year, there are roughly 1000 alternatives in each individual's choice set.¹² Rather than including a unique constant for each alternative, I take a middle ground by including a set of constants that identify vehicles according to seven different body types (2-door, 4-door, sport, minivan, pickup, small SUV, large SUV).

The large number of alternatives in each choice set also complicates alternative estimation techniques with more flexible IIA assumptions. Models such as the multinomial probit and the mixed logit are not subject to IIA, but are computationally intractable with the roughly 1000 alternatives available in this case.¹³ In fact, the IIA assumption presents a considerable advantage for estimation; since omitted alternatives are assumed irrelevant, coefficients in equation 4 can be estimated using a subset of alternatives for each individual (McFadden, 1978 and Train, 2009). I estimate the vehicle choice model using 10 random alternatives drawn without replacement from the set of all vehicles purchased in the same

¹¹ See Train (1993 and 2009) for discussion of alternative-specific constants and IIA.

¹² Depending on the specific variables included in the regression, the choice sets for each individual range in size from 670 to 1061 alternatives. For 2008 (January through April) the choice sets range from 120 to 200 alternatives.

¹³ See Train (2009) for a detailed discussion of various multinomial choice estimation techniques.

calendar year for each consumer.¹⁴ This decreases the number of vehicles from over 1,000 alternatives for each individual to only 10.¹⁵

1.4 Data

The data for this analysis come from five different sources: The National Household Transportation Survey, which provides information on household characteristics, vehicle holdings, and miles travelled; the Fatality Analysis Reporting System, which provides information on the number of deaths that occur in each specific vehicle model and vintage, along with some vehicle attributes; a dataset of vehicle attributes and prices provided by the National Automobile Dealers Association; additional vehicle attribute data from the EPA Fuel Economy Guides; and crash test results from the NHTSA New Car Assessment Program (NCAP). All five sources are combined to show individual vehicle purchase decisions, attributes of various alternative vehicles, and the number of deaths per mile of each age cohort in each make, model, and vintage combination.

The base unit of observation for my analysis is an individual household-automobile pair identified in the National Household Transportation Survey (NHTS), a nationally representative inventory of daily travel and vehicle holdings in the United States.¹⁶ I rely on the most recent round conducted between March 2008 and April 2009 with information on approximately 150,000 households and their vehicles. Rather than specific automobile sales transactions, I rely

¹⁴ The results of this analysis are robust to using five or 15 alternatives.

¹⁵ As a final note, it is not clear in practice whether IIA introduces a significant bias when the focus is on estimating parameters of the underlying utility model rather than generating substitution patterns, as is the case here. Recent studies have found that the multinomial logit (subject to IIA) performs as well or better than a multinomial probit or mixed logit (not subject to IIA) in terms of parameter estimation (Kropko, 2010 and Dahlberg and Eklöf, 2003).

¹⁶ Along with its predecessor survey, the NHTS has been conducted roughly every five to eight years since 1969.

on vehicle holdings in the NHTS to back-out information about vehicle acquisitions; for each household-automobile pair in the survey, I calculate the purchase date and the age of the primary driver at the time of purchase based on the date of the survey and number of months that the vehicle has been owned.¹⁷

I start by excluding all vehicles that were not registered for personal use (such as commercial and fleet vehicles) as well as medium- and heavy-duty trucks, regardless of their indicated use. I drop vehicles with extremely high annual mileage and rare vehicles selling for extremely high prices (the top one percent in terms of mileage and price).¹⁸ I also focus on automobiles that were acquired between January 2007 and April 2009, the three most recent calendar-years available.

The NHTS includes descriptive information about each household, including income bracket, intensity of automobile use (annual vehicle miles travelled), and annual fuel costs for each vehicle. Since past income is not available from the NHTS, I assume that real income at the time of purchase is equal to income reported in the survey.¹⁹ I also use this information to calculate an average operating cost for each make, model, and vintage combination.

Next, I calculate a per-mile fatality risk for each automobile. I use the National Highway Traffic Safety Administration's Fatality Analysis Reporting System (FARS), a census of all fatal crashes that occur in the United States each year, to calculate the number of fatalities by age

¹⁷ Because I don't have information on specifically who purchased the vehicle, I assume that the primary driver purchased the vehicle.

¹⁸ This results in dropping 330 vehicles, each with mileage greater than 60,349 miles per year and 205 vehicles that had used values greater than \$43,107.

¹⁹ I focus on recent purchases, limiting the sample to vehicles acquired between 2007 and 2009. Results are not significantly different if instead I adjust each household's income based on the average growth of real income from the US Census.

group for each make, model, and vintage automobile.²⁰ Then I use the NHTS to find the total number of miles traveled by each age group for each vehicle. I divide the number of deaths from the FARS by the number of miles driven from the NHTS to calculate the number of deaths per mile for each vehicle and age group. There are 1264 unique automobiles that appear in both the 2009 FARS and the 2009 NHTS and were purchased between 2007 and 2009, corresponding to 14,327 household-vehicle pairs. These make, model, and vintage combinations were associated with 9840 in-vehicle fatalities and had an average of 7.8 deaths and 741 million miles traveled.

To find automobile prices, I match each of the 14,327 household-vehicle pairs to the National Automobile Dealers Association (NADA) Used Car Guide. Based on the VIN numbers listed in the FARS and purchase year from the NHTS, I find year-specific prices for each observation.²¹ I rely on the estimated retail value for a used car in clean condition, adjusted for optional features based on VIN numbers. This price captures a typical price that would be paid for each car and is estimated by NADA based on “many sources of information, including wholesale and resale transaction data.”²² I assign a used price to all cars with the assumption that any premium paid for a new car compared to an exactly identical used car is unrelated to the (identical) safety features. Further, the used vehicle price represents the opportunity cost of holding a particular vehicle instead of re-purchasing a vehicle with different attributes. Table A.2 shows summary statistics for the sample of 14,327 household-vehicle pairs.

The VSL calculated in equation 5 is based on the ratio of the risk and cost coefficients, so it is necessary for these two variables to have intuitively comparable units. To achieve this I

²⁰ The FARS was formerly referred to as the Fatal Accident Reporting System.

²¹ The vehicle identification number (VIN) is standardized by the National Highway Traffic Safety Administration. The first 10 digits contain make, model, characteristics, and model year information.

²² See <http://www.nada.com/b2b/Support/Glossary.aspx>

calculate a per-year cost and per-year risk for each automobile based on households' vehicle use. Because each household drives a different number of annual miles, the per-year risk and cost for a given make, model, and vintage automobile will be specific to each household.

To annualize automobile cost I take the sum of two components, an annualized purchase price and an annual operating expense. The cost of vehicle j for household n is:

$$cost_j^n = (\delta + 0.15) * price^j + opexp_j * vmt^n \quad (6)$$

where δ is an annual discount rate, $price^j$ is the NADA used car price for automobile j , $opexp_j$ is the average per-mile operating expense for vehicle j based on the NHTS, and vmt^n is the number of annual miles driven by household n . I use a five percent discount rate, although the results are robust to using three or seven percent instead.²³ I also follow Atkinson and Halvorsen (1995) by including 15 of the purchase price to capture insurance and taxes.

To convert the per-mile risk measure to a per-year estimate, I multiply by each household's annual miles traveled. The riskiness of vehicle j for household n is:

$$risk_j^n = fatalities_j * vmt^n \quad (7)$$

where $fatalities_j$ is the fatalities per mile in automobile j and vmt^n is the number of annual miles driven by household n .

Annualizing the risk and cost metrics based on household miles traveled is necessary to obtain comparable units on the risk and cost variables, but in doing so I implicitly assume that drivers do not adjust their mileage based on the safety or cost of their cars. For both variables, this assumption may not be correct. For example, drivers may use small and less safe cars for local trips and larger safer cars for highway trips. Likewise for annual cost, drivers may use fuel

²³ The various discount rates lead to slight and insignificant shifts in scale for the age-VSL function.

efficient vehicles for commuting and use gas guzzlers for short trips that require a large vehicle. If in reality households tend to reduce their mileage in risky vehicles and also tend to reduce their mileage in costly vehicles, the risk and cost coefficients in equation 4 are biased in the same direction. Since the VSL is the ratio of the two coefficients, it is not possible to say how this will bias the final VSL estimate, except that it is likely less severe than the bias of either coefficient individually.

1.5 Baseline Results

I begin by ignoring the fact that the number of fatalities experienced in a particular automobile is a function of the type of people who choose to drive that car. I measure the riskiness of each vehicle based on the total number of fatalities (experienced by any age group) divided by the total number of miles (driven by any age group) and calculate an annual risk for each household by multiplying by annual miles traveled. I estimate the multinomial choice model based on pooled risk described in equation 3. To capture age variation in the risk and cost coefficients, I split the sample into seven age cohorts ranging from 18 to 85 and interact an indicator for each cohort with the risk and cost variables.

The first column of Table A.3 shows the results of estimating equation 3 using a pooled risk measure. Coefficients on the non-risk, non-cost vehicle attributes suggest a generally intuitive effect of these characteristics on consumer utility. Car buyers appear to gain utility from more engine displacement (0.16), higher MPG (0.33), more passenger space (0.14), domestic vehicles (-0.34 for imports), and newer vehicles (0.36 and 1.19 for recent model

years).²⁴ I also include an indicator for luxury brands and find that the effects depend on income; lower income consumers (below the 75th percentile) experience a negative marginal utility while higher income consumers (above the 75th percentile) experience positive marginal utility.²⁵

By interacting the risk variable with a cohort indicator, I calculate a separate risk and cost coefficient for each age group (even though the risk variable is the same for all groups). As expected, the cost and risk coefficients in the first two columns of Table A.3 are all negative, suggesting disutility from higher cost and higher fatality risk. These coefficients are also all significant, most at the one percent level.²⁶ With the exception of the 45-54 age group, the magnitude of the risk coefficient increases with age from -5.42 (in column 2) to -32.28 for the 65-74 age group, with a slight dip to -29.51 for the oldest group. The cost coefficients on the other hand are generally higher in magnitude for younger groups (e.g. -8.34 and -9.90) than for older groups (-3.40 and -3.06).

Because I haven't controlled for unobservable driver behavior, the risk coefficients in column 1 of Table A.3 are potentially biased. Nevertheless as a point of comparison I calculate the VSL for each group based on the ratio of the risk and cost coefficients multiplied by the average income for each group (equation 5). These values are shown at the bottom of Table A.3 and range from \$3.8 for the youngest group to \$130.4 million for the 55-64 age group, falling to \$61.2 million for the oldest age group.

²⁴ Engine displacement measures the volume swept by the pistons within the engine cylinders. In this case I use displacement as a measure of vehicle performance.

²⁵ The following brands are considered luxury: Acura, Audi, BMW, Cadillac, Infiniti, Jaguar, Land Rover, Lexus, Mercedes-Benz, and Porsche.

²⁶ Standard errors are calculated based on 1000 bootstrap repetitions. Because of the non-linear nature of the VSL (a ratio), outliers are trimmed from the bootstrap calculation.

The first way that I correct for the influence of drivers' behavior on vehicles' measured risk is to calculate fatality risk separately for each age cohort. The results of the automobile choice model using these age-specific risk factors are shown in the second column of Table A.3. The effects of non-risk vehicle characteristics on consumer utility are similar when I estimate the multinomial choice model using age-specific risk factors. The cost coefficients are larger in column 2 (except for the youngest group), but the pattern over the age distribution remains similar; younger groups experience much higher disutility from cost relative to middle-aged groups (-4.95 and -10.97 for the youngest two groups compared to -4.13 and -2.90 for the 45-54 and 55-64 age groups) and the oldest groups experience roughly similar disutility to one another (-4.88 and -4.03). The effects of other vehicle characteristics are also similar in columns 1 and 2. The coefficients on engine displacement and vintage decrease slightly in magnitude and the coefficients on vehicle width and weight increase in magnitude, but none are statistically different.

The most noticeable differences between columns 1 and 2 of Table A.3 are the risk coefficients, which decrease substantially for all age groups. In addition, the updated coefficients in column 2 reveal an inverted-U shape, increasing in magnitude from -0.76 for the youngest group to -11.37 for the 55-64 age group, and then falling again to -4.35 for the oldest group. The inverted-U shape can be seen in Figure A.3, which depicts the risk and cost coefficients from column 3 for each group.

VSL estimates for each group calculated from the coefficients on fatality risk and cost, and average income, are shown at the bottom of Table A.3. Similar to the pooled-risk scenario in the first column, the inverted-U shape of the risk coefficients implies a strong inverted-U

shape to the age-VSL function. The lower magnitude of the risk coefficients is reflected in smaller VSL estimates, which start at \$0.92 and \$0.87 million for the youngest two cohorts, increase to \$26.32 million for the 55-64 age group, and fall to \$6.70 million for the oldest group. This inverted-U pattern can be seen in Figure A.4, which depicts the VSL estimates for each group.

To test whether the VSL point estimates for each age group in Figure A.4 are significantly different from one another, Table A.4 presents a set of pairwise hypothesis tests for differences in VSL across age groups. The VSLs of the youngest two age cohorts are both significantly lower than the 55-64 age group (differences of 25.40 and 25.46, shown in column 4). Likewise, the oldest two cohorts are also significantly lower than the 55-64 age group (differences of -18.05 and -19.63, shown in columns 5 and 6). The diagonal values of Table A.4 show the differences between each group and the group that is immediately younger and confirm that the inverted-U pattern in Figure A.4 is statistically significant.

An advantage of the multinomial choice framework is that I can separately identify the disutility from risk and cost, which jointly determine the VSL. In this way, I can determine whether age variation in the VSL is due to changing disutility from risk or cost, or both. The age-VSL function in Figure A.4 has two regions. First, the region between 35 and 85 forms a peak that defines the inverted-U shape of the VSL. Over this age range the disutility from cost depicted in Table A.3 is roughly constant and the peaked shape is determined by the disutility from fatality risk, which forms a symmetric inverted-U around the 55-64 age group. For ages younger than 35, the disutility from risk is lower, but combines with a substantially higher disutility from cost to yield a much lower VSL. Both younger and older cohorts exhibit lower

VSLs than the middle-aged group, but their lower VSLs are the result of slightly different phenomena. Young consumers have a low VSL because of the combination of high marginal utility of money and low marginal utility of safety. In contrast, the lower VSL of older groups is driven almost entirely by lower marginal utility of safety.

The VSL estimates in Table A.3 control for age-related differences in driving ability by separately calculated fatality risk for each age group. It is likely however that even within age groups drivers sort into vehicle types according to unobserved behavior. Thus, even age-specific deaths per mile may not accurately reflect the perceived riskiness of each alternative automobile. To account for this, I estimate relative driver behavior using laboratory crash test results and real-world single-vehicle fatalities, and then I adjust the risk value for each alternative within each individual's choice set to capture the riskiness of each vehicle if she were actually the driver.

1.6 Correcting for Driver Behavior within Age Groups

Consumers choose automobiles based on perceived safety, but the number of fatalities per mile depends simultaneously on the behavior of individuals who choose to drive each vehicle. Laboratory crash tests conducted by NHTSA represent a parallel measure of vehicle safety that do not depend on driver behavior.²⁷ Nevertheless, using crash tests to measure risk in lieu of real-world fatalities is not feasible; in order to interpret the ratio of the risk and cost

²⁷ NHTSA conducts regular vehicle crash tests and publishes the results. In the 2011 model year roughly 60 percent of the light vehicle fleet was tested (NHTSA, 2012). Vehicles are retested after changes to an existing model's structure or safety features, and safety ratings for model years without substantive changes generally carry over (Hershman, 2001). I assume that once a vehicle is tested, all subsequent model years carry the same ratings until the model is tested again.

coefficients as a VSL, it is necessary to have risk measured in fatalities per year. Crash tests measure the force experienced by a crash-test dummy and cannot be intuitively compared to costs per year.

Rather than use the crash test results directly as a variable in the automobile choice decision, I use crash test results paired with single-vehicle fatality rates to estimate drivers' relative propensities for fatal accidents. Then I use these estimates to correct the risk variable for the effects of drivers' behavior. One measure of crash performance reported by NHTSA is the head injury criterion (HIC), which reflects head trauma experienced by the crash-test dummy and is strongly correlated with crash fatalities (Kahane, 1994). By comparing vehicles with similar HIC ratings but different real-world fatality rates, I estimate the relative propensity for experiencing a severe accident among the drivers of different vehicle types.

The first step is to use crash test results and fatality rates to estimate relative driver behavior for different types of automobiles. There is a large body of existing literature examining the relative safety of automobiles in the context of light duty CAFE standards. Recently, some studies, in particular Jacobsen (2013), have used crash test results and vehicle fatalities to estimate driver behavior directly.²⁸ In this paper I follow a similar approach and model the number of fatalities in a given automobile as a function of driver behavior and crash test performance.

To begin, I focus on single-vehicle fatalities where only one car is involved (for example, collision with a fixed object) so that only one individual's behavior is in question – the driver of

²⁸ Other studies often focus on the effects of relative vehicle weight on safety, controlling for unobserved differences using various techniques; see for example Anderson and Auffhammer (2011), Gayer (2004), and Sheehan-Connor (2012).

the car whose riskiness I am trying to assess. I model the number of single-vehicle fatalities per mile for each age group in any given make, model, and vintage automobile $Y_{i,c}$ as a function of the engineering safety of each vehicle, an age-group fixed effect that captures age-specific behavior and physical vulnerability, and the average behavior of drivers for each vehicle type:

$$E[Y_{i,c}] = \omega_{i,c} = \alpha_v \lambda_c X_i \quad (8)$$

where $\omega_{i,c}$ is the actual age-specific realization of single-vehicle fatalities for each make, model, and vintage, α_v is the average behavior of drivers for each type of vehicle, λ_c is an age-group fixed effect, and X_i is the crash-test HIC rating for each vehicle. In equation 8, α_v captures the relative likelihood of getting into a severe accident based on vehicle type and X_i captures the likelihood of a fatality conditional on being in a severe accident. λ_c captures the fact that age may affect both the likelihood of getting into an accident and the likelihood of a fatality conditional on being in an accident. The parameters α_v and λ_c in equation 8 can be estimated by fitting the model to a Poisson distribution.²⁹

Table A.5 shows the results of estimating relative driver ability based on equation 8 after dividing the sample into 11 vehicle types based on size and body style. The values are expressed as incident rate ratios and represent the proportional increase in fatalities due to driver behavior relative to compact passenger cars. The results are generally in line with the existing literature on driving ability (Jacobson, 2011). Minivan drivers are the safest, with an accident propensity of 0.64 relative to compact car drivers. Pickups and sports car drivers, on the other hand, are more than twice as likely to get into fatal accidents. The third riskiest driver group is large SUV

²⁹ Using a negative binomial distribution rather than a Poisson distribution yields similar results.

drivers, followed by heavy cars, medium cars, and small SUVs, although the last three are not statistically distinguishable from compact car drivers.

The next step is to use the estimates of $\widehat{\alpha}_v$ from equation 8 and Table A.5 to correct the risk variable for each vehicle within each individual's choice set to find a counterfactual risk for each individual if she were actually driving each alternative. To start, I model the total number of fatalities in a given automobile (for a given age group) as a function of driver behavior, age, the average behavior of all other drivers, and the relative robustness of an automobile relative to all other automobiles. The key assumption is that driver behavior affects single-vehicle fatality rates and multiple-vehicle fatality rates in a proportional manner. In this case, I model overall fatalities in a given automobile analogously to equation 8:

$$E[Z_{i,c}] = z_{i,c} = \alpha_v f(\lambda_c, \bar{\alpha}, \beta_i, \mu_{i,c}) \quad (9)$$

where α_v is the same measure of driver behavior found in equation 8, $\bar{\alpha}$ captures some measure of everyone else's ability, β_i measures the physical sturdiness of the vehicle relative to all others, and $\mu_{i,c}$ captures all other factors that may affect the number of fatalities per mile.

Based on equation 9, finding the counterfactual number of fatalities in vehicle type j – if *someone of type i were driving* – does not require estimating any parameter other than α_v . I simply multiply the fatality rate for vehicle j by the ratio of behavior for drivers of type i and j :

$$\tilde{z}_{j,c}|_i = \alpha_i f(\lambda_c, \bar{\alpha}, \beta_i, \mu_{i,c}) = \frac{\alpha_i z_{j,c}}{\alpha_j} \quad (10)$$

The parameters α_i and α_j in equation 10 can be replaced with their estimates from equation 8.

This does not require estimating any part of equation 9, only the assumption that the same α

enters equations 6 and 7, i.e. driver behavior affects the likelihood of a single-vehicle accident in the same way that it affects the likelihood of multiple-vehicle accidents.

Based on equation 10, I adjust the risk variable for each automobile within each individual's choice set by multiplying by the ratio α_i/α_j , where α_i is calculated for the chosen automobile and α_j is calculated for each specific alternative. Then I rerun the multinomial choice model from equation 4. These results are shown in the third column of Table A.3. The shape and scale of the estimates are consistent, except that the peak VSL for the 55-64 age group is much lower.³⁰ Figure A.5 shows the corrected VSL estimates (solid line) overlaid against the results from Figure A.4 (dotted line). The difference between the peak estimates is visually striking (19.23 million versus 26.32 million), but due to large standard errors the two cannot be distinguished statistically. For comparison, Figure A.6 shows the risk and cost coefficients from column 4 (solid line) overlaid against the coefficients from column 3 (dotted line). The lower peak of the corrected VSL estimates appears to be driven largely by smoother cost coefficients, as the dip in magnitude for the 55-64 age cohort is no longer observed in the adjusted setting.

VSL estimates based on behavior-adjusted age-specific vehicle fatality risk fall closely in-line with existing empirical estimates. The population-weighted average VSL based on vehicle choice is \$9.2 million, which is quite similar to the \$6.3 million used by NHTSA and the \$7.9 million used by the EPA.³¹ The results are also on par with studies that focus specifically

³⁰ Regressions in Table A.3 assume constant marginal utility for non-risk, non-cost covariates across age groups (these variables are not interacted with age-group indicators). For robustness, I repeat the regression in column 3 separately for each age group, allowing all coefficients to vary by age. The VSL results are not statistically distinguishable from column 3, although the peak occurs slightly earlier at \$21.3 million for the 45-54 age group.

