RISK, UNCERTAINTY, AND HETEROGENEITY: THREE AND A HALF ESSAYS IN ENERGY AND ENVIRONMENTAL ECONOMICS

By

Asa Watten

A DISSERTATION

Submitted to
Michigan State University
in partial fulfillment of the requirements
for the degree of

Agriculture, Food, and Resource Economics – Doctor of Philosophy

2021

ABSTRACT

RISK, UNCERTAINTY, AND HETEROGENEITY: THREE AND A HALF ESSAYS IN ENERGY AND ENVIRONMENTAL ECONOMICS

By

Asa Watten

This dissertation is on dynamic problems in energy and environmental economics. In it I model optimal subsidies for renewable energy when a regulator cannot commit to maintaining the policy in the future. I estimate the effect of political risk on rooftop solar adoption. I model the trade off between attributes in the production of cars and estimate how technology has changed the relationship between power and fuel economy. I model crop abandonment over a growing season. These dynamic problems are complicated by risk, uncertainty, and heterogeneity.

To hope. A sledgehammer, not a lottery ticket.

ACKNOWLEDGEMENTS

If you wish to make an apple pie from scratch, you must first invent the universe.

—Carl Sagan, *Cosmos.*

I wish to thank everyone who was happy to go on a tangent talking about ideas. In particular Abbie, Aman, Carolina, Eric, Gloria, Ora, Sophia, Soren, and Stephen. I hope we go on many more.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

CHAPTER 1

REBATES VS PERFORMANCE PAYMENTS WITH POLITICAL RISK

When should subsides be given? Should they be doled out over time as encouragement for producing social goods, or should they be given as a lump sum to induce costly investments with the expectation of future social benefits? I develop a model of a regulator choosing a policy with the knowledge that they may lose power. Policy uncertainty is modeled as an alternating Poisson process where policies may be repealed and then reinstated. The model captures key components of renewable energy investment and maintenance, and is adaptable to different technologies. I show that with political uncertainty, the optimal policy is not Pigouvian. Using data from a sample of American utility-scale solar projects, I calibrate key parameters of the model and simulate the relative social benefit of a Pigouvian rebate compared to Pigouvian payments, and find that rebates weakly dominate performance payments for all levels of political risk.

1.1 Introduction

What margin should regulators target for inducing pollution abatement? Pigouvian subsidies and taxes have come to mean inducements on the margin of output and not on capital investment, but Arthur Pigou himself was not clear on this point. Writing on the divergence between social and private benefits, he states:

> It is, however, possible for the State, if it so chooses, to remove the divergence in any field by "extraordinary encouragements" or "extraordinary restraints" *upon investments in that field*. The most obvious forms which these encouragements and restraints may assume are, of course, those of bounties and taxes. (1920, 111; emphasis added)

This endorsement for investment-margin inducements is then later contradicted in *The Economics of Welfare* by examples of corrective measures on the margin of consumption (e.g. alcohol).

In recent years, there has been a turn towards paying for output compared to paying for capital. For example, guidance given by the Climate Investment Funds, a renewable energy fund established by 14 donor countries and implemented by the multilateral banks, includes in their criteria for "smart" subsidies: "[subsidies] are linked to results and performance such as volume of green electricity and related emissions reduction, rather than to technology cost" (Climate Investment Funds 2011, 5). From this guidance, it seems they are concerned that recipients of an upfront grant or rebate will not optimally use or maintain the capital equipment without continued inducement on unit of output.

Another example of this turn to performance payments comes from rooftop solar incentives in the United States. Starting at the turn of this century, rebates for rooftop solar were introduced in many parts of the United States. By the 2010s, however, these rebates were on the decline while subsidies for renewable electricity output replaced them. Figure 1.1 shows the percent of solar panels in Tracking the Sun data. These data are compiled by Lawrence Berkeley National Laboratory (Barbose et al. 2019). It is an incomplete panel of states that received rebates or were in states with Renewable Portfolio Standards, which provide subsidies on electricity. Figure 1.2 shows how these policies have varied across states. Grayed states are not represented in the Tracking the Sun data.[1]

While governments may be concerned about an unreliable investor, investors may be concerned about an unreliable regulator. If one political party is supportive of renewable energy and another is not, then policies passed by the greener party may be repealed or reduced when the political tides turn.

Recent years have provided more than enough reason to support this concern. The Sabin for Climate Change Law at Columbia University (2020) tracked 176 environmental

[1] I also omit California from these data as that state offers a choice between an upfront payment and a performance payment contract.

Figure 1.1: Trends in incentive type

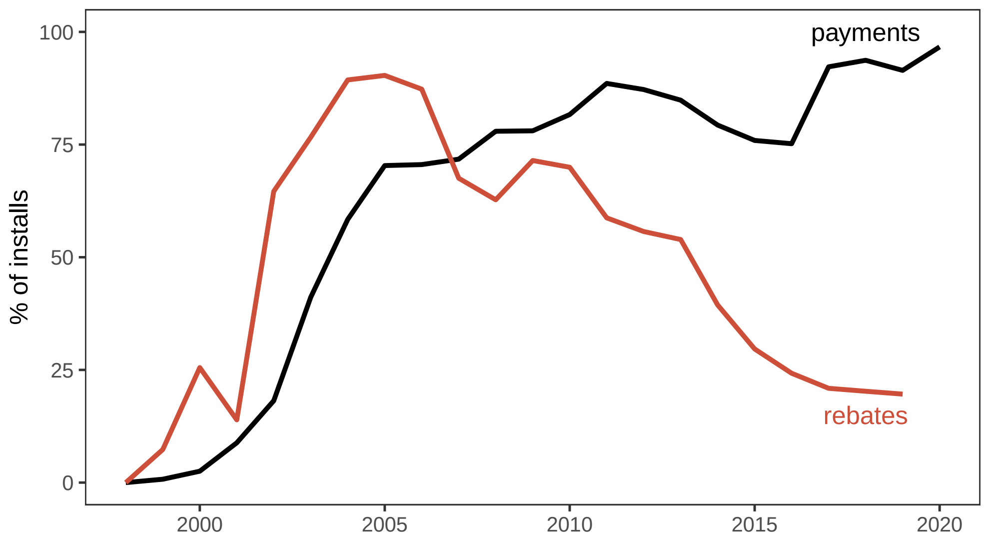

Figure 1.2: Incentives by state over time