³¹ These results are also similar to hedonic vehicle price studies such as Atkinson and Halvorsen (1990), who estimate a VSL of \$6.6 million, and Dreyfus and Viscusi (1995), who estimate a VSL of \$5.6 million. All figures are adjusted to 2009 dollars using the CPI-U.

on the effects of age. Figure A.1 shows the adjusted VSL estimates from Figure A.5 overlaid against recent age-VSL estimates from the hedonic and contingent valuation literature. The estimates from this study share the same general magnitude and strong inverted-U shape as their labor market counterparts, although the peak VSL occurs at later ages. A key attribute of this approach is that automobile-based VSL estimates extend beyond the working-aged population.

1.7 Caveats

Although these estimates comport closely with existing studies, there are several important caveats to keep in mind when interpreting these results. The first is that the FARS only includes accidents where at least one fatality occurred. As a result, I do not control for non-fatal injury risk when estimating age-specific VSLs. This is consistent with the vast majority of the age-VSL literature. Notable exceptions are Aldy and Viscusi (2003, 2008) who include on the job injury risk and worker's compensation replacement rates as additional controls in their hedonic wage frameworks. Viscusi and Aldy (2007) also include the same injury-risk control variables and estimate a VSL with and without controlling for non-fatal risk. Their VSL estimates decrease by roughly 30 percent for all age groups (and from an overall average of \$9.1 to \$6.3 million) when they control for on the job injury risk. If morbidity and mortality risk are positively correlated in automobile accidents, my estimates likely overstate the VSL.

I also have also not controlled for bias that may result from potentially omitted variables correlated with automobile cost. I assume that each consumer is a price-taker and unable to individually affect prices. But used car prices could be correlated with other desirable qualities that I cannot observe, such as comfort, style, build quality, or brand perception. If there is a

positive correlation between price and these omitted (desirable) qualities, the coefficients on the cost variables may be understated. For example, consumers may appear less averse to a price increase because they also receive desirable options such as leather seats (which are unobservable in the data). Nearly all hedonic wage studies are similarly susceptible to omitted variable bias due to unobservable job characteristics that are correlated with risk and wage. In the automobile choice context, omitting desirable vehicle characteristics will lead me to overstate the VSL and may affect the shape of the age-VSL relationship if there are systematic differences in the way that certain age groups value these unobserved features.

Another caveat related to automobile cost is that income in the NHTS might be the incorrect baseline for using cost as a share of income to calculate a VSL. Perhaps some measure of total wealth should be used instead of income. This may be important because seniors have a relatively low income-to-wealth ratio relative to younger groups. For example, according to the 2010 Survey of Consumer Finances, households with a head younger than 35 had average assets worth \$141,000; households with a head between the ages of 65 and 74 had assets averaging \$947,000.³² In my sample, seniors have similar income on average compared to other groups. If I am understating the wealth of seniors more severely than younger groups I could also be overstating the VSL for seniors, which would make my estimates of an inverted-u shape conservative.

An additional concern is that children may be affecting the observed VSL for their parents. For example, the peak at middle age might come from parents' concern for their

³² The Survey of Consumer Finances is a cross-sectional survey of US families' balance sheets, pensions, incomes, and demographics. The survey is conducted by the Federal Reserve Board and data are available at: <http://www.federalreserve.gov/econresdata/scf/scfindex.htm>.

children's safety as passengers rather than a valuation of own-mortality risk. To see whether concern for children is driving the inverted-U shape of the VSL function, I repeat the exercise in column 3 of Table A.3 using a restricted sample that excludes households with children under 18. The resulting VSL estimates shown in Figure 8 display a similar pattern and are not statistically different from those based on the full sample.³³ Parents may also affect measured VSL for younger drivers if a parent purchases a vehicle for their driving-aged child. To mitigate this I only include drivers over the age of 18.

The shape of the VSL function could also be driven in part by credit constraints in the market for automobiles. The cost coefficients implicitly reflect underlying differences in borrowing ability for different age groups. If a consumer is not fully able to borrow against her future earnings, she may appear to place a higher importance on cost relative to other desirable attributes, such as safety, size, performance, or vintage. Compared to a world with no credit constraints, I may be overstating the importance of cost relative to safety and thus understating her VSL. Disproportionately binding constraints on younger and older consumers could contribute to the inverted-U shape seen in Figure A.5.³⁴

A final concern involves the time period in question. Data for this analysis come from the 2009 cross-sections of the NHTS and FARS, collected during the midst of a substantial economic downturn in the US. This could influence the coefficient estimates in unpredictable ways. For example, car purchasers may have been more cost sensitive during the recession, or

³³ Having children could also increase parents' valuation of their own safety. This effect should be included in the VSL estimate. Aldy and Viscusi (2003) check for robustness in their labor market study by controlling for children via interaction terms and find no difference in the VSL of workers with or without children.

³⁴ Even without differences across age groups, credit constraints can contribute to age variation in the VSL. Shepard and Zeckhauser (1984) explore a lifecycle consumption-allocation model and find that a no-borrowing restriction leads to a much more pronounced inverted-U shaped WTP relative to perfect markets (though both scenarios still produce an inverted-U shape).

more pessimistic about future consumption, both of which would affect the relative importance of cost and risk. In addition, the auto industry was particularly affected by the economic downturn, which may have affected production, pricing, and marketing of new cars in ways that spilled over into the used car market. It is generally uncertain what affects the recession may have had on my final estimates.

1.8 Conclusion

Vehicle purchase decisions present a unique opportunity to examine the willingness to pay for reduced fatality risk among older individuals not typically represented in the labor force. Cost and safety are both important factors determining which automobiles consumers choose, and since different groups place varying emphasis on each of these attributes, we can observe how the willingness to pay for safety changes over the age distribution. Based on this, I estimate an age-VSL relationship for a wider age range and expand the revealed-preference estimates into a region not typically captured under prior labor market studies.

Initial estimates based on uniform fatality risk for all age groups (calculated based on total deaths per mile for each vehicle) suggest unreasonably high VSL estimates. After controlling for bias due to individuals sorting into different vehicle types, the automobile choice framework suggests an inverted-U shaped age-VSL relationship that is on par with existing empirical findings. The peak VSL for the 55-64 age group (\$19.23 million) is more than twice as large as the VSL for the 75-85 age group (\$8.15 million), and eight times as large as the VSL for the 18-24 age group (\$2.38 million). Pairwise tests for differences in the VSL across various age groups also confirm the inverted-U shape.

A specific policy question surrounding age and VSL is the appropriateness of a discount to the VSL for individuals over the age of 65. Controversy over the senior discount became a political issue in 2003 in response to the EPA's Clear Skies Initiative cost-benefit analysis. In one of several scenarios examined, the EPA employed a VSL for seniors that was 37 percent lower (EPA 2002). This provoked political backlash, and as a result, Federal agencies have since used a constant VSL across all age groups.³⁵ An automobile choice framework provides a unique opportunity to test the importance of a senior discount.

Using the estimates in column 3 of Table A.3, I calculate the population-weighted average VSL for the 18-64 age group and the 65-85 age group. The results are nearly identical (\$8.91 million versus \$8.80 million) and the difference is not statistically distinguishable from zero. In this case, the inverted-U shape leads to an average VSL that is the same for seniors and non-seniors. This is because the inverted-U shape in Figure A.5 peaks in the 55-64 age group, so taking averages above and below 65 – an already arbitrary cutoff – yields similar results for both groups.

Overall, VSL estimates based on vehicle choice reveal a strong inverted-U shape to the age-VSL function. Aggregating these figures into larger age groups (such as seniors and non-seniors) masks the underlying heterogeneity. The results of this study corroborate the growing consensus that the VSL first increases and then decreases with age and provide a basis for discussing age-VSL policy pertaining to individuals not captured in labor market data. These results also highlight the importance of considering the specific ages of affected individuals when evaluating the benefits of health and safety regulations.

³⁵ See, for example, Viscusi (2009).

CHAPTER II

ENVIRONMENTAL ENGEL CURVES

2.1 Introduction

This paper presents the first-ever estimates of household-level environmental Engel curves (EECs) showing the relationship between households' incomes and the amount of pollution embodied in the goods and services those households consume. Traditional Engel curves plot the relationship between income and consumption of a particular good. These original Engel curves are named for Ernst Engel, a German economist writing in the mid-1800s, who studied the degree to which household food expenditures increase with income. Engel curves have since been applied to all manner of goods and services (food, housing, transportation, etc.) and form the basis for "equivalence scales" that are used to determine eligibility for means-tested entitlements such as food stamps and Medicaid.

Environmental Engel curves describe how households' pollution changes with income. This calculation is less straightforward than traditional Engel curves because households generate pollution both directly as a consequence of their activities, such as driving cars, and indirectly as a consequence of consuming products whose production generates pollution, such as manufacturing the rubber and steel used to make those cars and refining the gasoline used to fuel them. We focus on the larger and less studied component, which is the indirect pollution generated through the production of goods and services consumed by households. We measure the effects of income growth on pollution by tracking the underlying changes in household consumption.

Economists have generally parsed the effect of economic growth on the environment into three components: scale, technique, and composition (Grossman and Kruger, 1993; Copeland

and Taylor, 2005; and Levinson, 2009). The scale component describes a proportional increase in economic activity, either from population growth or per capita economic growth. The technique component describes changes to the pollution intensity of any particular activity. For example, generating electricity is less pollution-intensive today than 30 years ago because more coal-fired power plants have abatement equipment. Last, the composition component describes changes to the mix of activities in the economy, such as changes in international trade patterns or in the mix of goods and services consumed by households. Much prior research has been devoted to estimating the effects of each of these various components, and this paper focuses on the relative importance of the scale and composition effects in relation to household income growth. Figure B.1 summarizes this project's relative position in that literature.

The obvious approach to estimating EECs would be to compare pollution, income, and consumption choices across countries at a point in time or across time within a country. This would be similar to the environmental Kuznets curve (EKC) literature that examines pollution and income on an aggregate level and typically finds an inverted-U relationship between pollution and economic activity (See, for example, Grossman and Krueger (1995) and Andreoni and Levinson (2001)). But EKCs are based on correlations with little economic interpretation. Engel curves, on the other hand, are based on the underlying structure of individual preferences and have a longstanding importance in policy analysis. Environmental Engel curves can be used to infer the future direction of environmental quality and have important ramifications for environmental policy.

Relying on variation across countries or over time to measure EECs is not viable because Engel curves must hold constant all other factors, including prices and characteristics of available goods. For example, richer countries may pass regulations making pollution-intensive

goods costlier or less desirable, causing households to consume proportionally fewer polluting goods. That difference in consumption of polluting goods could not be interpreted as the slope of an Engel curve because it does not represent the change in consumption that results from a change in household income, holding all else equal.

To examine the relationship between income and the pollution embodied in the goods and services consumed by households while holding all else constant, we estimate a set of EECs using household-level data within the United States at specific points in time. In this way, households within a given cross-section face the same relative prices, available products, and environmental regulations. We combine production-side pollution intensity data with detailed information on household consumption to calculate the total pollution created as a result of producing the goods and services that each household consumes. Using these combined data we estimate a set of annual EECs from 1984 to 2002.

One application for such EECs is to isolate the change in pollution associated with household income growth. For example, over the past 30 years total pollution emitted by US producers has declined considerably even though the real value of US production has increased.³⁶ This improvement has come in part from cleaner production technologies and changing patterns of international trade, but much has also come from US households consuming a cleaner mix of goods and services. Some analysts have pointed to these improvements as evidence that richer countries automatically pollute less, implying that economic growth alone will improve the environment.³⁷ But this inference combines the effects of income growth with the effects of changing prices and regulations within the US. Having a set of annual

³⁶ From 1980 to 2012, emissions of carbon monoxide and sulfur dioxide declined by roughly 80 percent, ground-level ozone by 25 percent, and nitrogen dioxide by 60 percent, even though real GDP and real personal consumption expenditures more than doubled (EPA 2014, FRED 2014a and 2014b).

³⁷ For example, Bartlett (1994) claimed that "existing environmental regulation, by reducing economic growth, may actually be reducing environmental quality."

environmental Engel curves allows us to separate the effects of regulation-induced price changes from coincidental preferences of richer households for cleaner goods

We find that EECs display three key characteristics. First, not surprisingly, richer households are responsible for more overall pollution (EECs are upward sloping). Second, although pollution increases with income, it does so at a decreasing rate (EECs are concave). And third, EECs shift down over time. Applying these characteristics to isolate the effects of household income growth, we find that between 1984 and 2002 pollution per household decreased by six percent while income grew by 13 percent. In addition, 60 percent of the offset pollution (relative to the 13-percent scale increase) was due specifically to richer households consuming less pollution-intensive goods, i.e. moving along a concave EEC, and 40 percent is due to downward shifts in the EECs over time.

2.2 Movement Along Versus Shifts in the EEC

Environmental Engel curves show the relationship between income and the pollution embodied in the goods and services consumed by households. A set of stylized EECs are depicted in Figure B.2, which shows pollution on the vertical axis plotted against income on the bottom axis. Under standard Engel curve definitions, goods whose consumption declines with income are "inferior;" goods whose consumption increases with income are "normal;" and among normal goods, those whose consumption increases less than one-for-one are considered "necessities." In Figure B.2, pollution is depicted as a necessity. Additionally, Engel curves require that all other factors are constant, including prices and regulations, so Figure B.2 shows two separate EECs for two separate time periods.

With a set of unique Engel curves for different periods, we can observe two simultaneous phenomena. By first focusing on a single EEC we can draw conclusions about how the pollution intensity of consumption is likely to evolve as household income changes but prices and environmental regulations remain fixed. This is tantamount to moving along a given Engel curve. In the hypothetical case depicted in Figure B.2, households with higher income are responsible for more pollution overall, but consume a less pollution intensive mix of goods and services.

In reality, individual EECs are likely to have a shape similar to those in Figure B.2. For example, richer households may consume relatively more accounting services or insurance and relatively less food or gasoline compared to poorer households.³⁸ If this is the case, then increases in household income should lead to relative reductions in the pollution intensity of consumption. It is also possible that pollution becomes more or less of a necessity at different income levels. The hypothetical scenario in Figure B.2 depicts relative increases in pollution that are smaller for higher income levels

A second phenomenon observable from EECs is the relative change in the position of each annual EEC over time. Since each individual EEC holds prices and regulations constant, a different set of prices leads to different EEC. If any prices or economy-wide factors change in such a way that leads households to consume a different mix of goods and services, the EEC will shift. For example, a downward shift in the EEC could occur if a country passes environmental regulations that make pollution intensive goods more expensive. Likewise, aggregate changes in consumer preferences towards cleaner goods could also cause a downward shift in the EEC.

³⁸ Richer households are also likely to save a larger share of their income, which would contribute to a concave EEC based on income.

The distinction between movements along the EEC and shifts of the EEC is key for understanding the effects of household consumption on the environment. Both movements along and shifts in the EEC characterize changes in the pollution intensity of household consumption that results from changing the mix of goods and services consumed by households. The difference is that shifts in the EEC are the result of economy-wide changes, separate from individual households' income growth, and are therefore not predictive of how pollution may evolve as a county grows richer unless those economy-wide trends continue. For example, tighter environmental regulations would lead to a downward shift in the EEC, but the EEC will not continue shifting downwards unless environmental regulations are continually tightened.

Movements along the Engel curve, however, are indicative of future pollution levels that may result from household income growth. The specific shape of an individual EEC depends only on the underlying preferences of households at different income levels and is independent of aggregate trends in prices or regulations. If household income growth were the only change, the EEC could be used to predict the pollution associated with the mix of goods and services consumed by households at their higher income levels.

The actual overall change in pollution intensity of consumption over time is a combination of movements along and simultaneous shifts in the EECs. A set of annual EECs allows us to decompose those changes into these two separate components. In this way we can determine what portion of the changes in overall pollution as well as pollution intensity of consumption is due directly to household income growth, and hence likely to continue as the economy grows, and how much is due to other factors such as changing prices or regulations, which would require continual tightening to maintain a trend of reduced pollution.

2.3 Data

Estimating EECs requires information on household income and the pollution attributable to each household's consumption. Since we are focusing on indirect pollution, we estimate the amount of pollution that was created in order to produce the specific goods and services consumed by each individual household in our sample. To do this we use information from the Consumer Expenditure Survey (CEX), which provides detailed information on itemized household consumption expenditure, paired with the Trade and Environmental Assessment Model (TEAM), which provides the pollution intensity of production for various goods and services.

Household consumption data come from the CEX, collected each quarter by the Census Bureau on behalf of the Bureau of Labor Statistics. The survey component of the CEX is a nationally representative sample of roughly 7,000 households selected on a rotating panel basis.³⁹ Households are tracked for five consecutive quarters, over which they provide information on their complete range of expenditures, income, and other demographics. Expenditure and income data in the CEX interviews are organized into roughly 700 separate universal classification codes (UCC) and capture roughly 80 to 95 percent of total household expenditure.⁴⁰

Each round of CEX contains households from every stage of the five-quarter interview process. These data have been combined and extracted into a publicly available, user-friendly

³⁹ The CEX is organized based on consumer units, rather than households. A consumer unit is smaller than a household and consists of: "(1) All members of a particular household who are related by blood, marriage, adoption, or other legal arrangements; (2) a person living alone or sharing a household with others or living as a roomer in a private home or lodging house or in permanent living quarters in a hotel or motel, but who is financially independent; or (3) two or more persons living together who use their incomes to make joint expenditure decisions" (Bureau of Labor Statistics, 2008). For convenience the BLS occasionally treats the terms "households" and "consumer units" interchangeably and we follow suit.

⁴⁰ The interview survey collects detailed data covering 60 to 70 percent of household expenditures, along with global estimates each period for food and related items that capture an additional 20 to 25 percent of expenditures. For detailed information on the CEX, UCC codes, and the structure of the survey, see Bureau of Labor Statistics (2008).

format by the National Bureau of Economic Research (NBER, 2000). These extracts consolidate family records across interview rounds and provide a single record for each family showing annual expenditure values. The end result is quarterly repeated cross-sectional data with unique observations by household. In addition, the roughly 700 UCC codes are collapsed into 109 spending and income categories that remain consistent across all sample rounds. Of these expenditure categories, 47 correspond to various goods and services while the remaining categories cover income, taxes and transfers, and measures of wealth. These consolidated data are a convenient starting place for our investigation.

We pair the NBER CEX extracts with pollution intensity data from the TEAM. The TEAM framework was originally developed for the EPA to highlight the environmental effect of US trade policies, but the model can be used to assess the effect of any change in economic activity (Abt Associates Inc., 2009). TEAM has as its core a list of emissions intensities by six-digit North American Industry Classification System (NAICS) codes.⁴¹ For each of over 1000 industries, TEAM reports the amount of pollution (sulfur dioxide, carbon monoxide, particulate matter, etc.) emitted per dollar of industrial output in three separate years: 1997, 2002 and 2007.

The TEAM data indicate the pollution generated during the production process of each good directly, but we also want to consider pollution from production that occurs further up the supply chain. For example, if a household purchases a sofa, we would want to know the pollution emitted while assembling of the sofa itself, but also the pollution from tanning the leather for its upholstery, milling the wood for its frame, and manufacturing the steel for its springs. Moreover, each of those inputs also required their own inputs which also resulted in

⁴¹ The North American Industry Classification System (NAICS) was developed jointly by the US, Mexico, and Canada to classify industries based on similarities in production processes. UCCs on the other hand categorize goods based on similarities in consumption patterns.

pollution. To fully capture the total pollution associated with each household's consumption, we want to include both direct pollution and pollution from all inputs further up the supply chain.

Upstream pollution from the entire chain of inputs for each item can be estimated using a Leontief (1970) input-output (IO) analysis based on IO tables published by the Bureau of Economic Analysis. These IO tables show the amount of each input that is necessary to produce a unit output of each other industry. Under linear production, this "direct requirements" matrix can be combined in series to give a "total requirements" matrix showing the total amount of each industry (through the entire supply chain) necessary to produce a unit of each other industry. Using this total requirements matrix we transform the TEAM emissions intensities into measures that include all of the inputs to each industry, the inputs to those inputs, and so on.

Using the total pollution intensity coefficients combined with the CEX we estimate the total amount of pollution that was created in order to produce the exact quantity of each of the goods and services consumed by every household in the survey. Adding up pollution for all items within a household gives the total amount of pollution attributable to the overall consumption of each household. Because the CEX and TEAM use different classification systems, we manually created a concordance to match consumption items in the CEX with pollution intensity in the TEAM.⁴² Since we are interested in the income-pollution relationship holding all else constant, we apply the 1997 TEAM values to all cross-sections of consumption data, essentially holding constant the pollution intensity of production.

In the same way we fix technology, we also exclude international trade. When calculating the total amount of pollution associated with each household's consumption, we

⁴² Since the TEAM coefficients (based on NAICS codes) consist of more categories than the CEX, most CEX codes were matched to several NAICS categories. In this case, we calculated the weighted average pollution intensity based on total output for each NAICS code. See Appendix I for a detailed discussion of matching between NAICS and CEX categories.

implicitly assume that the US-based TEAM coefficients characterize the pollution intensity at each stage of production. In this way, we measure the pollution that would have occurred if all production took place in the US. Although this does not capture the true amount of pollution created as a result of each household's consumption, it does allow us to separate the change in pollution due to compositional shifts in consumption from the simultaneous trends in offshoring of manufacturing and international transportation.

As a last step, we exclude any households with incomplete or partial-year income reporting and drop the top and bottom one percent of households based on total expenditure (to account for top-coding in the CEX survey). The final result is a sample of 57,704 household spread across 19 annual cross-sections from 1984 to 2002, in which each household has an estimated total pollution associated with its consumption expenditure.⁴³ At this point we focus on particulate matter less than 10 microns (PM10) because of its prominence in cost-benefit analysis and its observable public health consequences.⁴⁴ Table B.1 shows mean values for indirect PM10, income, and other variables for the 1984 and 2002 cross-sections of data.

2.4 Estimating Environmental Engel Curves

Before estimating EECs, it is worthwhile to acknowledge several challenges to estimating traditional Engel curves and how they apply to estimating environmental Engel curves. First, traditional Engel curves are susceptible to bias due to unobserved variation in the quality of goods that households consume. If richer households purchase more expensive, higher quality goods, it may appear as though these households spend a larger share of their income on a

⁴³ We exclude CEX rounds prior to 1984 because this is the first year with integrated diary and interview data, and the first year with both urban and rural households included (See BLS 2014). 2002 is the most recent CEX round with all four quarters available as NBER extracts.

⁴⁴ This analysis can be repeated with other pollutants to find similar results.

particular type of good (or purchasing a larger volume). For example, if rich households purchase expensive wine and poor households purchase cheap wine, an estimated Engel curve for wine will be biased.

Although pollution is not consumed directly, we attribute pollution to specific households based on the estimated amount of pollution generated through the manufacturing of the goods and services each household consumes. Because we estimate overall pollution by multiplying itemized expenditure by per-dollar pollution intensity coefficients, expensive items are assigned more pollution than inexpensive items. For example, an expensive bottle of wine that costs twice as much as a cheap bottle will be assigned twice as much pollution, even if it came from the same winery. If richer households consume more expensive products and poorer households consume cheaper products, we may overstate the pollution intensity of consumption for richer households relative to poorer households, and vice versa.

A second challenge to estimating traditional Engel curves is determining the appropriate degree of aggregation. Demand for narrow categories, such as a specific type of wine, can vary sharply across households and over time, making patterns difficult to discern. But broader categories, such as all beverages, may combine inferior and normal goods and cancel out different shapes to the underlying Engel curves. This is not an issue when estimating EECs. Even if some high-income households spend their extra money on fancy wine, while others buy expensive concert tickets, what matters for this analysis is the overall pollution created indirectly as a result of each household's consumption, not the specific consumption of individual goods or services. Using a more disaggregate set of expenditure categories to estimate EECs will not separate normal and inferior goods, as in traditional Engel curves, but may provide a finer measure for capturing compositional changes in consumption across different households. It will

also make it more difficult to match pollution intensity coefficients to narrowly defined expenditure categories.⁴⁵

Since no theory dictates the form of the income-pollution relationship, the goal is to first get an idea of the shape and structure of the EECs with as few restrictions as possible.⁴⁶ To begin, we separate households in the 1984 cross-section of the CEX into 50 groups based on income, where each group represents two percent of the overall 1984 income distribution. Then within each group we calculate the average level of PM10 associated with each household's consumption. Plotting these 50 points with income on the horizontal axis and pollution on the vertical axis yields a non-parametric EEC for PM10 in 1984. This EEC is displayed as the top line in Figure B.3. Since each of the 50 groups captures two percent of households, there are an equal number of households represented by each point. Further, average pollution is calculated using 1997 TEAM coefficients, so this EEC represents the pollution associated with household consumption 1984 if all goods and services were produced using 1997's production technology. For example, households that fall into the 25th income bin in (\$32,128 to \$33,999 measured in 1997 dollars) were indirectly responsible for an average of 147.23 tons of PM10.

To observe how the EEC relationship may be evolving over time, Figure B.3 also depicts a second EEC estimated based on the 2002 cross-section of the CEX paired with the 1997 TEAM coefficients. In order to keep the two curves directly comparable, we use the same income bin cutoff values in the 2002 EEC as are used in the 1984 EEC.⁴⁷ Households with

⁴⁵ We rely on the NBER CEX extracts that use a relatively aggregate set of expenditure categories, but acknowledge the additional precision and difficulties that may come with using the disaggregate set of expenditure categories in the original CEX data.

⁴⁶ Common approaches range from simply plotting the data, to nonparametric kernel estimation (Lewbel, 1991; Hausman, *et al.* 1995).

⁴⁷ In this case, each point of the 2002 EEC does not represent an equal number of households. Income growth between 1984 and 2002 led to a rightward shift of the income distribution, but since income bin cutoff values are determined based on the 1984 income distribution there are relatively fewer households in low bins and relatively more households in higher bins.

income in 2002 between \$32,128 and \$33,999 were indirectly responsible for 131.77 tons of PM10 on average, 15.46 fewer tons than households earning the same income in 1984.

Three phenomena are apparent from the set of EECs shown in Figure B.3. First, richer households are responsible for more overall pollution. This is not surprising since richer households spend more on consumption and are therefore expected to have more pollution created as a result of the goods and services they consume. Second, EECs are concave, meaning that incremental increases in household pollution become smaller as income grows larger. This means that richer households consume less pollution-intensive mixes of goods and services, even if they are responsible for more overall pollution.⁴⁸ This effect may be understated if richer households actually consume a similar quantity of goods that happen to be more expensive. In that case, EECs are likely to be more concave than those depicted in Figure B.3.

Third, the set of EECs in Figure B.3 also suggest that EECs are shifting down over time. The shape and concavity are generally consistent in both years, but households represented by the 2002 EEC are generally responsible for less pollution relative to their 1984 counterparts across the entire income distribution (and particularly above the median). These EECs suggest that households adjusted the composition of their consumption towards a less pollution intensive mix of goods and services over time. In addition, this downward shift is not due to improvements in technology or abatement since we have fixed the pollution intensity of production using the 1997 TEAM coefficients for both years. The shift over time observed in Figure B.3 is due to changing prices, regulations, and social norms.

The EECs in Figure B.3 are based on total household income measured in the CEX. One concern is that a progressive US tax system (in which higher income households pay a higher

⁴⁸ Under standard Engel curve definitions, goods whose consumption increases with income are considered “normal;” and among normal goods, those whose consumption increases less than one-for-one are considered “necessities.” Pollution, according to these EECs, is a necessity.

share of their income in taxes) may exaggerate the concavity of the EECs. In that scenario, higher income households would appear to consume less pollution intensive bundles, but this would be driven by the fact that higher income households have less disposable income relative to their overall income, not by differences in the composition of goods and services they consume. Additionally, consumption relative to income may be affected by underlying changes in the saving rate not observed up in Figure B.3. To account for these factors, we repeat the analysis but measure after-tax income on the horizontal axis in Figure B.4a and total consumption expenditure on the horizontal axis in Figure B.4b. As expected, the EECs in 4a and 4b appear slightly less concave and somewhat closer together. Overall, the same key characteristics of Figure B.3, that EECs are increasing, concave, and shifting over time, appear to hold in Figures 4a and 4b as well.

One drawback of the otherwise flexible approach to estimating EECs depicted in Figures 3 and 4 is that they do not account for additional factors that may affect household consumption. For example, it appears as though households consumed a less pollution-intensive mix of goods and services in 2002 than in 1984, but an alternative explanation may be that households consumed the same mix of goods while average household sizes decreased. Or perhaps the changes were due to general migration patterns as the US population shifted towards regions with different climates and transportation infrastructures. Table B.1 shows the change in average indirect pollution and income for US households between 1984 and 2002, along with changes in demographic variables. Over this period the average indirect PM10 emissions decreased 6.1 percent (from 155.1 tons to 145.6 tons) while average real income increased 13.2 percent (from \$40,970 to \$46,370). At the same time, the average household became older, smaller, better educated, and shifted geographic regions.

To include other factors affecting the quantity and mix of goods and services consumed by households we use a parametric regression with total PM10 on the left-hand side and income, along with other covariates, on the right-hand side. We use a linear regression framework and account for concavity by including income squared on the right-hand side. To begin, column 1 of Table B.2 shows an OLS regression of indirect pollution on income and income squared using the 1984 cross-section of the CEX. Coefficients on both terms are significantly different from zero (19.04 and -0.39, respectively) and corroborate the increasing and concave EECs depicted in Figure B.3.

The second column of Table B.2 adds additional control variables for age, family size, marital status, indicators for race and education of the household head, and regional indicators. Nearly all covariates are significantly correlated with total PM10 at the one percent level, except indicators for some college (relative to less than high school) which is significant at the five percent level and indicators for Asian and Other races (relative to White), which are indistinguishable from zero. Overall the results suggest that larger families, older households with more education, and non-Black households are indirectly responsible for more pollution. For example, better educated households spend more money on food, airfare, and clothing. In addition, households located in the Eastern region consume a more polluting mix of goods and services, driven partially by differences in natural gas, electricity, and food consumption.

Including a set of additional covariates also has a substantial effect on the measured impact of income on total PM10. The magnitude of the coefficients on both income and income squared decrease (to 10.28 for income and to -0.13 for income squared). After controlling for

other factors, the effect of income remains significant at the one percent level and income squared is significant at the ten percent level.⁴⁹

To compare the relationship across time, the columns 3 and 4 of Table B.2 repeat the regression from column 2 using the 1993 and 2002 cross-sections of the NBER CEX extracts (the middle and last years of available extracts). Indicators for high school and some college education are less significant in 2002 relative to 1984, but the remaining covariates generally maintain a consistent statistical significance across the columns. The coefficients on income and income squared in 2002 relative to 1984 also suggest an EEC that is lower and more concave in recent years. The final column in Table B.2 shows the difference between coefficients in 1984 and 2002 (from column 2 and 4, respectively) and indicates whether there is a significant difference. With the exception of age, region, and high school education, the coefficients in 2002 are statistically indistinguishable from those in 1984.

Figure B.5 shows the relationship between income and total PM10 based on the parametric EECs estimated in columns 2 and 4 of Table B.2 (which are based on the 1984 and 2002 cross-sections of the CEX). For each EEC we fix the other covariates at their average value (the 1984 EEC uses average covariates from 1984 and the 2002 EEC uses average covariates from 2002) and map out the relationship between income and total PM10 based on the coefficients for income and income squared in Table B.2. Both parametric and non-parametric EECs describe a relationship between income and PM10 from goods and services that is increasing, concave (at least weakly), and may shift over time. With this in mind, two phenomena appear to occur as income grows over time. The first is that richer households move

⁴⁹ To account for potential non-linear EECs beyond a quadratic form we also ran regressions including higher order polynomial terms (such as income cubed and income raised to the fourth). The higher order terms captured much of the influence of the income variable, which was no longer individually significant, but the joint significance of income terms, fitted values, and goodness-of-fit remained essentially the same.

upward along a given EEC towards a less pollution-intensive mix of consumption. Second, households that fall into a given income bin consume a less pollution-intensive mix over time. One application for EECs is to decompose the overall compositional shifts in household consumption into movement along and shifts of EECs. The following section provides a concrete example based on the EECs in Figure B.3 and Table B.2.

2.5 An Application: Decomposing the Effects of Income Growth

The EECs in Figure B.3 and Table B.2 demonstrate two factors that contribute to compositional shifts in household consumption over time. The first is that income growth moves households upwards along a given EEC under a constant set of prices. This has two opposing effects on pollution: richer households have higher overall consumption (a scale effect), but they also consume a less pollution intensive mix of goods and services (a composition effect). The extent to which households with different income are responsible for different pollution levels is characterized by the shape of individual EECs (which can also change over time). Second, aggregate prices evolve as regulations and changing social norms make pollution intensive goods more expensive. This causes individual households at all income levels to consume a less pollution intensive mix of goods and services, which is reflected in a downward shift in EECs over time.

Both movements along and shifts in the EECs affect the level of pollution embodied in the goods and services consumed by households, but there is an important policy distinction between the two effects. Movements along the EEC depend on underlying preferences of richer households relative to poorer households. The environmental consequences of movement along the EEC are independent of any particular policy intervention. In this sense, movements along

the EEC are predictive of future levels of pollution under status quo environmental regulations. This also means that movements along the EEC provide information about what to expect as income increases in less developed countries. In contrast, shifts in the EEC are the direct result of evolving aggregate preferences or environmental policies that increase the relative price of pollution intensive goods. There is no reason to expect the environmental benefits of downward shifting EECs to continue without the accompanying change in preferences or tightening of environmental policy.

An annual set of EECs allows us to decompose the changes in the mix of goods and services consumed by households into a component due to income growth (a movement along the EEC) and a component due to aggregate prices (a shift in the EEC). For example, we could use the 1984 EEC to assign a hypothetical level of total PM10 to each household in 2002. This would tell us how much pollution to expect if the EEC was fixed based on 1984 prices, but and households could move along the EEC as their incomes changed. The difference between this hypothetical level of PM10 and the actual emissions (after holding technology constant) is due to shifts in the EEC between 1984 and 2002.

In a parametric setting, an Oaxaca (1973) Blinder (1973) decomposition provides a formalized way of separating these components.⁵⁰ To begin, define the average level of pollution in a given year based on the OLS model from Table B.2:

$$\bar{P}_t = \alpha_t \bar{Y}_t + \beta_t \overline{Y_t^2} + \bar{X}_t \delta_t \quad (11)$$

where \bar{P}_t is average indirect pollution, \bar{Y}_t and $\overline{Y_t^2}$ are income and income squared, and \bar{X}_t captures other included covariates.. The error term disappears because the average error in OLS is zero by construction.

⁵⁰ For additional discussion of decomposition techniques, see also Fortin, Lemieux, and Firpo (2010).

Then based on equation 11, the change in average pollution between 1984 and 2002 can be written as:

$$\begin{aligned}\bar{P}_{02} - \bar{P}_{84} = & \alpha_{02}\bar{Y}_{02} + \beta_{02}\overline{Y_{02}^2} + \bar{X}_{02}\delta_{02} \\ & - \alpha_{84}\bar{Y}_{84} - \beta_{84}\overline{Y_{84}^2} - \bar{X}_{84}\delta_{84}\end{aligned}\quad (12)$$

and by adding and subtracting $\alpha_{84}\bar{Y}_{2002} + \beta_{84}\overline{Y_{2002}^2} + \bar{X}_{2002}\delta_{84}$ and grouping terms, we have:

$$\begin{aligned}\bar{P}_{02} - \bar{P}_{84} = & \alpha_{84}(\bar{Y}_{02} - \bar{Y}_{84}) + \beta_{84}(\overline{Y_{02}^2} - \overline{Y_{84}^2}) \\ & + (\alpha_{02} - \alpha_{84})\bar{Y}_{02} + (\beta_{02} - \beta_{84})\overline{Y_{02}^2} \\ & + \bar{X}_{02}(\delta_{02} - \delta_{84}) + (\bar{X}_{02} - \bar{X}_{84})\delta_{84}\end{aligned}\quad (13)$$

The first two terms in equation 13 capture the effect of changing income on total pollution, holding constant the 1984 OLS coefficients. This is equivalent to a movement along the 1984 EEC. The second two terms capture the effect of different OLS coefficients on income and income squared in 2002 relative to 1984. This is equivalent to a shift (or change in shape) of the EEC. Finally, the last two terms account for changes in all other covariates, including demographics, migration, and household size, and their changing coefficients.

Table B.3 presents the results of this decomposition between 1984 and 2002. The first three columns show the mean values of total PM10, income, and other demographic variables in 1984 and 2002 and the change between the two years from Table B.1. Column 4 shows the OLS coefficients based on the 1984 cross-section from Table B.2. The last column is calculated by multiplying the change in average values (column 3) by the 1984 OLS coefficients (column 4) and represents the change in pollution attributable to movement along the EEC (the first three terms of equation 13).

The level of total PM10 embodied in household consumption decreased 9.5 tons between 1984 and 2002 (using the 1997 pollution coefficients). Changes in average income (and income

squared) led to a hypothetical increase of 4.17 tons (5.55 increase from income and 1.39 decrease from income squared), which combines the increased pollution due to a scale effect with an offsetting compositional shift. At the same time, changing demographics offset this increase by 2.71 tons (the sum of column 4 values associated with non-income variables). The remaining difference, 10.97 tons, is attributable to shifts in the EEC.

The hypothetical 4.17 ton increase in pollution due to changes in household income can be further decomposed into separate scale and composition components. Along a given EEC, richer households consume more goods and services overall, but they also consume a less pollution-intensive mix relative to poorer households. The balance of these two effects depends on the shape of the EEC. To the extent that EECs are more concave, the compositional component is prominent and households with higher income are responsible for less than proportionally higher levels of pollution. On the other hand if EECs are perfectly straight there is no compositional component and pollution will increase at the same rate as income growth.

By comparing the hypothetical change associated with movement along the EEC with the rate of income growth, we can observe the extent to which changes in the composition of consumption across rich and poor households offsets the increase in pollution due to the scale effect from household income growth (while maintaining constant prices – i.e. along the 1984 EEC). Between 1984 and 2002 average household income increased 13 percent. With no compositional shift in consumption, we would expect total PM10 emissions to increase proportionally by 20.4 tons. The difference between this number and our movement-related estimate of 4.16 represents the mitigating effect of compositional shifts along the 1984 EEC. In this case, compositional changes in consumption along the EEC offset 16.3 tons of PM10 from the scale effect. In total, the sum of the compositional offsets (-16.3 from movement along and -

10.9 tons from shifts of the EEC) together with the effects of demographics (-2.71 tons) counteract the scale effect (20.4 increase) to equal the overall predicted change of -9.5 tons.

Figure B.6 depicts the relative magnitude of these effects over time between 1984 and 2002 by applying the same decomposition to all interim years. The top line depicts the level of pollution that would occur if households did not adjust the mix of goods and services they consumed other than proportional increases with income. The second line in Figure B.6 captures the hypothetical effect of movements along the 1984 EEC. The vertical difference between these two lines (16.3 tons in 2002) is the offsetting compositional effect reflected in the concavity of EECs. The third line shows the contribution of changing demographics in addition to changing income and falls below the second line because the balance of other factors, such as household size and geography, led to a net decrease in pollution intensity of consumption.

The fourth line of Figure B.6 shows the level of pollution in each year calculated by pairing the 1997 TEAM pollution coefficients with each round of the NBER CEX extracts. This is the level of pollution that would occur if technology were fixed based on 1997 pollution intensities, but we account for the true mix of goods and services consumed by households in each period. The vertical distance between the third and fourth lines (10.9 tons in 2002) is due to downward shifts in the EEC over time.

As Figure B.6 demonstrates, compositional shifts in the mix of goods and services consumed by the average household have more than offset any increases in PM10 due to the scale effect of income growth. To this end, compositional changes due to movements along the EEC have offset roughly one-and-a-half times as much pollution as corresponding shifts in the EEC. Whether the magnitude of compositional offsets due to movements along the EEC continues into the future depends on the shape of EECs, whereas continued compositional

changes due to shifts in the EEC hand depend entirely on continued tightening of environmental policy.

2.6 Conclusion

Over the past 30 years overall pollution in the US has declined despite increases in total production. Some of this improvement has come from employing cleaner production technologies in cars and factories, but much of it comes as a result of consuming a cleaner mix of goods and services. We do not know, however, whether this cleaner consumption has been a consequence of economy-wide trends such as regulation-induced price changes, or of an underlying and possibly coincidental preference by richer households for cleaner goods. Environmental Engel curves modeling the relationship between income and the pollution-intensity of household consumption and provide a means for disentangling the effects of household income growth from economy-wide changes.

Whether estimated parametrically or non-parametrically, EECs display three key characteristics: they are increasing, concave, and shifting over time. These characteristics allow us to decompose changes in the pollution associated with household consumption into movements along the EEC and shifts in the EEC. The total amount of pollution embodied in household consumption increases with income, but it does so at a decreasing rate. This is captured by movement along a concave EEC. Along a given EEC, the extent to which pollution grows less than income represents compositional shifts in the mix of goods and services that offset increases in pollution due to higher overall consumption. At the same time, the entire EEC shifts down as economy-wide factors drive households to consume less pollution-intensive mixes of goods and services.

Between 1984 and 2002 we find that movements along the EEC account for a 4.17 ton increase in total PM10, holding technology constant using the 1997 TEAM coefficients. This represents a hypothetical increase of 20.4 tons due to pure scale effect combined with an offset of 16.27 tons due to compositional changes in consumption along the EEC. At the same time, the shifting EEC between 1984 and 2002 is alone responsible for a 10.93 ton reduction in PM10 emissions. The pollution offset by downward-shifting EECs is roughly two-thirds of the magnitude of pollution offset due to compositional changes as household move along the EEC.

The fact that compositional shifts in consumption along the EEC appear to have done more to offset the increase in PM10 emissions due to the scale effect does not mean that movements along the EEC are necessarily more effective at reducing pollution overall. The reduction in PM10 due to shifts in the EEC was the result of aggregate changes in prices that resulted from increased environmental regulations, changes in social norms, or other economy-wide factors. There is no reason why that figure can't be larger or smaller in the future depending on public policy. Indeed if no additional policy actions are implemented, the future gains from shifting EECs are likely to be zero.

In the end, decomposing income growth into movements along and shifts in the EEC represents just one aspect of the environmental consequences of economic growth. A large portion of the cleanup in the US comes from changes in technology, changes in the composition of production, and changes in the pollution intensity of US imports and exports. Nevertheless, isolating the consumption-related compositional changes in pollution suggests that composition alone has more than offset the scale effect from household income growth. The degree to which this continues into the future depends both on the level of income growth and on the continued enactment of environmental policy.

CHAPTER III

RURAL PUBLIC WORKS AND CHILDHOOD SCHOOLING INTENSITY

3.1 Introduction

Rural public works programs are used as social safety nets in developing countries around the world. In India, the National Rural Employment Guarantee Act (NREGA) guarantees up to 100 days of temporary employment for any household upon request. Public works constitute a substantial portion of the poverty alleviation strategy in India, and the NREGA is one of the largest public works programs in history (see, for example, Reddy et al., 2010). NREGA expenditures in fiscal year 2010-2011 were roughly nine billion dollars and the program provided emergency employment to 54 million households.⁵¹ In this paper I focus on the NREGA to empirically measure the effects of public works on the schooling intensity of children in participating households.

Although the effects of NREGA on childhood schooling may be secondary to the immediate benefits of poverty alleviation, these effects have important implications. Childhood schooling captures an intergenerational aspect of participation. If additional income allows households to increase schooling intensity, then public works may mitigate the transmission of poverty from parent to child. If participation causes households to decrease schooling, then these programs may work against their own objective by reinforcing the intergenerational transmission of the poverty they are meant to fight. In this case, public works may make it more difficult for poor households to escape poverty in the long run, since children of participating households are more likely to be poor themselves, in turn relying on the program, which then affects the schooling of the next generation.

⁵¹ Ravallion et al., 2013. The average number of days per participating household in fiscal year 2010-2011 was 47. See Drèze and Oldiges (2007) and the Department of Rural Development (2013b).

Previous literature on public works focuses on whether such programs are well targeted and whether they alleviate poverty.⁵² Few studies have examined the spillover effects to other household members, and fewer still concentrate on intergenerational effects. Datt and Ravallion (1994) show that other household members take up abandoned home production when one member joins a public works project, but their study is strictly limited to adult members. Brown, Yohannes, and Webb (1994) focus on children while examining the effects of a public works program in Niger, but do not address schooling effects. Chumacero and Paredes (2002) estimate the effects of an public works program in Chile on young adults and find that the high-school dropout rate increased by 1.1 percent. Last, Afridi, Mukhopadhyay and Sahoo (2012) examine NREGA implementation in several districts of Andhra Pradesh and find that mothers' participation increases schooling, but fathers' participation decreases schooling.

The challenge to estimating the effects of public works on childhood schooling is a potential spurious correlation between those two activities; lower schooling intensity and higher public works participation may appear correlated because they are both driven by the same unobserved factors. The standard solution is to find an instrument or exclusion restriction that is correlated with public works participation, but does not directly affect the schooling decision. Such an instrument is hard to find. Instead, I rely on a new econometric technique introduced by Klein and Vella (2010) that exploits heteroscedasticity to identify treatment effects in the absence of an exclusion restriction.

I model the schooling decision and the public works decision as a triangular system where schooling depends on public works and a set of explanatory variables, and the public works decision itself also depends on the same set of explanatory variables. I find that an

⁵² See, for example, see Dev (1995) and Devereau and Solomon (2006) for reviews of the literature surrounding rural public works programs.

increase in household public works participation leads to a decrease in the likelihood that a child is in school fulltime. This effect persists even after controlling for the endogeneity of public works participation and the schooling decision.

3.2 The National Rural Employment Guarantee Act

Various piecemeal rural public works programs existed in India prior to the NREGA, with the most notable early example being the Maharashtra Employment Guarantee Scheme (EGS) in the 1970's. Starting in 2006, the NREGA institutionalized public works on a national scale by introducing a federal mandate that rural local governments provide working-age citizens up to 100 days of annual employment per household upon request. Multiple individuals within each household are eligible to seek employment through the program, but cumulative household hours cannot exceed the 100-day limit. To participate, a household member must apply for a job card with the Gram Panchayat, the village-level government representative.

Employment under the NREGA is unskilled manual labor that anyone can perform.⁵³ Workers are required to report to the job site, and their compensation must match the prevailing minimum wage.⁵⁴ Manual labor and low wages in the NREGA act as a self-targeting mechanism so there is no means-test for participation. The NREGA also affords other benefits to rural workers, such as unemployment benefits if the guarantee of work is not met within 15 days, and additional travel subsidies if work occurs more than five kilometers from the applicant's village.

Financing for the NREGA comes from the central government, but projects are planned and implemented on the village level. Further, existing public employment programs were subsumed under the financial umbrella of the NREGA as their coverage areas overlapped. A key

⁵³ For example, permissible works include water conservation and harvesting, digging irrigation canals, drought proofing. For a full list of permissible projects, see Department of Rural Development, 2013a.

⁵⁴ For more information on the salient features of the NREGA, see Mahapatra, et al. (2008).

aspect of NREGA is the timing of implementation. The program has universal coverage within eligible areas, but only certain districts were included at certain times. The program was rolled-out in three phases beginning in February 2006. The first phase included 200 districts and gave priority to the most backward areas. For political reasons, some more-developed districts were included, along with at least one district from every state. The second phase began in April 2007 when coverage was expanded by 130 districts. Finally, Phase III began in April 2008 when the remaining 266 districts were included. In total, 596 of the 640 Indian districts are eligible.

3.3 The Effects of Parent Participation on Child Schooling

The effects of public works participation on children's schooling can be separated into two components. The first is a pure income effect. Public works represent an additional outside employment option, and higher earning potential due to NREGA will lead to more consumption of normal goods. To the extent that richer households prefer more schooling for their children, an increase in household income will have a positive effect on schooling intensity.

The second component is a substitution effect. Because the NREGA requires the worker to physically attend a job site, a worker must abandon her current household activities, which creates a vacancy in home production. To the extent that children can substitute for adult labor in the household, adult participation in the NREGA increases the opportunity cost of schooling for children in participating households. This should lead to lower school intensity. The overall effect on schooling depends on the relative magnitude of the income and substitution effects.

To demonstrate that the net effects are unclear, consider the following illustrative model. Suppose each household consists of only two members, an adult and a child, that there is no

saving, and that households care about consumption (c) and the schooling of the child (s).

Households maximize a concave, additively separable utility function:

$$U(c, s) = u(c) + \pi v(s) \quad (14)$$

where π captures the relative weight of utility from schooling relative to utility from consumption.

In addition, suppose there are two states, one in which a public works project is available and pays a fixed wage w , and one in which there is no public works project. Income, and hence consumption, can be generated by allocating resources to household production or to public works, if available. Both members can contribute to household production, but child labor substitutes less than one-for-one with adult labor. Only the adult is eligible for public works employment, and the household can allocate adult labor l_A and child labor l_C within the home such that:

$$\begin{cases} c(l_A, l_C) \leq f(l_A + \phi l_C), \text{ if public works are not available} \\ c(l_A, l_C) \leq w(1 - l_A) + f(l_A + \phi l_C), \text{ if public works are available} \end{cases} \quad (15)$$

$$S(l_C) \equiv (1 - l_C) \quad (16)$$

where f is a concave production function with $0 < \phi < 1$.⁵⁵

Without public works, the adult member devotes all of her time to household production and $l_A = 1$. Let \bar{l}_C represent the allocation of the child's time to production under this scenario, which is implicitly defined by the first order condition:

$$u'[f(1 + \phi \bar{l}_C)] * [\phi f'(1 + \phi \bar{l}_C)] = v'(1 - \bar{l}_C) \quad (17)$$

This condition balances the marginal utility from additional consumption due to child labor on the left side with the marginal utility of schooling on the right.

⁵⁵ The function f can be interpreted generally to include any income generating activity in which child labor can substitute for adult labor, as well as utility-enhancing activities such as childcare or household maintenance.

When public works are available, an interior solution involves $l_A \leq 1$ and the adult worker chooses an allocation such that the marginal benefit of working in public works equals the marginal benefit of household production:

$$f'(l_A + \phi \tilde{l}_C) = w \quad (18)$$

Since children are not eligible for public works employment, the allocation decision for child labor is similar to equation 17, but both l_A and l_C can vary:

$$u'[f(l_A + \phi \tilde{l}_C) + w(1 - l_A)] * [\phi f'(l_A + \phi \tilde{l}_C)] = v'(1 - \tilde{l}_C) \quad (19)$$

Equations 18 and 19 implicitly define \tilde{l}_C and l_A in the presence of public works by equating the marginal benefits of each individual's activities (production or schooling for the child and production or public works for the adult). The relative magnitude of \tilde{l}_C in equations 17 and \tilde{l}_C in equations 18 and 19 depends on the functional forms of u and f and the value of ϕ .

3.4 Econometric Model

To measure the effects of public works participation on schooling intensity of workers' children, I estimate a reduced form model where the hours of schooling for a child depend on a series of personal and household characteristics, including the number of hours that her family participates in public works. In addition, the parents' participation decision also depends on household characteristics. Together, these two factors lead to a triangular structure:

$$s_i = \beta x_i + \delta pw_i + u_i \quad (20)$$

$$pw_i = \gamma z_i + v_i \quad (21)$$

where s_i denotes the average school intensity for child i , pw_i denotes the average public work intensity for adults in her household, and x_i and z_i capture personal and household characteristics that affect participation and schooling.

When unobserved factors affect both schooling and participation, the correlation between u_i and v_i induces bias in the OLS estimates of equation 20. If there are valide instruments available to use in equation 21 so that z_i contains some variables not included in x_i , the system can be estimated using standard instrumental variable or sample selection techniques. In the case of public works and schooling, it is unlikely that any factors affect the participation decision without also affecting the schooling decision. Instead, I rely on a new technique introduced by Klein and Vella (2010) that identifies the coefficients in equation 20 and 21 by exploiting heteroscedasticity in the error terms.

The Klein and Vella (2010) estimator is a refinement of a control function approach. Intuitively, the standard control function approach takes the bias in equations 20 and 21 and recasts it as an omitted variable bias. To see this, express the error in equation 20 as a linear function of the error in equation 21:

$$u_i = \rho v_i + \varepsilon_i \quad (22)$$

Then by substituting this back into equation 20 we have:

$$s_i = \beta x_i + \delta p w_i + \rho v_i + \varepsilon_i \quad (23)$$

where ε_i is independent from all other explanatory variables. With the addition of the control function term ρv_i , the model in equation 23 is purged of endogeneity. The value of v_i can be approximated by its estimate \hat{v}_i based on equation 21, and the coefficient ρ is estimated along with the other coefficients in equation 23.⁵⁶

A limitation of the standard control function approach is that \hat{v}_i is a linear function of $p w_i$ and z_i . When $z_i = x_i$, introducing \hat{v}_i into equation 23 results in perfect multicollinearity between \hat{v}_i and the other explanatory variables. Klein and Vella (2010) derive a control function

⁵⁶ See Heckman, LaLond, and Smith (1999) for a technical discussion of control function techniques.

technique that includes \hat{v}_i along with a non-linear function of x_i based on heteroscedasticity in the error terms in equations 20 and 21. To accomplish this, they assume that both u_i and v_i are heteroscedastic multiples of homoscedastic components u_i^* and v_i^* , and that those components have a constant correlation. Formally:

$$u_i = S_u(x_i)u_i^* \text{ and } v_i = S_v(x_i)v_i^* \quad (24)$$

$$E[u_i^*v_i^*] = E[u_i^*v_i^*|x_i] = \rho_0 \quad (25)$$

where $S_u^2(x_i)$ and $S_v^2(x_i)$ denote the conditional variance functions for u_i and v_i , and u_i^* and v_i^* are zero-mean homoscedastic error terms. They also assume that $S_u(x_i)$ and $S_v(x_i)$ do not have a constant ratio. In a manner analogous to equation 22, let:

$$u_i = A(x_i)v_i + \varepsilon_i \quad (26)$$

where $A(x_i) = \rho_0 S_u(x_i)/S_v(x_i)$ and $\rho_0 = [cov(u_i, v_i|x_i)/(S_u(x_i)S_v(x_i))]$. Substituting equation 26 into equation 20 yields:

$$s_i = \alpha + \beta x_i + \delta p w_i + A(x_i)v_i + \varepsilon_i \quad (27)$$

In this case, ε is uncorrelated with v , $p w$, and x . Further, since $A(x_i)$ is no longer a linear combination of other variables, equation 27 does not contain perfect multicollinearity. In this way, Klein and Vella (2010) exploit the assumption embodied in equations 24 and 25 (heteroscedasticity and a constant correlation between the homoscedastic portions of the error terms) to allow for identification of the coefficients in equations 20 when no valid instruments are available.

A convenient and intuitive interpretation of the homoscedastic portion of the error terms in the context of public works and schooling is the unobserved ability of adults and children. Although we can observe many factors that affect schooling and public works participation, each individual has some innate ability that we are not able to observe. The decision whether to

participate in school or public works depends on each individual's privately-known ability. If the incremental effect of ability on the participation decision is correlated with other observable factors (for example, if the effect of ability is correlated with age), and if ability can be passed genetically from parent to child, then interpreting u_i^* and v_i^* as unobserved ability (rename them a_i^* and A_i^*) satisfies the requirements embodied in equations 24 and 25.

The first of these assumptions (embodied in equation 25) is that the effect of ability depends on other observable characteristics. There are many reasons why this may be the case for public works participation. For example, participation rates vary widely by district and the impact of ability may depend on location characteristics such as public infrastructure, attitudes towards public works, bureaucracy, or local corruption. Household characteristics also determine the marginal impact of ability on participation. The variance is likely to depend on self-employment, land holdings, and whether the household earns wage or agricultural income. For example, agricultural households may be more susceptible to random shocks such as drought or crop failure. The marginal impact of ability may also be correlated with individual factors such as education and age; manual labor is more difficult for elderly workers than for young workers and higher-educated workers may have better access to alternative employment.

The same variables that lead to heteroscedasticity in the parent participation decision can also cause heteroscedasticity in the schooling decision. The effect of ability on schooling will be different for agricultural households than for wage-earning households, and will also be different for self-employed households. In addition, the marginal effect of ability on schooling may also be correlated with child-specific factors. Older students can easily substitute for adult labor whereas younger students cannot, so ability will affect the schooling decision differently for different ages.

The assumption in equation 25 requires a constant correlation between the error terms in equation 20 and 21. This is also plausible when I interpret the error term as capturing unobserved ability. For example, suppose the error structure satisfies the relationship:

$$u_i = S_u(x_i)a_i^* \quad (28)$$

$$v_i = S_v(x_i)A_i^* \quad (29)$$

where a_i is child's ability and A_i is the ability of her parent. In this case, the constant correlation assumption is satisfied if I assume that ability is inherited from parent to child. Suppose ability passes down through generations in a stochastic process where the child receives her parent's ability, multiplied by a random shock τ_i :

$$a_i^* = \tau_i \cdot A_i^* \quad (30)$$

If τ does not depend on x , as would be the case if ability is genetic, the error structure in 20 and 21 satisfies the constant-correlation assumption.

3.5 Estimation

To estimate the coefficients in equations 20 and 21 I use a three-step approach to implement a parameterized version of the Klein and Vella estimator.⁵⁷ I also make the additional simplifying assumption that the unobserved error in equation 20 is homoscedastic, implying that $S_u(x_i)$ in equation 24 and 28 is constant, S_u . The ultimate goal is to estimate the effects of public works participation by substituting the functional form of $A(x_i)$ into equation 27 and running an OLS:

$$s_i = \alpha + \beta x_i + \delta p w_i + \rho_0 \left(\frac{S_u}{S_v(x_i)} \hat{v}_i \right) + \varepsilon_i \quad (31)$$

⁵⁷ Examples of other studies that have implemented the Klein and Vella estimator include Farré, Klein, and Vella (2012), Millimet and Roy (2012), Millimet and Tchernis (2012), and Schroeder (2010).

To estimate equation 31 in the final stage, I first need to calculate \widehat{v}_i and $\widehat{S}_v(x_i)$ in the first and second stages.

To begin, I run a first-stage ordinary least squares regression to estimate the coefficients in equation 21 and use those results to predict \widehat{v}_i . Next, I assume that $S_v(x_i)$ has an exponential form and run a second-stage non-linear least squares regression to estimate the equation:

$$S_v(x_i) = \sqrt{\exp(x_i \theta_v)} \quad (32)$$

I use the results of this second-stage regression to predict $\widehat{S}_v(x_i)$. Finally, by substituting \widehat{v}_i and $\widehat{S}_v(x_i)$ into equation 31, I run an OLS regression that jointly estimates β , δ , and the product ρS_u . The resulting coefficient δ from the third stage is an unbiased estimate of the effect of public works on schooling.

3.6 Data and Results

Information on individuals and families that participate in rural public works in India is obtained from household surveys conducted by the National Sample Survey Organization (NSS), a division of the Ministry of Statistics. The NSS conducts annual surveys that capture a nationally representative sample of the Indian population. On an irregular basis, the NSS adds a module to their survey that elicits information on employment and unemployment activity. I rely on three rounds of NSS survey data that include the employment module, spanning 2004 through 2006. These data capture the initial rollout of the NREGA, as well as a small amount of public works participation in the piecemeal programs that existed before NREGA implementation.

The employment modules of the NSS surveys provide detailed information about time-use over the past seven days for all household members. This includes working hours for adults, along with specific indicators for public works, and school attendance for children. Time-use is

measured in “intensity units,” where an individual can participate in up to two tasks per day for a total weekly intensity limit of fourteen. I characterize children as being in school full time if they indicate a schooling intensity of at least ten, meaning their primary task was school for at least five full days per week. Table 1 shows summary statistics for the estimation sample.

To begin, I regress an indicator for fulltime schooling status on the intensity of household participation in public works using ordinary least squares. This naïve regression does not control for the endogeneity of public works participation. The first column of Table 2 shows the results of this regression. I restrict the sample in this case to school-aged children from households that participate in public works, so this regression measures the effects of public works participation intensity conditional on being from a household that participates. In this case, schooling intensity is strongly correlated with adult participation in public works (-0.002). Several other factors are also significantly correlated, including age (0.204), an indicator for boys (0.112), self-employed households (-0.053), education and age of the household head (0.040 and 0.045), household size (0.045) and whether the household is in a scheduled caste or tribe (-0.233). These correlations are generally intuitive.

For comparison, the second column of Table 2 shows a similar regression with monthly per-capita household consumption expenditure on the left hand side. In this case, average household public works intensity is not significantly correlated with consumption expenditure. Agricultural households and scheduled castes or tribes have lower consumption (-167.63 and -243.47), while consumption increases with education of the household head, with household size, and for self-employed households (62.28, 246.54, and 304.10, respectively). The regression in column 2 also produces intuitive correlations, which corroborate the correlations in the first column.

The ordinary least squares regression in column 1 does not account for the endogeneity of the public works decision so the coefficient estimates may be biased. To correct for this I implement the multi-stage Klein and Vella approach. The first stage is to estimate the public works participation decision. Column 1 in Table 3 presents OLS estimates of equation 21 using the same sample and explanatory variables as Table 2. The intensity of public works participation is significantly correlated with many of the explanatory variables; larger households have higher participation (9.09), but less so if they consist of more children (-5.73), and agricultural, self-employed, and better educated households participate with lower intensity (-15.74, 8.94, and -1.26, respectively). Both Breusch-Pagan and White tests for heteroscedasticity strongly rejecting the null hypothesis of homoscedasticity (test statistic of 252.9 and 148.6, respectively).

In the second stage I use the residuals from the first column to estimate $S_v(x)$ by assuming the variance of the residuals follow an exponential form. Column 2 of Table 3 shows the results of a non-linear least squares regression with $\ln(\hat{v}_i^2)$ on the left-hand side and all the same explanatory variables on the right-hand side.⁵⁸ Several of the explanatory variables are strongly correlated with the standard deviation of the error from the participation decision. The variance is bigger for larger households (0.045), agricultural households (0.07), and for households with older heads (0.018), which is consistent with the NREGA rules that limit eligibility to workers over the age of 18. The variance is decreasing in education of head (-0.022) and lower for self-employed households (0.067), which is consistent with fewer outside employment options for those groups.

⁵⁸ Standard errors are calculated based on 1000 bootstrap repetitions.

Using the results of columns 1 and 2 of Table 3, I substitute \hat{v}_i and $\widehat{S_v(x_i)}$ into equation 31. The third column of Table 2 shows the results of a linear probability estimate of the coefficients in equation 31. Even after correcting for endogeneity, the coefficient on participation is negative and statistically significant (-0.007). The magnitude of the coefficient increases after controlling for endogeneity and the difference is significant at the ten percent level; children in households with higher participation in public works projects attend school with lower intensity.

The coefficient on the control function term in column 3 measures the effect of the endogenous portion of the error term in equations 20 and 31. In this case, the coefficient is positive (0.014) and statistically significant at the ten percent level. The sign of this coefficient is determined by the covariance between u_i^* and v_i^* , and the positive sign is consistent with the interpretation of unobserved ability that is passed from parent to child. A statistically significant coefficient on the control function term also supports the hypothesis that the public works decision is endogenous to the schooling decision and that the control function approach has captured at least some of that endogeneity. Even after correcting for the spurious correlation between public works participation and schooling, it appears as though increased public works intensity results in lower schooling intensity among participating households.

Despite statistically significant results using a control function to correct for bias from spurious correlation, these results should be interpreted with caution. The regressions in Table 2 only included households that participated in public works during the survey period. As such, I am only measuring the effects of increasing public works participation on schooling for those households that already participated in public works; I am measuring the effect on the intensive margin. One reason for this is that public works take-up rates in the NSS data are extremely

small. For example, there are roughly 310,000 households surveyed in the 60-62 rounds of the NSS. Of these, 1065 (0.3 percent) indicate participation in public works. This may be because widespread public works participation in India began with the NREGA implementation, and the NSS data used in this paper extend only several months after NREGA rollout. In addition, I have simplified the estimation relative to the full Klein and Vella approach by parameterizing the model and assuming homoscedastic errors in the participation decision. The significance of the negative coefficient in Table 2 is not robust to these simplifying assumptions.

3.7 Conclusion

The effects of public works participation on schooling intensity of children in workers' households may have important intergenerational implications for evaluating public works as long-term anti-poverty programs. Participation appears to negatively affect education, so public works projects may create self-perpetuating cycles in which the children of participating households are more likely to require public works in the future.

Focusing on the National Rural Employment Guarantee Act in India, I find that public works participation is strongly correlated with lower schooling intensity for the children in workers' households. This negative effect persists even after controlling for the endogeneity of public works. Each additional unit of intensity of public works is associated with a 0.007 decrease in the probability that a student is in school fulltime. Although significant, this effect is small in magnitude. Nevertheless, it is important to consider the education spillover effects of the NREGA when evaluating the poverty-alleviation benefits of rural public works in India.

APPENDIX A
TABLES AND FIGURES FOR CHAPTER I

Table A.1
Selected Recent Age-VSL Studies

Study	Age-VSL relationship	Framework	Relevant Age Groups	VSL Range (millions, 2009\$)
Alberini, et al. (2004)	Flat, modest decline for 70+	Contingent valuation	40 to 75 (in Canada) or 80 (in US)	1.92 to 2.16
Kneisner, Viscusi, and Ziliak (2006)	Inverted-U, peak around 50	Hedonic wages	18 to 65	8.42 to 19.92
Aldy and Viscusi (2008)	Inverted-U, peak around 46	Hedonic wages	18 to 62	4.22 to 9.7
Viscusi and Aldy (2007)	Inverted-U, peak around 35-44	Hedonic wages	18 to 62	7.30 to 15.35
Evans and Schaur (2010)	Decreasing	Quantile wages	50 to 60	<0 to 19.34
Aldy and Viscusi (2003)	Inverted-U	Hedonic wages	18 to 62	4.0 to 10.42
Anderson (2007)	Decreasing (weekly)	Contingent Valuation (road safety)	17 to 74	
Smith, et al. (2004)	Increasing	Hedonic wages	26 to 65	<0 to 18.29
Krupnick, et al. (2002)	Flat, modest decline for 70+	Contingent valuation	40 to 75	1.24 (average)
Evans and Smith (2008)	Inverted-U	Hedonic wages	51 to 65	13.41 to 34.27
DeShazo and Cameron (2004)	Inverted-U (insignificant)	Contingent valuation	25 to 65	0.42 to 4.75
Anderson (2008)	None	Hedonic vehicle prices	18 to 74	2.07 (average)
Johannesson, et al. (1997)	Inverted-U	Contingent valuation	18 to 74	4.83 to 7.48
Persson, et al. (2001)	Inverted-U	Contingent valuation	18 to 74	3.48 (average)

Table A.2
Summary Statistics For
the Full Sample

	Mean	SD
Age	45.08	16.01
Income (2009\$)	56,831	29,642
Deaths per Make/Model/Vintage	13.78	12.58
VMT (miles per year)	12,021	2,360
Deaths per 100 Million Miles	1.15	1.29
Used Price (2009\$)	14,364	7,034
Annualized Cost (2009\$)	1,741	1,128
Engine Displacement (liters)	3.10	0.97
Cylinders	5.50	1.27
Passenger Volume (cubic feet)	96.29	7.81
Gross Vehicle Weight (1000 lbs.)	3.50	0.79
Share of Luxury Vehicles	0.07	0.25
Share of Imported Vehicles	0.41	0.49
Share of Vintage 2003-2005	0.33	0.47
Share of Vintage 2006-2008	0.29	0.45
Shares by Class		
Two Door	0.06	0.24
Four Door	0.54	0.50
Sport	0.03	0.16
Minivan	0.16	0.37
Pickup	0.01	0.11
Small SUV	0.08	0.27
Large SUV	0.12	0.32
Sample Size	14,321	-

Notes: All values except sample size are calculated using sample weights. Engine displacement measures the volume swept by the pistons within the engine cylinders. Luxury brands include: Acura, Audi, BMW, Cadillac, Infiniti, Jaguar, Land Rover, Lexus, Mercedes-Benz, and Porsche.

Table A.3
Vehicle Choice Model and VSL Estimates For Seven Age Cohorts

	<u>Pooled Risk</u>	<u>Age-Specific Risk</u>	<u>Behavior-Adjusted Risk</u>
	(1)	(2)	(3)
Risk * I(age 18-24)	-5.42* (2.16)	-0.76* (0.37)	-2.40** (0.41)
Risk * I(age 25-34)	-13.00** (3.01)	-1.66† (0.97)	-2.79† (1.49)
Risk * I(age 35-44)	-23.92** (4.21)	-3.81* (1.65)	-4.28* (2.01)
Risk * I(age 45-54)	-17.84** (2.87)	-5.74** (1.46)	-5.75** (1.66)
Risk * I(age 55-64)	-27.38** (4.42)	-11.37** (2.33)	-11.23** (2.56)
Risk * I(age 65-74)	-32.28** (10.92)	-6.72** (1.70)	-7.54** (1.88)
Risk * I(age 75-85)	-29.51** (6.30)	-4.35** (0.91)	-4.59** (1.13)
Cost * I(age 18-24)	-8.34** (2.28)	-4.95** (1.65)	-6.41** (2.01)
Cost * I(age 25-34)	-9.90** (2.01)	-10.97** (2.05)	-10.55** (2.11)
Cost * I(age 35-44)	-3.72** (1.11)	-4.61** (1.36)	-4.38** (1.45)
Cost * I(age 45-54)	-2.34 (1.52)	-4.13† (2.22)	-3.66 (2.46)
Cost * I(age 55-64)	-1.45† (0.78)	-2.90** (1.09)	-3.62** (0.99)
Cost * I(age 65-74)	-3.40** (1.18)	-4.88** (1.12)	-4.68** (1.11)
Cost * I(age 75-85)	-3.06** (1.03)	-4.03** (1.04)	-3.49** (0.93)

Continued on next page...

Table A.3 Continued

	Pooled Risk	Age-Specific Risk	Behavior-Adjusted Risk
	(1)	(2)	(3)
Engine Displacement	0.16† (0.09)	-0.03 (0.10)	-0.02 (0.11)
Cylinders	-0.10 (0.06)	-0.08† (0.06)	-0.13* (0.06)
Cylinders*Pickup		0.11** (0.04)	-0.12** (0.04)
MPG	0.33** (0.06)	0.49** (0.06)	0.52** (0.07)
Width	0.13 (0.09)	0.33* (0.10)	0.24* (0.10)
Passenger Volume	0.14* (0.05)	0.06 (0.06)	0.09 (0.06)
Luxury Brand	-0.62** (0.15)	-0.70** (0.18)	-0.70** (0.15)
Luxury*(Income>75%ile)	0.93** (0.19)	0.89** (0.18)	0.88** (0.19)
Gross Vehicle Weight	-0.08 (0.07)	-0.23* (0.12)	-0.18 (0.11)
Imported	-0.34** (0.07)	-0.27** (0.08)	-0.26** (0.08)
Model Year 2003-2005	0.36** (0.07)	0.19* (0.08)	0.17** (0.07)
Model Year 2006-2008	1.19** (0.08)	0.59** (0.09)	0.55** (0.09)
Class Constants	Yes	Yes	Yes
n	8503	8486	7638
Likelihood ratio index	0.11	0.09	0.07
VSL Estimates (million 2009\$)			
Ages 18-24	3.82* (1.66)	0.92* (0.44)	2.38* (0.82)
Ages 25-34	7.90** (2.40)	0.87 (0.53)	1.54† (0.92)
Ages 35-44	40.27** (12.61)	4.94* (2.39)	5.95† (3.19)
Ages 45-54	60.12 (80.22)	10.53† (6.06)	13.89 (10.26)
Ages 55-64	130.37 (142.78)	26.32* (11.43)	19.23** (6.67)
Ages 65-74	59.27* (23.66)	8.27** (2.44)	9.63** (2.72)
Ages 75-85	62.52* (25.29)	6.70** (2.22)	8.15** (3.02)

Note: Significant at the one (**), five (*), or 10 percent (†) level. Standard errors are calculated based on 1000 bootstrap repetitions. Because of the non-linear nature of the VSL (a ratio), outliers are trimmed from the bootstrap calculation. Risk and cost are measured in expected deaths per year and cost per year, respectively. Pooled risk is calculated by dividing the total number of deaths by total VMT for each vehicle, and then per-mile risk is multiplied by each household's annual VMT. Age-adjusted risk is calculated in a similar manner, but separately for each age cohort. Behavior-adjusted risk uses the results of Table A.5 to reweight the measured riskiness of each vehicle within each individual's choice set. Engine displacement measures the volume swept by the pistons within the engine cylinders. Luxury brands include Acura, Audi, BMW, Cadillac, Infiniti, Jaguar, Land Rover, Lexus, Mercedes-Benz, and Porsche.

Table A.4
Testing Differences in VSL Between Age Cohorts
Based on Column 3 of Table A.3

	Age 25-34	Age 35-44	Age 45-54	Age 55-64	Age 65-74	Age 75-85
	(1)	(2)	(3)	(4)	(5)	(6)
Age 18-24	-0.85** (0.15)	3.56 (2.38)	11.51 (9.46)	16.85** (5.85)	7.25** (1.90)	5.77** (2.21)
Age 25-34		4.41† (2.28)	12.35 (9.36)	17.70** (5.75)	8.10** (1.81)	6.61** (2.11)
Age 35-44			7.94 (7.16)	13.29** (3.48)	3.69** (0.53)	2.21** (0.24)
Age 45-54				5.35 (4.00)	-4.26 (7.67)	-5.74 (7.36)
Age 55-64					-9.60* (3.96)	-11.08* (3.65)
Age 65-74						-1.48** (0.34)


Note: Significant at the one (**), five (*), or 10 percent (†) level. Standard errors are calculated based on 1000 bootstrap repetitions. Because of the non-linear nature of the VSL (a ratio), outliers are trimmed from the bootstrap calculation. Differences are based on VSL estimates from column 3 of Table A.3. Values along the diagonal correspond to the difference between the VSL of a given cohort and the VSL of the next youngest cohort.

Table A.5
Driver Effects: Proportional Increases in
Fatalities Relative to Compact Passenger Cars

Vehicle Class	Alpha
Heavy Car	1.11 (0.07)
Light Car	0.78** (0.06)
Medium Car	0.99 (0.04)
Mini Car	1.29 (0.35)
Minivan	0.64** (0.04)
Pickup	2.18** (0.09)
Sports Car	2.35** (0.18)
Large SUV	1.41** (0.07)
Small SUV	0.96 (0.05)
Full-Sized Van	2.17** (0.13)
<hr/>	
n	10,382

Note: **Significantly different from one, at the one percent level. Figures are expressed as incident rate ratios, which represent the proportional increases in fatalities due to driving behavior relative to compact passenger cars. I assume that once a vehicle is tested by NHTSA, all subsequent model years have the same results until the model is re-tested.

Figure A.1
Selected Age-VSL Results in Recent Literature



Note: Solid lines are revealed preference studies; dotted lines are contingent valuation studies. All values have been adjusted to 2009 dollars using the CPI-U. The VSL estimates for each age cohort are plotted as single values at the midpoint of each age range. Smoothed connecting lines are added for visual clarity using Bezier splines.

Figure A.2
Highest and Lowest Number of Fatalities per Mile




Figure A.3
Average Disutility from Risk and Cost Using Age-Adjusted Risk Measures

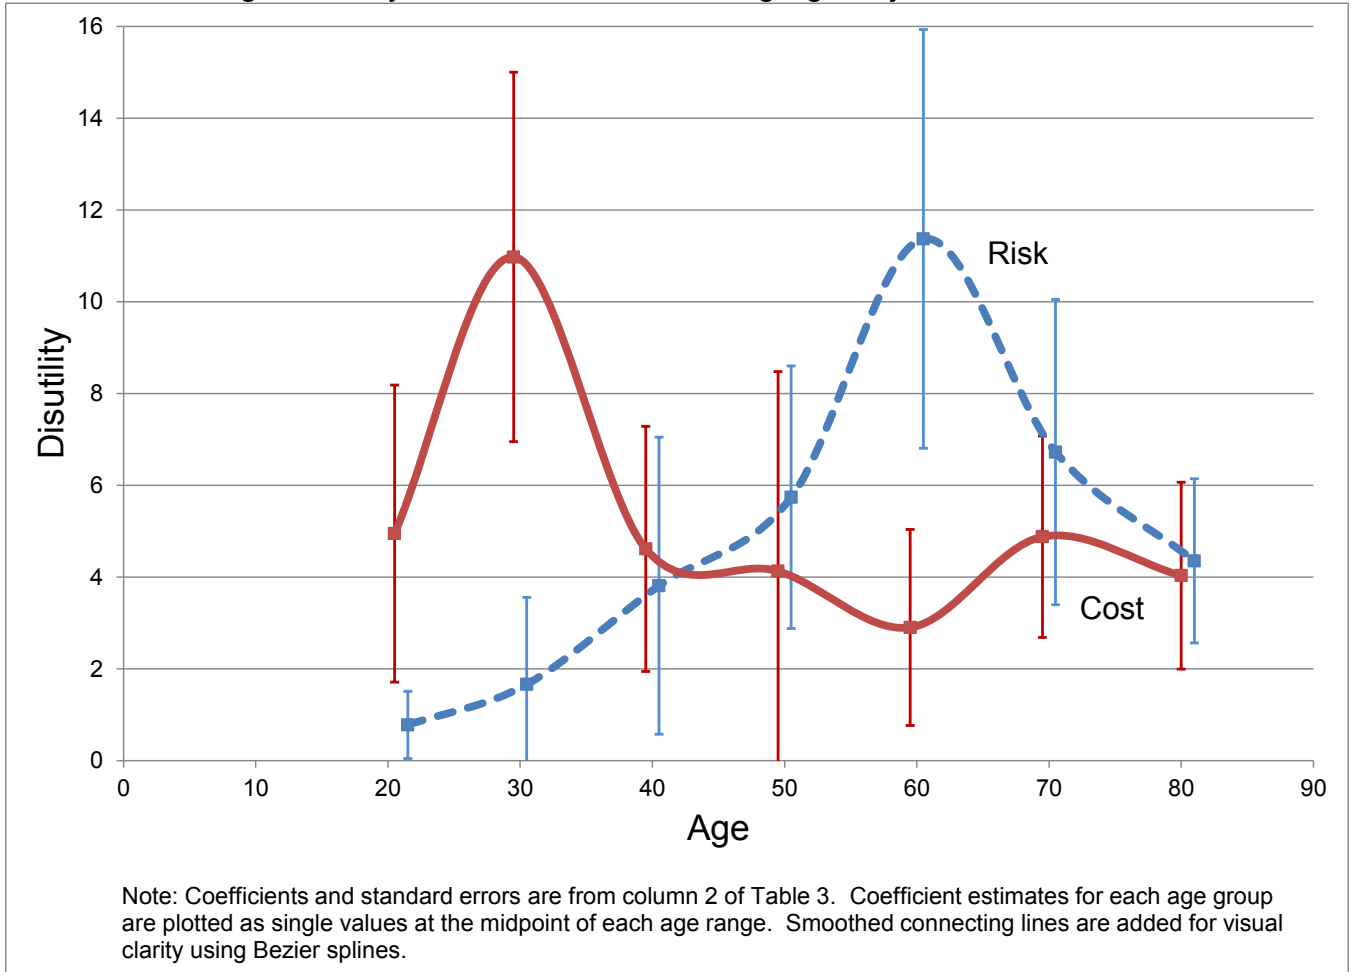


Figure A.4
Average VSL by Age Group Using Age-Specific Risk Measures




Figure A.5
 Average VSL by Age Group Using Behavior-Adjusted Age-Specific Risk Measures

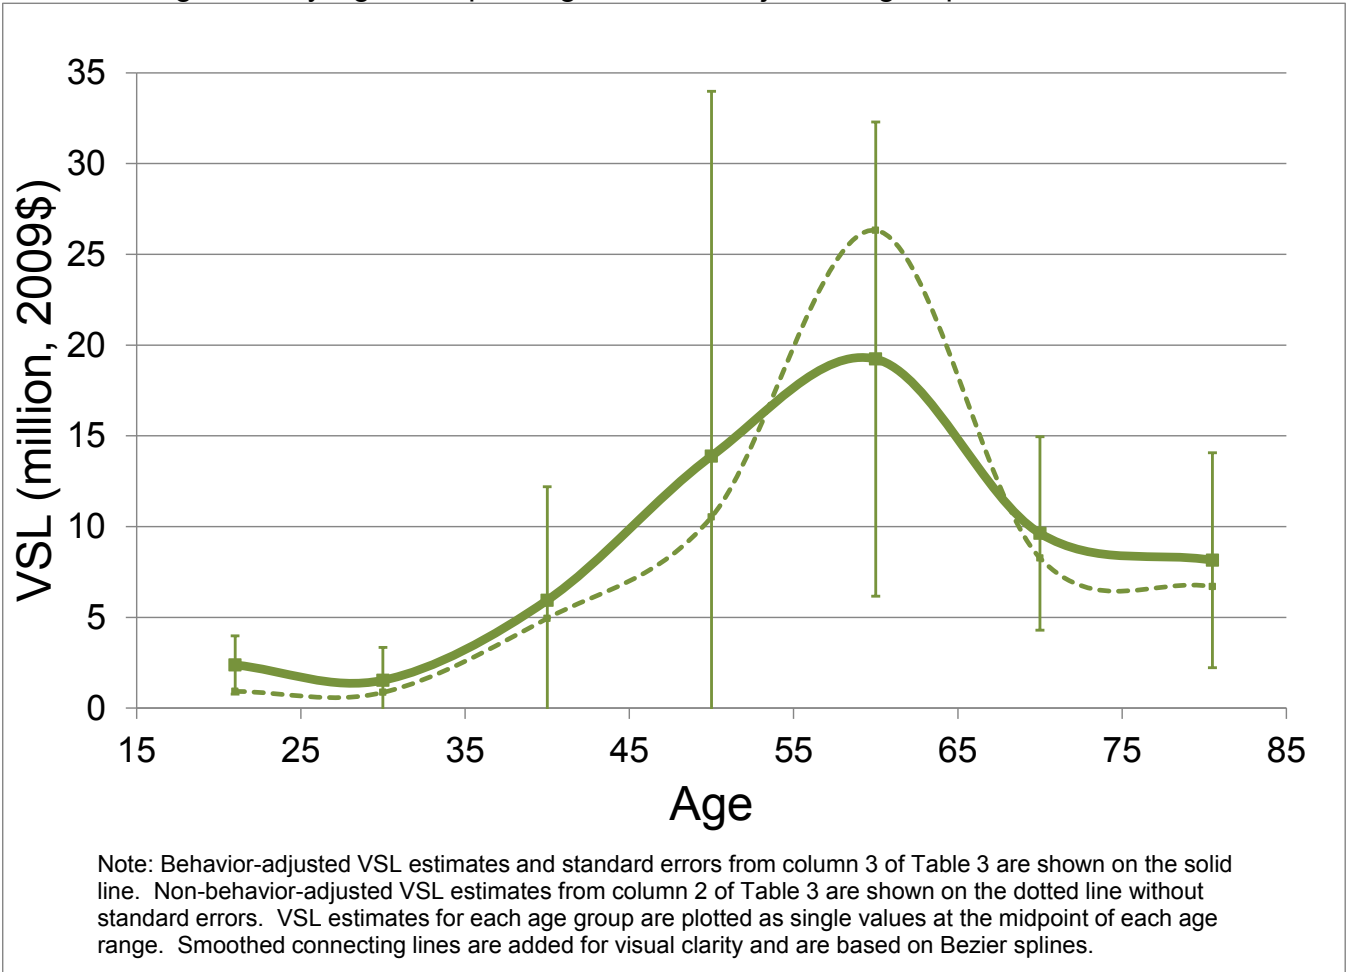


Figure A.6

Average Disutility from Risk and Cost Using Behavior-Adjusted Risk Measure


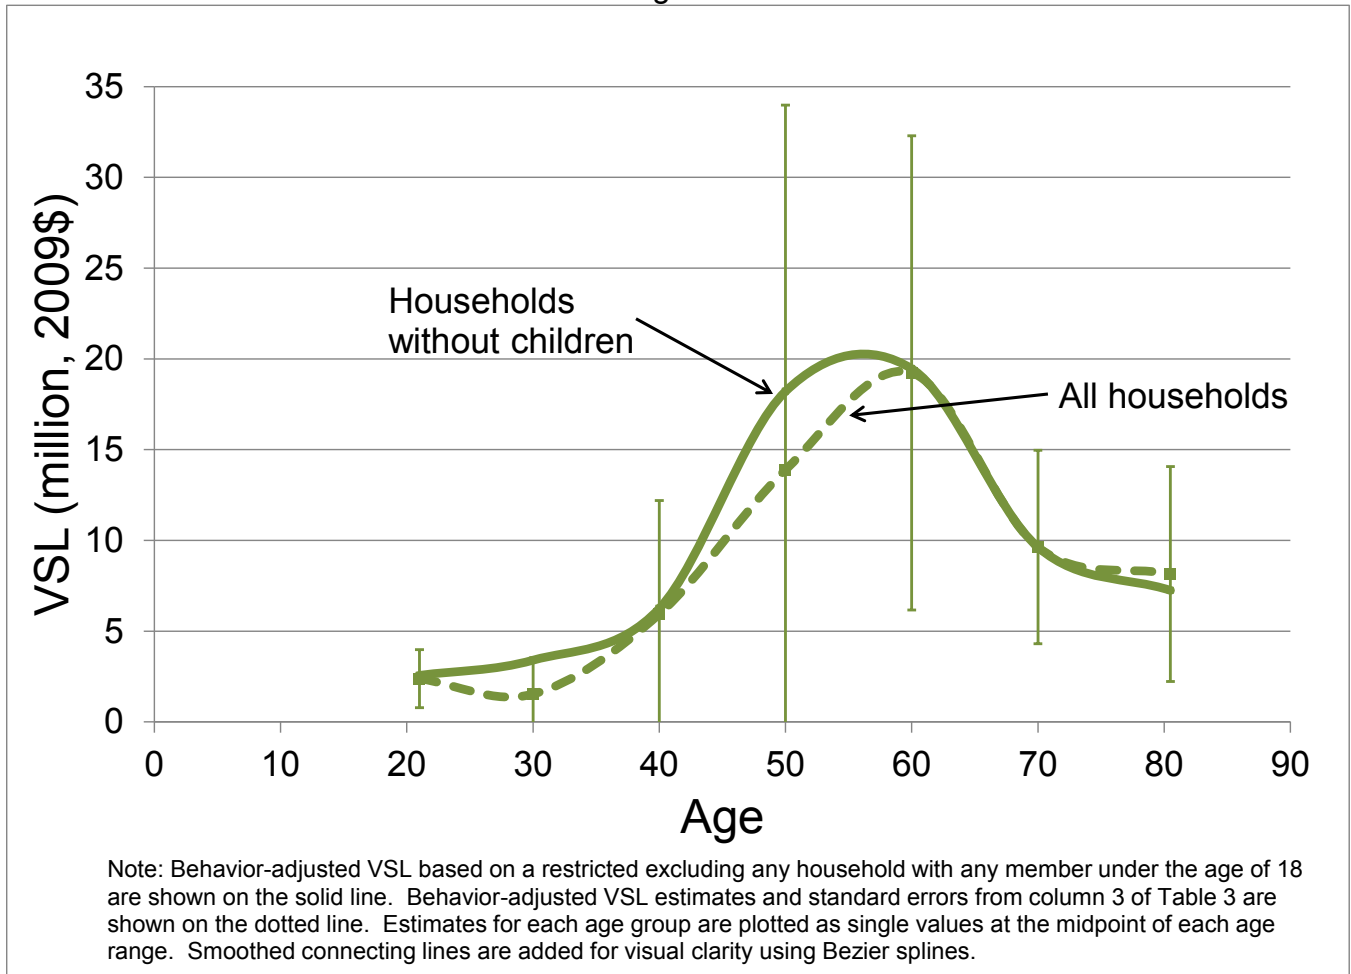


Figure A.7
VSL Estimates Excluding Households with Children



APPENDIX B

TABLES AND FIGURES FOR CHAPTER II

Table B.1
Average Values for Selected Variables
1984 and 2002

Variable	Cross-Section		Difference
	1984	2002	
Total PM10 (tons, 1997 technology)	155.1 (1.74)	145.6 (1.24)	-9.48 (2.14)
Income, (ten thousand \$1997)	4.1 (0.07)	4.6 (0.07)	0.54 (0.10)
Family Size	2.7 (0.03)	2.6 (0.02)	-0.17 (0.04)
Age of Head	46.7 (0.40)	48.4 (0.30)	1.67 (0.50)
Head is Married	0.6 (0.01)	0.5 (0.01)	-0.08 (0.01)
<u>Race of Head (fraction of population)</u>			
White	0.87 (0.007)	0.84 (0.006)	(0.03) (0.010)
Black	0.10 (0.006)	0.11 (0.006)	0.01 (0.009)
Asian or Pacific Is.	0.01 (0.003)	0.04 (0.003)	0.03 (0.004)
Other	0.02 (0.003)	0.01 (0.002)	-0.01 (0.004)
<u>Education of Head (fraction of population)</u>			
Elementary Only	0.27 (0.010)	0.16 (0.006)	-0.11 (0.01)
High School	0.32 (0.010)	0.27 (0.008)	-0.04 (0.01)
Some College	0.19 (0.009)	0.30 (0.008)	0.11 (0.01)
College	0.12 (0.007)	0.17 (0.006)	0.05 (0.01)
More than College	0.11 (0.007)	0.10 (0.005)	-0.01 (0.01)
<u>Region (fraction of population)</u>			
Northeast	0.18 (0.008)	0.19 (0.007)	0.01 (0.01)
Midwest	0.22 (0.009)	0.23 (0.007)	0.01 (0.01)
South	0.27 (0.010)	0.36 (0.008)	0.09 (0.01)
West	0.17 (0.008)	0.21 (0.007)	0.04 (0.01)
Rural	0.16 (0.009)	0.13 (0.006)	-0.04 (0.01)
Total Households	88,312,263	110,554,781	
Observations	3,187	4,363	

Table B.2
Parametric Environmental Engel Curves
1984, 1993, and 2002

Dependent Variable:	1984		1993	2002	Change between
	(1)	(2)	(3)	(4)	1984 and 2002
Total PM10 per household					(5)
Income (ten thousand dollars)	19.04*** (1.309)	10.28*** (1.214)	9.462*** (1.973)	8.280*** (0.550)	-1.998 (1.332)
Income squared	-0.388*** (0.0838)	-0.125* (0.0706)	-0.118 (0.118)	-0.0885*** (0.0248)	0.0369 (0.0748)
Family size		39.10*** (2.938)	37.37*** (4.656)	38.71*** (2.645)	-0.395 (3.953)
Family squared		-1.930*** (0.375)	-2.010*** (0.703)	-2.362*** (0.351)	-0.432 (0.514)
Age		3.430*** (0.399)	2.569*** (0.472)	1.905*** (0.324)	-1.525*** (0.514)
Age squared		-0.0292*** (0.00383)	-0.0211*** (0.00438)	-0.0159*** (0.00299)	0.0133*** (0.00486)
Married		8.585*** (2.940)	7.928** (3.559)	6.080** (2.554)	-2.505 (3.895)
Race (White omitted)					
Black		-26.88*** (3.682)	-17.76*** (3.737)	-20.89*** (2.415)	5.986 (4.402)
Asian		6.455 (17.95)	-6.903 (10.97)	-10.43* (5.547)	-16.89 (18.78)
Other		2.177 (11.13)	-6.223 (15.58)	-7.642 (6.386)	-9.820 (12.83)
Education (< high school omitted)					
High school		9.423*** (3.281)	6.574** (3.215)	-0.132 (2.514)	-9.554** (4.133)
Some college		9.672** (3.860)	15.96*** (4.131)	3.627 (2.639)	-6.045 (4.675)
College		18.10*** (4.917)	16.50*** (4.722)	9.922*** (3.317)	-8.182 (5.930)
Graduate		24.47*** (5.288)	19.93*** (4.908)	15.71*** (3.977)	-8.762 (6.615)

Continued on next page...

Table B.2 Continued

	1984		1993	2002	Change between 1984 and 2002
	(1)	(2)	(3)	(4)	(5)
Region (Northeast omitted)					
Midwest	-15.44***	-15.44***	-26.13***	-8.783***	6.658
	(3.533)	(3.533)	(3.646)	(2.873)	(4.553)
South	-11.62***	-11.62***	-17.72***	-6.598**	5.025
	(3.437)	(3.437)	(3.619)	(2.741)	(4.396)
West	-8.046**	-8.046**	-7.729*	4.003	12.05**
	(3.878)	(3.878)	(4.095)	(3.068)	(4.944)
Rural	-28.56***	-28.56***	-26.76***	-9.922***	18.63***
	(3.844)	(3.844)	(4.047)	(2.443)	(4.553)
Constant	87.90***	-58.59***	-27.18**	-18.40**	40.19***
	(3.685)	(10.82)	(12.61)	(9.115)	(14.14)
Observations	3,185	3,185	3,203	4,363	
R-squared	0.317	0.581	0.502	0.513	

Note: Total household pollution is calculated by multiplying itemized household consumption with the pollution intensity of production for each type of good and summing for each household. Income is measured in tens of thousands of 1997 dollars. Coefficients are significant at the one (***), five (**), or ten (*) percent level.

Table B.3
Oaxaca-Blinder Decomposition into Movement Along
And Shifts in the EEC: 1984 to 2002

	Mean values (Table B.1)			EEC Coefficients (Table B.2)	Movement Along EEC (3) x (4)
	1984 (1)	2002 (2)	Diff (3)	1984 (4)	(5)
Total PM10 (tons)	155.1	145.6	-9.5		
	(1.74)	(1.24)	(2.14)		
Income	4.10	4.64	0.54	10.28***	5.55***
	(0.07)	(0.07)	(0.10)	(1.214)	(1.196)
Income squared	27.83	38.91	11.08	-0.125*	-1.39
	(1.07)	(1.54)	(1.87)	(0.0706)	(0.816)
Family size	2.73	2.56	-0.17	39.10***	-6.61***
	(0.03)	(0.02)	(0.04)	(2.938)	(1.692)
Family squared	9.78	8.64	-1.14	-1.930***	2.20***
	(0.24)	(0.17)	(0.29)	(0.375)	(0.709)
Age	46.70	48.38	1.67	3.430***	5.74***
	(0.40)	(0.30)	(0.50)	(0.399)	(1.845)
Age squared	2,488	2,630	142	-0.0292***	-4.16***
	(38.44)	(29.73)	(48.60)	(0.00383)	(1.521)
Married	0.61	0.54	-0.08	8.585***	-0.65**
	(0.01)	(0.01)	(0.01)	(2.940)	(0.251)
Race (White omitted)					
Black	0.10	0.11	0.01	-26.88***	-0.17
	(0.006)	(0.006)	(0.009)	(3.682)	(0.232)
Asian	0.01	0.04	0.03	6.455	0.21
	(0.003)	(0.003)	(0.004)	(17.95)	(0.575)
Other	0.02	0.01	-0.01	2.177	-0.01
	(0.003)	(0.002)	(0.004)	(11.13)	(0.0635)
Education (< high school omitted)					
High school	0.32	0.27	-0.04	9.423***	-0.41**
	(0.010)	(0.008)	(0.013)	(3.281)	(1.129)
Some college	0.19	0.30	0.11	9.672**	1.05**
	(0.009)	(0.008)	(0.012)	(3.860)	(1.392)
College	0.12	0.17	0.05	18.10***	0.94***
	(0.007)	(0.006)	(0.010)	(4.917)	(1.033)
Graduate	0.11	0.10	-0.01	24.47***	-0.21
	(0.007)	(0.005)	(0.008)	(5.288)	(0.639)

Continued on next page...

Table B.3 Continued

	Mean values (Table B.1)			EEC Coefficients (Table B.2)	Movement Along EEC (3) x (4)
	1984 (1)	2002 (2)	Diff (3)	1984 (4)	(5)
Region (Northeast omitted)					
Midwest	0.22 (0.009)	0.23 (0.007)	0.01 (0.011)	-15.44*** (3.533)	-0.22 (0.183)
South	0.27 (0.010)	0.36 (0.008)	0.09 (0.013)	-11.62*** (3.437)	-1.10*** (0.358)
West	0.17 (0.008)	0.21 (0.007)	0.04 (0.010)	-8.046** (3.878)	-0.36* (0.194)
Rural	0.16 (0.009)	0.13 (0.006)	-0.04 (0.011)	-28.56*** (3.844)	1.05*** (0.339)
Total change due to income (movement along EEC) (tons)					4.16
Total change due to other demographics (movement along EEC) (tons)					-2.71
Unexplained difference (shift in EEC) (tons)					-10.93

Note: Total household pollution is calculated by multiplying itemized household consumption with the pollution intensity of production for each type of good and summing for each household. Income is measured in tens of thousands of 1997 dollars. Coefficients are significant at the one (***), five (**), or ten (*) percent level.

Figure B.1

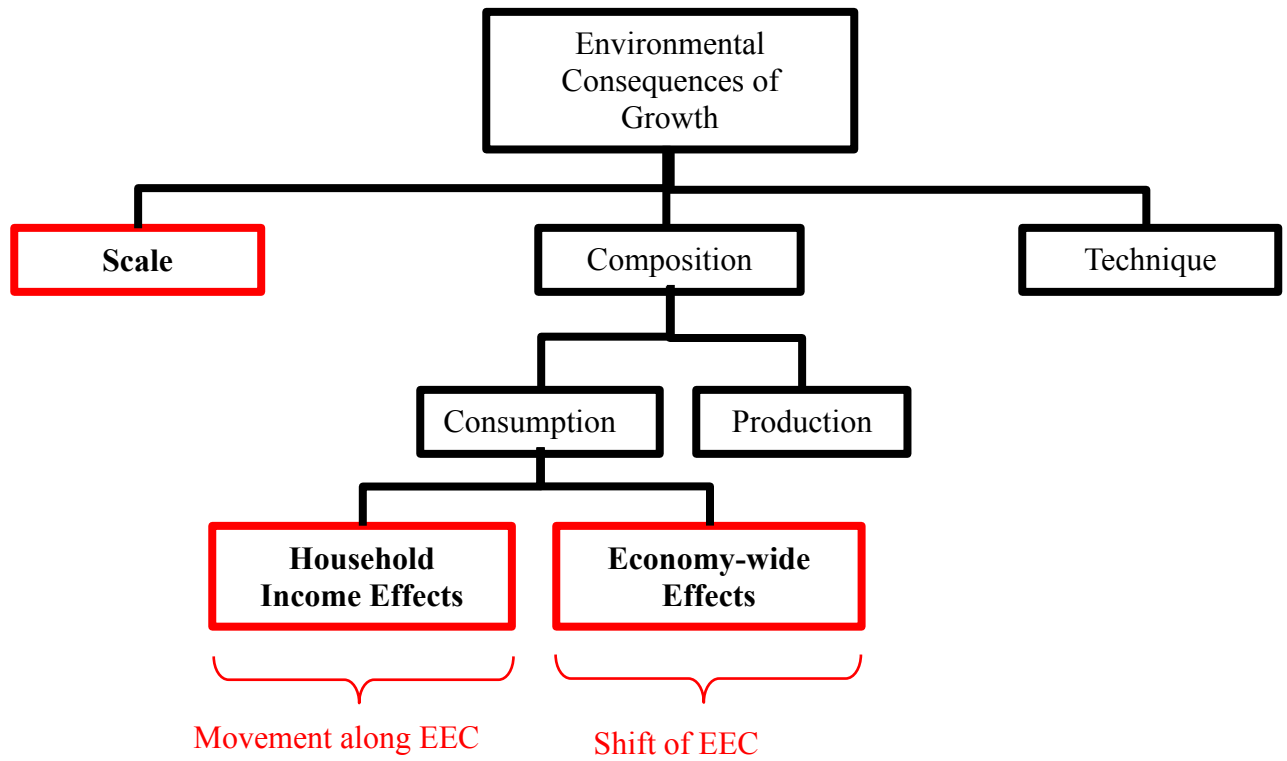


Figure B.2

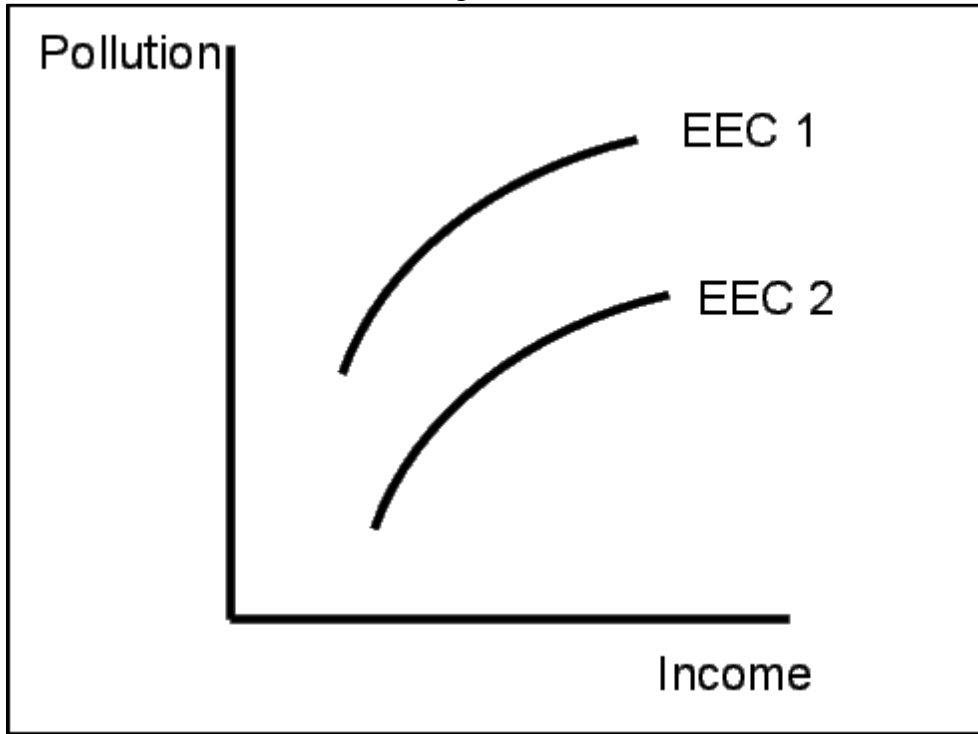


Figure B.3

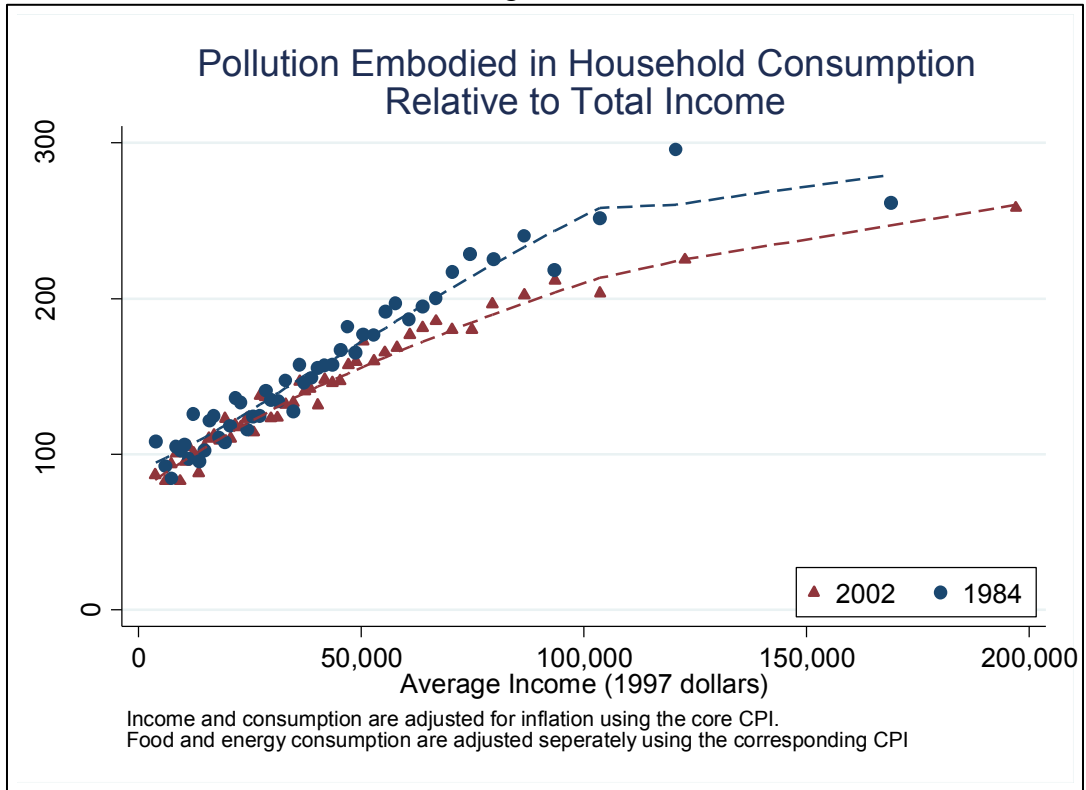


Figure B.4a




Figure B.4b




Figure B.5





Figure B.6



APPENDIX C
TABLES FOR CHAPTER III

Table 1
Summary Statistics for Estimation Sample

Variable	Sample Mean
Child in school fulltime (%)	0.71 (0.45)
Monthly per-capita consumption	2978.5 (1467.6)
Household PW intensity	77.53 (53.96)
Age	10.34 (2.81)
Gender (% male)	0.51 (0.50)
No. school-aged children	2.80 (1.45)
Agricultural household (%)	0.45 (0.50)
Self-employed household (%)	0.41 (0.49)
Education of head	3.16 (2.55)
Age of head	44.15 (10.72)
Gender of head (% male)	0.94 (0.23)
Household size	6.56 (2.38)
Land ownership (hectares)	1.05 (2.32)
Scheduled caste or tribe (%)	0.80 (0.40)
Number of observations	1306

Note: standard deviations are shown in parentheses.

Table 2
Child Schooling Intensity and Parent Participation in Public Works

Model	OLS	OLS	Klein and Vella
Dependent variable:	Child in school fulltime	Monthly per-capita consumption	Child in school fulltime
	(1)	(2)	(3)
Household PW intensity	-0.002** (0.0003)	0.394 (0.596)	-0.007* (0.003)
Age	0.204** (0.061)		0.200** (0.063)
Age squared	-0.012** (0.003)		-0.012** (0.003)
Gender	0.112* (0.044)		0.076 (0.047)
No. school-aged children	-0.064 (0.021)	-34.76 (41.76)	-0.079** (0.025)
Agricultural household	-0.044 (0.048)	-167.63* (86.03)	-0.080 (0.066)
Self-employed household	-0.053* (0.044)	304.10** (92.94)	-0.208* (0.091)
Education of head	0.040** (0.010)	68.28** (16.59)	0.030* (0.014)
Age of head	0.045* (0.018)	10.88 (23.12)	0.060* (0.025)
Age of head squared	-0.0005* (0.0002)	0.07 (0.25)	-0.0006* (0.0003)
Gender of head	-0.051 (0.012)	-161.64 (233.03)	0.051 (0.124)
Household size	0.045** (0.012)	246.54** (25.81)	0.142** (0.053)
Land ownership (area)	0.035** (0.012)	117.39** (17.51)	-0.018 (0.026)
Scheduled caste or tribe	-0.233** (0.084)	-243.47* (119.55)	-0.203** (0.072)

Continued on next page...

Table 2 Continued

Model	OLS	OLS	Klein and Vella
Dependent variable:	Child in school fulltime	Monthly per-capita consumption	Child in school fulltime
	(1)	(2)	(3)
Constant	-0.860† (0.445)	723.44 (604.39)	-1.47* (0.618)
Control Term: $A(x_i)v_i$			0.014† (0.008)
State effects	Yes	Yes	Yes
Seasonal effects	Yes	Yes	Yes
Year effects	Yes	Yes	Yes
Observations	1306	1306	1306
Adj. R-squared	0.320	0.626	0.261

Note: Significant at one (**), five (*) and ten (†) percent levels. The sample is restricted to households with non-zero public works participation. Standard errors in column 3 are based on 1000 bootstrap repetitions.

Table 3
First and Second Stage Regressions for the
Klein and Vella Estimator

Model	OLS	Exponential
Dependent Variable	Public works intensity	Ln[v ²]
	(1)	(2)
No. of school-aged children	-5.73** (1.04)	-0.007 (0.011)
Agricultural household	-15.74** (3.42)	0.070* (0.029)
Self-employed household	-8.94** (3.31)	-0.067** (0.024)
Education of head	-1.26 (0.48)**	-0.022** (0.008)
Age of head	1.57† (0.87)	0.018** (0.006)
Age of head squared	-0.02† (0.01)	-0.0002** (0.0001)
Gender of head	0.14 (4.53)	-0.022 (0.046)
Household size	9.09** (1.17)	0.045** (0.006)
Land ownership (area)	-1.75** (0.54)	0.006 (0.005)
Schedule caste or tribe	-10.86† (3.70)	0.016 (0.034)
Constant	22.60 (24.03)	1.16** (0.17)
State effects	Yes	Yes
Seasonal effects	Yes	Yes
Year effects	Yes	Yes
Observations	1306	1306
Adj. R-squared	0.271	0.925
Breusch-Pagan test	252.9	
White test	148.6	

Note: Results are significant at one (**), five (*) and ten (†) percent levels. The sample is restricted to households that had non-zero public works participation.

APPENDIX D

DATA APPENDIX FOR CHAPTER II

D.1 Introduction

The purpose of the CEX/pollution program is to find the level of pollution that is implicitly created as a result of individual household spending. This household pollution contribution is calculated by combining the CB/BLS Consumer Expenditure Survey (CEX) with per-dollar pollution intensities provided by the EPA Trade and Environmental Assessment Model (TEAM). The TEAM data provide pollution intensities for 1,148 industries, organized by NAICS code. The NBER CEX extracts provide consumer spending behavior for 109 various spending categories, of which 52 are relevant for finding household pollution contributions. Most of this program focuses on aggregating the 1,148 pollution intensities to match the 52 relevant expenditure categories. This requires finding appropriate weights and passing through several intermediate levels before finally obtaining CEX-level pollution intensities.

The aggregation procedure is complicated by two issues. First, each level of data uses a different coding system with ambiguity and overlap between codes. Second, the aggregation itself should be weighted by final consumption, as opposed to (readily available) total output. We use personal consumption expenditure (PCE) from the Bureau of Economic Analysis (BEA) “Use” table to proxy for final consumption. These figures are not available on a NAICS level. Instead, total output and pollution intensities for the 1,148 NAICS codes must first be aggregated to match the 471 input-output (IO) codes used in the BEA tables. These IO-level personal consumption expenditures then provide the necessary weights to complete the final stage of aggregation to CEX-level pollution intensities. In short, the overall procedure involves aggregating the pollution intensities from NAICS to IO weighted by total output, and then again from IO to CEX weighted by personal consumption expenditure.

After aggregation, we multiply the household expenditure for each CEX category by the per-dollar pollution intensity of that category to find individual household pollution for each spending category. Pollution from all categories is then combined to find the total pollution associated with each household’s spending. Although the resulting CEX pollution intensities are based on the 1997 TEAM data, we also apply these intensities to all available years of expenditure data. Lastly, we combine the modified NBER CEX extracts from each year into a

single file containing the original expenditure data from all years along with 33 additional pollution-related variables for each household. These 33 pollution variables represent three separate intensity measures for each of the eleven available pollutants.

The first step of this procedure is to consolidate the necessary data into a single master file. This also involves combining several different systems of industry classification and building a concordance from NAICS to IO to CEX codes. Early on we take a brief sidestep to manually calculate total pollution coefficients based on the TEAM direct coefficients. After that, we return to the original procedure and match the data to the personal consumption expenditure figures available from the BEA, which are used to weight the last stage of aggregation. This step in itself requires some initial aggregation because the total output data are in terms of NAICS codes whereas the BEA tables are in terms of IO codes. This also requires the creation of a sub-IO code system to capture certain IO codes that get split across separate CEX categories. After personal consumption expenditure figures have been calculated on a sub-IO level, the last stage of aggregation is a straightforward weighted average of sub-IO-level pollution intensities. Finally, the pollution intensities can be applied to the Consumer Expenditure Survey and we attribute an overall pollution contribution to each household's expenditure.

D.2 Data Consolidation

Combining the NAICS-level TEAM pollution data with CEX-level consumer expenditure data involves unifying several sources of information. In addition to the NBER CEX extracts, the primary supporting raw data come from the 1997 U.S. Economic Census, the 1997 TEAM study, and 1997 BEA "Use" table. The procedure hinges upon building and utilizing several concordances between the three separate classification systems used in these data: the North American Industry Classification System (NAICS), the BEA input-output codes (IO), and the NBER CEX expenditure categories (CEX).

The most specific classifications in the data are the NAICS codes, which characterize total output in the Census and pollution intensities in the TEAM data. These 6-digit codes drill down to specific industries starting from broad 2-digit classifications. In total, there are approximately 1,140 unique NAICS codes that describe all facets of the U.S. economy. The BEA input-output codes classify industries in a similar manner, although these 6-digit codes are defined somewhat differently from the NAICS system. With only 471 unique codes, the IO

codes are less specific, but capture largely the same information as NAICS, albeit in a more general setting. Lastly, the CEX codes used in the NBER CEX extracts are the most general classification system used in this program. The extracts are divided into 109 separate spending categories that are intended to provide a simple and consistent classification across multiple cross-sections of data. Of these 109 codes, about 52 are pertinent to finding household pollution contributions.

To begin, we combine the primary supporting data into one master file. The goal is to have a single file containing both the NAICS-level TEAM data and the NAICS-level Census data, where each record also indicates an IO code, sub-IO code, and CEX code.¹ After obtaining the initial raw data in an appropriate format, the first step is to compile a complete list of 6-digit NAICS codes. This list will provide the basic universe of codes for building the master file. A list of NAICS codes used in the 1997 Economic Census is available from the Census Bureau website.² From this list, we create a sub-concordance that links 2-, 3-, 4-, and 5-digit NAICS codes to their 6-digit components. This is a self-contained NAICS-NAICS concordance that we use when it is necessary to link specific NAICS to broader NAICS categories.

The next step is to combine in a separate file the complete list of 6-digit NAICS codes from the Census with a mapping from NAICS to IO codes. An official concordance is available on the BEA website, but it frequently links IO codes to broader NAICS categories (such as 2-, 3-, 4-, or 5-digit NAICS codes).³ Thus we combine the BEA NAICS-IO concordance with the NAICS-NAICS concordance to create a complete list containing all IO-NAICS pairs on a 6-digit level. Then we re-merge this with the original full list of NAICS codes to obtain a set of all possible NAICS codes along with the corresponding IO code for each record (every NAICS, except Public Administration, gets mapped to some IO code). This list provides the foundation for the data consolidation.

The master file also includes a corresponding CEX code for each 6-digit NAICS code. We manually constructed a NAICS-CEX concordance for this project by individually inspecting each NAICS code and determining which CEX code, if any, is most appropriate. The CEX codes themselves are based on underlying Universal Classification Codes (UCC), which make up the

¹ As will be discussed, sub-IO codes are created because certain IO codes are split across sufficiently different CEX categories and thus need to be divided into separate sub-IO components. Details are discussed later.

² See <http://www.census.gov/epcd/naics/naicscod.txt>

³ See http://www.bea.gov/industry/io_benchmark.htm

raw classification system used by the Consumer Expenditure Survey. Assignment of NAICS codes to CEX codes was primarily based on similarities between NAICS and the underlying UCC codes.

A list of UCC codes that correspond to each CEX category is available in the NBER CEX Extract documentation.⁴ Each individual 6-digit NAICS code was compared to this list to determine which CEX code best applies. There is no available mapping directly from UCC to NAICS and the two systems do not overlap exactly. For each record, a judgment was made regarding which CEX code best captures the subject of the NAICS code. Usually the connections were straightforward. Occasionally, however, there was room for interpretation.

As an example of the necessary interpretation, we assign NAICS 339113 “Surgical Appliance and Supplies Manufacturing” to CEX code 45 “Ophthalmic Products and Orthopedic Appliances” because the NAICS code encompasses “orthopedic devices, prosthetic appliances, surgical dressings, crutches, and surgical sutures.” In another example, we assign NAICS 314912 “Canvass and Related Product Mills” to CEX code 29 “Clothing and Shoes.” In this case, there is no direct indication of the appropriate category. Although the individual code itself does not fall into the “apparel manufacturing” section of the NAICS system, the broader CEX description specifically mentions “travel items, including luggage, and luggage carriers” and “[men's, women's, and children's] accessories.” For this reason, CEX 29 is the most appropriate category to capture canvass and related products.

Another difficulty we encountered while creating the NAICS-CEX concordance was separating the “input” NAICS from relevant NAICS codes. The overall objective is to find pollution intensities for consumers’ actual expenditures, so NAICS codes that do not correspond to consumer purchases should not be included directly. Those NAICS categories may affect household pollution if they are used as inputs to production of goods that households do buy, but for this reason, the TEAM data include two sets of pollution intensities: direct and total coefficients. The total coefficients capture pollution from direct production as well as production of inputs and inputs to inputs. Thus, the pollution intensity of input-related NAICS codes is already implicitly captured in the TEAM total coefficients. As a result, we must determine which NAICS categories may correspond to direct household purchases and which categories are

⁴ See http://www.nber.org/ces_cbo/Cexfam.pdf.

entirely input-related, the latter being excluded from the aggregation to CEX pollution intensities.

Many NAICS codes clearly stand out as intermediate products. For example, “Commercial Lithographic Printing” and “Metal Heat Treating” are not services that an individual consumer would purchase directly. Others are more ambiguous. If we determine a NAICS code to be primarily an input, it does not get mapped to a CEX code. For example, “Yarn Spinning Mills” and “Paint and Coating Manufacturing” are used as production inputs, but small amounts could be attributed to final household consumption as well. Still, in these cases, as in similar cases, the NAICS codes do not correspond to goods or services that are generally consumed by a household directly. As such, we treat these NAICS as inputs.

In addition to individual codes that are treated as inputs, several broad NAICS categories do not get mapped to CEX codes. These general categories represent goods and services that are not directly purchased by households, or do not capture the actual production of goods, which is what drives the inherent pollution intensity. All NAICS codes falling under the categories “Agriculture, Forestry, Fishing, and Hunting,” “Wholesale Trade,” “Retail Trade,” and “Public Administration” are not assigned to CEX codes.

As an additional check, we also compared the newly created NAICS-CEX concordance with the BEA data on personal consumption expenditure. These data come from the BEA “Use” table and are used later in the procedure to provide weights for aggregating to CEX coefficients. Here, they provide a handy reality check because they indicate a dollar amount for each industry that goes towards personal consumption expenditure. The comparison is not perfect, because PCE are only available by IO code, not NAICS code, but they do provide some indication about which general NAICS categories should be considered inputs.

As a general rule, if an IO code does not have any personal consumption expenditure associated with it, we ensure that all NAICS codes flowing into that IO code are also considered inputs to production. In the other direction, however, classifying NAICS codes associated with positive PCE is done on a case-by-case basis. Although some IO codes have non-zero PCE, logic demands that these codes remain classified as inputs. For example, “Poultry and Egg Production” has a PCE of \$2,673 million, but the pollution is already implicitly captured under the food manufacturing NAICS codes. Likewise, “Other Industrial Machinery Manufacturing”

and “Management Consulting Services,” for example, both have positive PCE, although both should remain classified as inputs.

In nearly every instance, each relevant NAICS code was assigned to a specific CEX code that we determined to be the most appropriate. For NAICS codes that had a small portion attributable to some other CEX category, we assigned the entire NAICS code to the best fitting CEX code. There are two exceptions, however. We assigned NAICS 524126 “Direct Property and Casualty Insurance Carriers” to both CEX 57 “Auto Insurance” and CEX 43 “Domestic Service, Other Household Operation.” The latter because CEX 43 includes homeowner’s and renter’s insurance. We also assigned NAICS 336991 “Motorcycle, Bicycle, and Parts Manufacturing” to both CEX 52 and CEX 63. My intention was to capture the recreation and motor-vehicle components separately. This multiplicity is a result of assigning NAICS codes directly to CEX codes, as opposed to assigning NAICS codes to IO codes and then IO codes to CEX codes. It does not induce double-counting because the procedure to capture split IO codes (discussed later) also intrinsically captures these split NAICS codes.

After having finishing researching, inspecting, and assigning each individual NAICS code, we re-merged the NAICS-CEX concordance into the master file. Thus the master file contains the full universe of NAICS codes, along with IO and CEX codes, and provides the basis for the remainder of the analysis. Rather than combining the remaining raw data together directly, each data source is independently merged into this master file. This way, we can account for any possible discrepancies between the codes used in different data sources. For example, the NAICS codes in the TEAM data and those used in the Census data do not match exactly. Thus, the master file provides the benchmark for merging. Any discrepancies that arise are resolved by visually checking any rogue records and consolidating where necessary.

Most inconsistencies between the TEAM and Census data occur because there are NAICS codes present in the TEAM data that are not present in the Economic Census. Of these codes, most have pollution intensities listed as zero or missing. For example, the “Public Administration” NAICS codes are not captured by the Economic Census, but are present in the TEAM data, although the pollution intensities for all NAICS in this category are zero. Similarly, the Census does not cover “Agriculture, Forestry, Fishing, and Hunting,” but there are non-zero pollution intensities available for these codes in the TEAM data. None of these NAICS codes get mapped to a CEX code. There are also four power generation NAICS codes used in the TEAM

data that are not used in the Census data. These receive special treatment during the aggregation process, and are discussed separately, below.

Since higher-level pollution coefficients will be found based on weighted averages of information in the master file, it is also necessary to include a consistent weighting metric for later aggregation. While the final aggregation to CEX-level coefficients will be weighted by personal consumption expenditure, which is introduced later, the initial aggregation to IO (and later sub-IO) level pollution coefficients requires a manually created metric to properly weight each base-level NAICS code. These weights are closely related to NAICS-level total output, but we cannot simply use the Economic Census total output because of slight incongruencies between the Census (providing the total output data) and the TEAM data (providing the pollution intensities).

Although the Census and the TEAM data characterize many of the same industries, the set of NAICS codes used by each do not match exactly, and data are missing for certain codes in one or both sources. For example, the Economic Census does not provide information on “Agriculture, Forestry, Fishing, and Hunting” NAICS codes, and several power generation NAICS are treated differently in the TEAM data than in the Census data. Thus we manually assign weights for the intermediate aggregation (from NAICS to IO) based on total output whenever possible, but still account for any discrepancies. We rely on the following rubric for assigning weights:

1. When data are available from both sources for all NAICS codes within a given IO code, the aggregation is weighted by total output.
2. If pollution data are missing entirely for a given NAICS code, that NAICS is assigned a weight of zero. TEAM data with values of “0” are not generally considered missing.
3. If total output data are missing for all NAICS codes within a given IO code, an equal weight is given to all NAICS codes.

There are also four additional pseudo-NAICS representing power generation in the TEAM data that require special attention. The TEAM study created these pseudo-NAICS in order to more acutely capture air pollution from fossil fuel generation. When the pollutant in question is air related, the TEAM data separate fossil fuel energy production into oil (221114), coal (221115), and natural gas (221116) components. When the pollutant is not air-related, no distinction is drawn between fuel types, and pseudo-NAICS 221117 captures the pollution from

all energy production, including fossil fuels, nuclear, hydroelectric, and other sources. In this way, pseudo-NAICS 221114-221116 and 221117 are mutually exclusive. Since my analysis is not concerned with separately capturing fuel source, we consolidate the TEAM pseudo-NAICS before aggregating to IO-level coefficients.

Combining the power generation pseudo-NAICS involves creating a new pseudo-NAICS 221100 to capture the production intensity for all types of pollutants. When 221117 is present in the TEAM data, the new pseudo-NAICS simply re-states the same information. When 221117 is not present (and the TEAM data separate fuel type) the new pseudo-NAICS represents a weighted average of the existing TEAM data. Since the existing pseudo-NAICS are not based on census data, the weighting is instead based on the Energy Information Administration's Annual Energy Review of 2008, which includes 1997 output data itemized by fuel type.⁵ During the consolidation, we drop NAICS 221112, which represents fossil fuel power generation, because this pollution is already captured by the TEAM pseudo-NAICS. Also, the TEAM study is not able to distinguish between power generation and distribution, so the distribution-related NAICS 221121 and 221122 are dropped from the analysis. The resulting new pseudo-NAICS 221100 corresponds one-for-one with the power generation IO code 221100.

The above rubric, along with the special treatment of power generation codes, captures all situations that arise during the intermediate aggregations from NAICS to IO and NAICS to sub-IO levels. These intermediate aggregations are necessary to find pollution coefficients that match the BEA personal consumption expenditure and total requirements data that are used later. These raw data are supplied by the BEA on an IO-level and cannot be further disaggregated (below sub-IO codes, discussed later).

Thus the master file contains all the basic components necessary for aggregating the NAICS-level pollution intensities to CEX level. It contains the base NAICS-level pollution intensities and total output data, complete mappings from NAICS to IO and from NAICS to CEX codes, and appropriate weights for aggregating the NAICS-level data to IO and sub-IO codes. This file is the foundation for manually calculating total coefficients based on direct coefficients and for the intermediate aggregation from NAICS- to sub-IO-level pollution intensities. This intermediate aggregation is necessary so that the pollution intensities match the personal

⁵ Available at <http://www.eia.doe.gov/emeu/aer/elect.html>.

consumption expenditure data that are used to weight the final aggregation to CEX-level intensities.

D.3 Manually Calculate Total Pollution Intensities

The TEAM pollution data contain pollution intensities measured in direct and total terms. The direct coefficients measure the pollution intensity associated with each NAICS category, but only include the pollution generated from direct production. The total coefficients capture the same information, but also include additional pollution from the production of inputs and inputs to inputs, etc. In addition to the original total coefficients included in the TEAM data, it is possible to manually re-calculate total coefficients based on the provided direct coefficients. Doing so requires supplementing the data with a Total Requirements matrix included as part of the BEA input-output tables.

The BEA regularly publishes input-output tables covering industries and commodities in the U.S. economy. One of these tables, the “Industry-by-Industry Total Requirements” matrix, shows the dollar amount of each industry that is needed to create one dollar of output for each other industry. These figures include the production of inputs and inputs to inputs, etc. In other words, the Total Requirements matrix shows the overall dollar amount of each industry that is needed, in total, to create one dollar of output for any other given industry. Thus for each industry, we find the dollar amount of each other industry that is used as an input to that industry and we multiply by the per-dollar direct pollution coefficient. Summing across inputs gives the total pollution created both directly and as a result of inputs and inputs to inputs for each dollar of output of that industry.⁶ This captures the same information as the total pollution coefficient in the original TEAM data, although on an IO level (as opposed to NAICS).

Since the BEA tables are organized by IO code, implementing this procedure requires first finding direct pollution intensities on an IO level. Thus we use the master file to aggregate the TEAM pollution coefficients to a compatible IO level. For each IO code, we average the pollution intensity for each of the NAICS-level components, weighted by the NAICS-level total output. The BEA provides a one-to-one mapping of most NAICS to IO codes, so it is a

⁶ In matrix terms, this is: Total Coefficients = (Total Requirements)' (Direct Coefficients)

straightforward aggregation for most codes. IO codes that do not correspond to any NAICS codes are not included in this part of the analysis.⁷

The IO-NAICS mapping provided by the BEA is nearly complete, but does not differentiate construction-related codes. The BEA maps all such IO codes to a general “23*” NAICS code (23 is the overarching 2-digit “Construction” category in the NAICS system). Since there are several construction related IO codes and we cannot tell which specific NAICS should be associated with each one, we consolidate all construction related IO codes into a single IO code, which we also label 23*. Before the total pollution coefficients can be found manually, the same adjustment must be made to the Total Requirements matrix.

Since the matrix shows the amount of each industry necessary to produce each other industry, all IO codes are present in both the column and row headings. Specifically, the matrix shows the amount of each row code that is necessary to produce one dollar of output in each column code. Thus, to consolidate the construction-related IO codes within the Total Requirements matrix, we sum across rows and average across columns for all rows/columns that correspond to the newly defined IO code 23*. The implicit assumption is that a dollar output of 23* uses each of the 23xxxx IO codes in equal proportion, which justifies consolidating columns by simple averaging. After the construction codes have been consolidated, the Total Requirements matrix conforms to the direct pollution coefficients in the TEAM data and these two sources can be combined to calculate IO-level total pollution intensities.

D.4 Calculate sub-IO-Level Pollution Intensities and Household Expenditure

The main purpose of the procedure is to aggregate the NAICS-level pollution coefficients to match the CEX-level Consumer Expenditure Data, the end result being a set of coefficients that measure the per-dollar pollution intensity of each CEX expenditure category and thus overall household pollution contribution. These CEX-level pollution intensities ultimately represent a weighted average of the original NAICS-level figures. However, we weight the aggregation by final consumption expenditure rather than total output. Because of this, we cannot proceed directly from NAICS to CEX code. Instead, an intermediate aggregation to IO

⁷ There are 28 additional codes that represent final uses, government uses, and special industries (the F-, S-, and V-IO codes). These do not correspond to any NAICS industries.

code must first occur so that we can match the data to the personal consumption expenditure (PCE) figures published by the BEA.

We use personal consumption expenditure to weight the final aggregation in order to more closely capture the relationship between pollution and overall household expenditure. Households are end-users and increases in household expenditure directly correspond to increases in final demand. A single dollar increase in household expenditure requires more than a single dollar increase in total output, since some of the total output is needed as an input for production. Similarly, since household expenditure represents an increase in final demand, when we consolidate the pollution intensities from several IO we should weight the aggregation by final demand. To this end, personal consumption expenditure acts as a proxy for final demand.

Since the total pollution coefficients already capture pollution from direct production, along with the pollution from producing inputs and inputs to inputs, weighting the aggregation by total output may mis-state pollution intensities. For example, pollution intensities would be overstated if there are high-polluting goods that are used primarily as inputs, or vice-versa. Including input-production in the aggregation over-weights the pollution from inputs because this pollution is already counted in the total coefficient for the final product.

We use personal consumption expenditure from the BEA “Use” table to proxy for final demand as weights in the aggregation. The BEA regularly publishes the “Make/Use Tables” as part of the Benchmark Input-Output Accounts of the United States.⁸ The “Use” table shows the amount of each commodity that is consumed by each industry and final consumer. The tables specifically quantify personal consumption expenditure in terms of commodities. These commodities are simply assigned the IO code of the industry for which they are the primary output, and adjustments are made to capture secondary output. In this way, the PCE data do not exactly match the industry-level classification of the pollution coefficients, but for the purposes of weighting the aggregation, the commodity-level PCE are adequate. Since the data are in terms of IO codes, not NAICS codes, it is not possible to calculate personal consumption expenditure on a NAICS level. Instead, we first use the master file to aggregate the output and pollution data to match the BEA IO codes.

The first step in finding the appropriate weights for aggregation actually involves some *disaggregation*. We have already organized the data into a master file that matches NAICS-level

⁸ Available from the BEA website at http://www.bea.gov/industry/io_benchmark.htm.

data to corresponding IO and CEX codes. In most cases, there is a clear link from NAICS to IO to CEX and a related group of NAICS codes all correspond to a specific IO code *and* a specific CEX code. In some cases, however, the NAICS codes that compose a given IO code do not all get mapped to the same CEX code, so the IO code is split across multiple CEX categories.⁹ This is a result of separately mapping NAICS and IO codes to CEX codes, which allows for greater accuracy in aggregation. For example, IO code 339910 “Jewelry and Silverware Manufacturing” is mapped to both CEX code 31 “Jewelry and Watches” and CEX code 36 “Furniture and Durable Household Equipment” (for the silverware). Two approaches are taken to resolve the split IO codes: we) IO codes are disaggregated into sub-IO codes so that a portion of each split code can be attributed to a separate CEX category and ii) sufficiently similar CEX categories are consolidated to encompass split codes. The second approach is implemented along with the final aggregation to CEX level, as discussed later.

For the first, we create sub-IO codes to categorize the individual components of each split IO codes. These sub-IO codes can be uniquely attributed to CEX categories. The majority of IO codes are not split across CEX categories and can thus have sub-IO codes that are identical to the original IO codes. For those that are split (and not captured later by the CEX consolidation), we create a unique sub-IO code by appending a letter suffix to the original IO code, thus creating a new 7-digit code. Unlike the original IO codes, every 7-digit alphanumeric sub-IO code can be uniquely mapped in both directions, to NAICS and CEX codes. The sub-IO total output is found by summing the NAICS-level total output and the sub-IO pollution intensities are found by taking the weighted average over the relevant NAICS codes.

The weights that we use for calculating sub-IO pollution intensities are nearly identical to the weights that we originally created in the master file. These weights were used for aggregating from NAICS to IO code, but since sub-IO codes and IO codes are identical in most situations, the original weights generally apply. In the special cases where sub-IO codes and IO codes are not the identical, the same weighting rubric from the master file is applied to create new weights for sub-IO codes. Since individual NAICS codes are not split across IO or sub-IO codes (with

⁹ In three instances, the underlying NAICS code is actually split actually across CEX categories. These are 336991 “Motorcycle, Bicycle, and Parts manufacturing” (Recreation and Sports Equipment, New and Used Motor Vehicles); 324110 “Petroleum Refineries” (Gasoline and Oil, Fuel Oil and Coal); and 522292 “Non-depository Credit Intermediation and Related Activities” (Mortgage Interest, Mortgage Principle, Lump-Sum Mortgage Payments). 336991 is manually disaggregated. The multiple CEX codes that contain 324110 and 522292 are later combined to avoid dividing more IO codes.

one exception, discussed below), creating weights for sub-IO codes basically involves splitting the weights of specific IO codes into their relevant sub-IO components. Simply re-creating sub-IO-level weights by following the original weighting rubric from the master file, this time applied to sub-IO codes, captures this effect.

After creating weights and calculating sub-IO-level pollution intensities, we also disaggregate personal consumption expenditure to match the newly found pollution coefficients. We do this by dividing PCE into separate sub-IO components, using NAICS-level total output to weight the disaggregation. This time, however, only “relevant” NAICS codes contribute to the weighting, we include only NAICS codes that eventually flow into a CEX category. Since these weights are used to disaggregate final consumption, it doesn’t make sense to base those weights on output that is used primarily as an input.

As an example, consider IO code 811200 “Electronic Equipment Repair and Maintenance.” This code is split into two sub-IO codes: 811200D, which captures NAICS 811212 “Computer and Office Machine Repair,” and 811200R, which captures NAICS 811211 “Consumer Electronic Repair.” There are also two other NAICS codes contained within the IO code that do not get mapped to any CEX code, namely 811213 “Communication Equipment” and 811219 “Precision Equipment” maintenance and repair. Personal consumption expenditure from the overarching IO code is split between the two sub-IO codes based on the relative shares of NAICS-level total output. The latter two NAICS codes, which do not get mapped to any CEX category, are not used to calculate the relative total output for the sub-IO codes. Thus, the two related sub-IO items are also not assigned any personal consumption expenditure; all PCE is distributed to either sub-IO code 811200D or 811200R.

The last step is to manually adjust for IO code 336991 “Motorcycle, Bicycle, and Parts Manufacturing.” In this special case, an individual NAICS code is split across two separate sub-IO codes. Final consumption is divided based on relative market shares using supplemental data from the U.S. Census Bureau’s Annual Survey of Manufacturers.¹⁰ Based on this data, 58 percent of the final consumption is attributed to 336991M (for motorcycles) and 42 percent to 336991B (for bicycles). Once this final adjustment has been made, the data are organized and weighted appropriately for the final aggregation to CEX level.

¹⁰ See <http://www.census.gov/prod/2000pubs/m98-as2.pdf>.

The construction-related NAICS also warrant special mention, although we do not follow any special procedure while addressing these codes. Recall that the BEA does not map construction-related IO codes to specific NAICS codes, and as a result, we also consolidate construction-related IO codes into a new 23* IO code. Only two of the NAICS codes that flow into IO code 23* correspond to a CEX code (CEX 86, Construction). These are 233220 (“Multifamily Housing Construction”) and 233210 (“Single Family Housing Construction”). The pollution intensities from the TEAM study for these particular NAICS codes are zero for all eleven pollutants. Further, there is no personal consumption data available for any of the construction-related IO codes. Subsequently, these two NAICS, and the general 23* IO category, are excluded from the analysis (if nothing else, by receiving a zero weight in the aggregation based on PCE). The Construction CEX category (CEX 86) thus consists of IO codes that come from outside 23* IO category. These codes primarily originate from the manufacturing sector and all include non-zero personal consumption expenditure.

The result thus far is a properly weighted set of sub-IO-level pollution intensities derived from the TEAM data along with corresponding sub-IO-level personal consumption expenditure from the BEA “Use” table. These two datasets can be combined to calculate CEX-level pollution intensities, which are simply the average of the sub-IO-level intensities weighted by PCE. The aggregation is weighted by PCE rather than total output in order to capture the “final demand” nature of household expenditure. Indeed, this is the reason for the initial aggregation to sub-IO code (to match the PCE data) rather than simply aggregating directly to CEX code based on total output.

D.5 Aggregate from sub-IO to CEX

After finding pollution intensities and personal consumption expenditure on the sub-IO level, we are nearly ready to complete the aggregation to CEX level. Before the final averages can be found, however, we must first consolidate a handful of CEX categories. Recall that several IO codes are split across multiple CEX codes and that we did not disaggregate these IO codes when the CEX categories were sufficiently similar. Instead, we now combine these CEX codes into new, composite codes that capture the split IO codes, along with any other relevant IO codes. For example, “Expenses of Handling Life Insurance,” “Auto Insurance,” and “Health Insurance” are each unique CEX categories. It is unclear which CEX should capture IO code

524100 “Insurance Carriers.” Rather than dividing the IO code, we create a new CEX category “Insurance” to capture the related CEX codes. This removes any ambiguity about the mapping of IO codes into CEX categories. We make similar adjustments for publications, fuel oil, transportation, debt services, and food/alcohol on-premise.

After consolidating similar CEX codes, there is a clear and unique mapping from each sub-IO code to a specific CEX category. At this point, the pollution intensities are easily aggregated to CEX level. To do this, we take the weighted average of all sub-IO-level pollution intensities within a particular CEX category using personal consumption expenditure as the weighting metric. The resulting data represent the average per-dollar pollution intensity of each relevant CEX expenditure category based on the 1997 NAICS level pollution intensities found in the TEAM study.

D.6 Combine Aggregate Pollution Intensities with NBER CEX Extracts

After having obtained CEX-level pollution intensities, we combine this data with the NBER CEX extracts to create the final product, which is a single file containing the original NBER CEX extracts along with the amount of pollution implicitly created by each household’s annual expenditure. We do this for all relevant years of extracts and combine these to create a large file containing all necessary household expenditure data with the associated levels of pollution.

Each record of the NBER CEX extract contains expenditure data for an individual household. The expenditure is divided into 109 categories that are kept consistent across years. Thus to combine these data with the CEX-level pollution intensities, we multiply the expenditure from each spending category by the per-dollar pollution intensity of that specific category to find the level of pollution implicitly created as a result of each spending category. Then we sum the category specific pollution within households to find the total amount of pollution implicitly created as a result of the individual household’s entire annual expenditure. For each record, this procedure is repeated for all 33 pollution measures (three types of coefficients capturing eleven pollutants). Finally, we merge the household-level overall pollution contributions back into the original NBER CEX extracts.

We repeat this process for all available years of the NBER CEX extract. Each time, we apply the 1997 CEX-level pollution intensities to the household expenditure from a given year.

Having completed the procedure for each individual year, we combine all of these modified NBER CEX extracts into a single file. The end result contains the original NBER CEX extracts from all available years along with 33 variables corresponding to the levels of pollution implicitly created as a result of each household's yearly total consumption expenditure.

BIBLIOGRAPHY

- Abt Associates Inc. (2009). Trade and Environmental Assessment Model: Model Description. Prepared for the U.S. Environmental Protection Agency National Center for Environmental Economics, Cambridge, MA.
- Afridi, F., Mukhopadhyay, A., & Sahoo, S. (2012). *Female labour force participation and child education in India: The effect of the national rural employment guarantee scheme* (No. 6593). Discussion Paper series, Forschungsinstitut zur Zukunft der Arbeit.
- Alberini, A., Cropper, M., Krupnick, A., & Simon, N. B. (2004). Does the value of a statistical life vary with age and health status? Evidence from the US and Canada. *Journal of Environmental Economics and Management*, 48(1), 769-792.
- Aldy, J. E., & Viscusi, W. K. (2003). Age variations in workers' value of statistical life (No. w10199). *National Bureau of Economic Research*.
- Aldy, J. E., & Viscusi, W. K. (2007). Age differences in the value of statistical life: revealed preference evidence. *Review of Environmental Economics and Policy*, 1(2), 241-260.
- Aldy, J. E., & Viscusi, W. K. (2008). Adjusting the value of a statistical life for age and cohort effects. *The Review of Economics and Statistics*, 90(3), 573-581.
- Anderson, M., & Auffhammer, M. (2011). Pounds that kill: the external costs of vehicle weight. Working Paper No. w17170. National Bureau of Economic Research.
- Andersson, H. (2007). Willingness to pay for road safety and estimates of the risk of death: Evidence from a Swedish contingent valuation study. *Accident Analysis & Prevention*, 39(4), 853-865.
- Andersson, H. (2008). Willingness to pay for car safety: evidence from Sweden. *Environmental and Resource Economics*, 41(4), 579-594.
- Andreoni, J., & Levinson, A. (2001). The simple analytics of the environmental Kuznets curve. *Journal of public economics*, 80(2), 269-286.
- Arrow, K., Solow, R., Portney, P. R., Leamer, E. E., Radner, R., & Schuman, H. (1993) *Report of the NOAA Panel on Contingent Valuation*.
- Ashenfelter, O., & Greenstone, M. (2002). Using mandated speed limits to measure the value of a statistical life (No. w9094). *National Bureau of Economic Research*.

- Atkinson, S. E., & Halvorsen, R. (1990). The valuation of risks to life: evidence from the market for automobiles. *The Review of Economics and Statistics*, 133-136.
- Bartlett, B. (1994, September 14). The High Cost of Turning Green. *The Wall Street Journal*.
- Blomquist, G. C., Miller, T. R., & Levy, D. T. (1996). Values of risk reduction implied by motorist use of protection equipment: new evidence from different populations. *Journal of Transport Economics and Policy*, 55-66.
- Brown, L., Yohannes, Y., and Webb, P. (1994). "Rural Labor-Intensive Public Works: Impacts of Participation on Preschooler Nutrition: Evidence from Niger." *The American Journal of Agricultural Economics*, Vol. 76, No. 5, Proceedings Issue.
- Chumacero, R. and Paredes, R. (2003). "Does Centralization Imply Better Targeting?: Evaluating Emergency Employment Programs in Chile." Manuscript, IV Reunión sobre Pobreza y Distribución del Ingreso, Universidad Nacional de La Plata.
- Copeland, B., & Taylor, S. M. (2005). *Trade and the Environment: Theory and Evidence*. Princeton, NJ: Princeton University Press.
- Dahlberg, M., & Eklöf, M. (2003). Relaxing the IIA assumption in locational choice models: a comparison between conditional logit, mixed logit, and multinomial probit models. Working Paper. Uppsala University. Available at: <http://uu.diva-portal.org/smash/record.jsf?pid=diva2:129254>
- Datt, D. and Ravallion, M. (1994). "Transfer Benefits From Public-Works Employment: Evidence for Rural India." *The Economic Journal*, Vol 104, No. 427.
- Department of Rural Development. (2013a). "Mahatma Gandhi National Rural Employment Guarantee Act, 2005 (Mahatma Gandhi NREGA) Operational Guidelines, 4th Edition." Ministry of Rural Development, Government of India. New Delhi.
- Department of Rural Development. (2013b). "MGNREGA Outcomes: Physical Employment." Ministry of Rural Development, Government of India. Available at: http://164.100.12.7/Netnrega/mpr_ht/nregampr_dmu.aspx?fin_year=2010-2011&month=Latest&flag=1&page1=S.
- DeShazo, J. R., & Cameron, T. A. (2004). Mortality and morbidity risk reduction: an empirical life-cycle model of demand with two types of age effects. *Unpublished paper*, Department of Policy Studies, University of California at Los Angeles.
- Dev, S. M. (1995). India's (Maharashtra) employment guarantee scheme: lessons from long experience. *Employment for poverty reduction and food security*, 108-143.

- Devereux, S. and Solomon, C. (2006). "Employment Creation Programmes: The International Experience." *Issues in Employment and Poverty Discussion Paper*, Economic and Labor Market Analysis Department, International Labor Office, Geneva.
- Diamond, P. A., & Hausman, J. A. (1994). Contingent valuation: Is some number better than no number?. *The Journal of economic perspectives*, 8(4), 45-64.
- Dreyfus, M. K., & Viscusi, W. K. (1995). Rates of time preference and consumer valuations of automobile safety and fuel efficiency. *Journal of Law and Economics*, 38, 79.
- Dreze, J., & Oldiges, C. (2007). How is NREGA doing?. In *Presentation made in a seminar at the Institute of Human Development, IHD, New Delhi*. Available at: http://knowledge.nrega.net/651/1/How_is_NREGA_Doing.doc.
- Evans, M. F., & Smith, V. K. (2006). Do we really understand the age-VSL relationship?. *Resource and Energy Economics*, 28(3), 242-261.
- Evans, M. F. & Smith, V. K. (2008). Complementarity and the Measurement of Individual Risk Tradeoffs: Accounting for Quantity and quality of life effects. *Environmental and Resource Economics*, 41(3): 381-400.
- Evans, M. F., & Schaur, G. (2010). A quantile estimation approach to identify income and age variation in the value of a statistical life. *Journal of Environmental Economics and Management*, 59(3), 260-270
- Farré, L., Klein, R., & Vella, F. (2012). Does Increasing Parents' Schooling Raise the Schooling of the Next Generation? Evidence Based on Conditional Second Moments*. *Oxford Bulletin of Economics and Statistics*, 74(5), 676-690.
- Gayer, T. (2004). The fatality risks of sport-utility vehicles, vans, and pickups relative to cars. *Journal of Risk and Uncertainty*, 28(2), 103-133.
- Grossman, G., & Krueger, A. (1993). Environmental Impacts of a North American Free Trade Agreement. The Mexico-U.S. Free Trade Agreement. (P. M. Garber, Ed.) Cambridge, MA: MIT Press.
- Grossman, G. M., & Krueger, A. B. (1995). Economic growth and the environment. *The quarterly journal of economics*, 110(2), 353-377.
- Hausman, Jerry. (2012). "Contingent Valuation: From Dubious to Hopeless." *Journal of Economic Perspectives*, 26(4): 43-56.

- Hausman, J., Newey, W., & Powell, J. (1995). Nonlinear Errors in Variables: Estimation of Some Engel Curves. *Journal of Econometrics*, 50, 205-234.
- Heckman, J., LaLond, R., and Smith, J. "The Economics and Econometrics of Active Labor Market Programs." Prepared for The Handbook of Labor Economics, Vol. III, Ashenfelter, O. and Card, D. editors; 1999.
- Hershman, L. (2001). The US new car assessment program (NCAP): Past, present and future. In *International Technical Conference on the Enhanced Safety of Vehicles*.
- Jacobsen, M. R. (2013). Fuel economy and safety: the influence of vehicle class and driver behavior. *American Economic Journal*, 5(3).
- Jaffe, A. N. (2003). Technological Change and the Environment. In K. M. Vincent (Ed.), *Handbook of Environmental Economics* (Vol. 1, pp. 461-516). Amsterdam: North-Holland/Elsevier.
- Johannesson, M., Johansson, P. O., & Löfgren, K. G. (1997). On the value of changes in life expectancy: Blips versus parametric changes. *Journal of Risk and Uncertainty*, 15(3), 221-239.
- Johansson, P. O. (2002). On the definition and age-dependency of the value of a statistical life. *Journal of Risk and Uncertainty*, 25(3), 251-263.
- Kahane, C. J. (1994). Correlation of NCAP performance with fatality risk in actual head-on collisions. NHTSA Technical Report number: DOT HS 808 061. Washington, DC.
- Klein, R., & Vella, F. (2010). Estimating a class of triangular simultaneous equations models without exclusion restrictions. *Journal of Econometrics*, 154(2), 154-164.
- Kniesner, T. J., Ziliak, J. P., & Viscusi, W. K. (2006). Life-cycle consumption and the age-adjusted value of life. *Harvard Law School John Olin Center for Law, Economics and Business Discussion Paper Series, Paper 459*. http://lsr.nellco.org/harvard_olin/459.
- Kropko, J. (2010). A comparison of three discrete choice estimators. *Unpublished paper*. Available from: <http://www.unc.edu/kropko/paper1.pdf> [Accessed 17 June 2011].
- Krupnick, A. (2007). Mortality-risk valuation and age: stated preference evidence. *Review of Environmental Economics and Policy*, 1(2), 261-282.
- Krupnick, A., Alberini, A., Cropper, M., Simon, N., O'Brien, B., Goeree, R., & Heintzelman, M. (2002). Age, health and the willingness to pay for mortality risk reductions: a contingent valuation survey of Ontario residents. *Journal of risk and Uncertainty*, 24(2), 161-186.

- León, G. & Miguel, E. (2013). Transportation choices and the values of statistical life. Working Paper 19494. National Bureau of Economic Research.
- Leontief, W. (1970). Environmental Repercussions and the Economic Structure: An Input-Output Approach. *Review of Economics and Statistics*, 52(3), 262-271.
- Levinson, A. (2009). Technology, International Trade, and Pollution from U.S. Manufacturing. *American Economic Review*, 99(5), 2177-92.
- Lewbel, A. (1991). The Rank of Demand Systems: Theory and Nonparametric Estimation. *Econometrica*, 59, 711-730.
- Mahapatra, R., Sakhuja, N., Das, S., & Singh, S. (2008). The National Rural Employment Guarantee Act (NREGA): Opportunities and Challenges. *New Delhi: Centre for Science and Environment*.
- McFadden, D. (1974). The measurement of urban travel demand. *Journal of Public Economics*, 3(4), 303-328.
- McFadden, D. (1978). Modelling the choice of residential location. In: Karlqvist, A., Lundqvist, L., Snickars, F., Weibull, J. (Eds.), *Spatial Interaction Theory and Residential Location*. Amsterdam, North-Holland, 75–96.
- Millimet, D. L., & Roy, J. (2012). Three New Empirical Tests of the Pollution Haven Hypothesis When Environmental Regulation is Endogenous.
- Millimet, D. L., & Tchernis, R. (2012). Estimation of treatment effects without an exclusion restriction: With an Application to the Analysis of the School Breakfast Program. *Journal of Applied Econometrics*.
- Mount, T., Weng, W., Zhang N., & Chestun, L. (2004). The effects of age and family status on the value of a statistical life – evidence from the automobile market and a national survey of automobile use. Report C.R. 824393-01-0. Prepared for the US Environmental Protection Agency. Washington, DC.
- NADA. (2012). NADA Used Vehicle Price Forecast: Winter/Spring 2012-2013. www.nada.com/b2b. Mclean, VA.
- National Bureau of Economic Research (NBER). (2000). Consumer Expenditure Survey Family-Level Extracts -- 1980:1 - 2003:2. Last updated July 2011. Cambridge, MA.

- National Highway Traffic Safety Administration. (2010). Notice of Proposed Rulemaking: Federal Motor Vehicle Standard, Rearview Mirrors. 49 CFR Parts 571 and 585. Docket No. NHTSA-2010-0162. RIN 2127-AK43.
- National Highway Traffic Safety Administration. (2012). 5-Star Safety Ratings: Frequently Asked Questions. Safercar.gov.
- New York City Department of Health and Mental Hygiene. (2011). Air Pollution and the Health of New Yorkers: The impact of Fine Particles and Ozone. Stevens, L. M. (ed.). www.nyc.gov/html/doh/downloads/pdf/eode/eode-air-quality-impact.pdf.
- OECD. (2012). Mortality Risk Valuation in Environment, Health, and Transport Policies. OECD Publishing. <http://dx.doi.org/10.1787/9789264130807-en>.
- Persson, U., Norinder, A., Hjalte, K., & Gralén, K. (2001). The value of a statistical life in transport: Findings from a new contingent valuation study in Sweden. *Journal of Risk and Uncertainty*, 23(2), 121-134.
- Petrin, A., & Train, K. (2010). A control function approach to endogeneity in consumer choice models. *Journal of Marketing Research*, 47(1), 3-13.
- Portney, P. R. (1981). Housing prices, health effects, and valuing reductions in risk of death. *Journal of Environmental Economics and Management*, 8(1), 72-78.
- Portney, P. R. (1994). The contingent valuation debate: why economists should care. *The Journal of Economic Perspectives*, 8(4), 3-17.
- Ravallion, M., van de Walle, D., Dutta, P., & Murgai, R. (2013). Testing Information Constraints on India's Largest Antipoverty Program. *Policy Research Working Paper*, 6598.
- Reddy, D. N., Tankha, R., Upendranadh, C., & Sharma, A. N. (2010). National rural employment guarantee as social protection. *IDS Bulletin*, 41(4), 63-76.
- Rohlf, C., Sullivan, R., & Kniesner, T. (2012). New estimates of the value of a statistical life using air bag regulations as a quasi-experiment. *IZA Discussion Paper No. 6994*.
- Schnier, K. E., Horrace, W., & Felthoven, R. (2009). The Value of Statistical Life: Pursuing the Deadliest Catch. *Center for Policy Research Working Paper*, (117).
- Schroeder, E. (2010). The Impact of Microcredit Borrowing on Household Consumption in Bangladesh. *Department of Economics, Georgetown University*.

- Scotton, C. R., & Taylor, L. O. (2011). Valuing risk reductions: Incorporating risk heterogeneity into a revealed preference framework. *Resource and Energy Economics*, 33(2), 381-397.
- Sheehan-Connor, D. (2012). Life and Death at the CAFE: Predicting the Impact of Fuel Economy Standards on Vehicle Safety. Wesleyan Economic Working Papers, No. 2012-002. Wesleyan University, Middletown, CT.
- Shepard, D. S., & Zeckhauser, R. J. (1984). Survival versus consumption. *Management Science*, 30(4), 423-439.
- Smith, V. K., Evans, M. F., Kim, H., & Taylor Jr, D. H. (2004). Do the near-elderly value mortality risks differently?. *Review of Economics and Statistics*, 86(1), 423-429.
- Train, K. (1993). *Qualitative Choice Analysis: Theory, Econometrics, and an Application to Automobile Demand*. MIT Press: Cambridge, Massachusetts.
- Train, K. (2009). *Discrete Choice Models with Simulation*, 2nd ed. Cambridge University Press, New York, New York.
- US Bureau of Labor Statistics. (2008). *Consumer Expenditure Survey Anthology*. Washington, DC.
- US Environmental Protection Agency. (2002). *Technical Addendum: Methodologies for the Benefit Analysis of the Clear Skies Initiative*. Washington, DC.
- US Environmental Protection Agency. (2010a). *Guidelines for Preparing Economic Analyses: Appendix B*. EPA 240-R-10-001. Office of the Administrator. Washington, DC.
- US Environmental Protection Agency. (2010b). *Valuing Mortality Risk Reductions for Environmental Policy: A White Paper*. SAB Review Draft. Washington, DC.
- US Environmental Protection Agency. (2011). *The Benefits and Costs of the Clean Air Act from 1990 to 2020: Final Report*. Office of Air and Radiation. Washington, DC.
- US Office of Management and Budget. (2003). *Circular A-4*. Washington, DC. Available at: http://www.whitehouse.gov/omb/circulars_a004_a-4#19.
- Viscusi, W. K. (2009). The Devaluation of Life. *Regulation & Governance*, 3(2), 103-127.
- Viscusi, W. Kip. (2011). "Policy Challenges of the Heterogeneity of the Value of Statistical Life." *Published Articles & Papers*. Paper 95. http://research.create.usc.edu/published_papers/95

Viscusi, W. K., & Aldy, J. E. (2003). The value of a statistical life: a critical review of market estimates throughout the world. *Journal of risk and uncertainty*, 27(1), 5-76.

Viscusi, W. K., & Aldy, J. E. (2007). Labor market estimates of the senior discount for the value of statistical life. *Journal of Environmental Economics and Management*, 53(3), 377-392.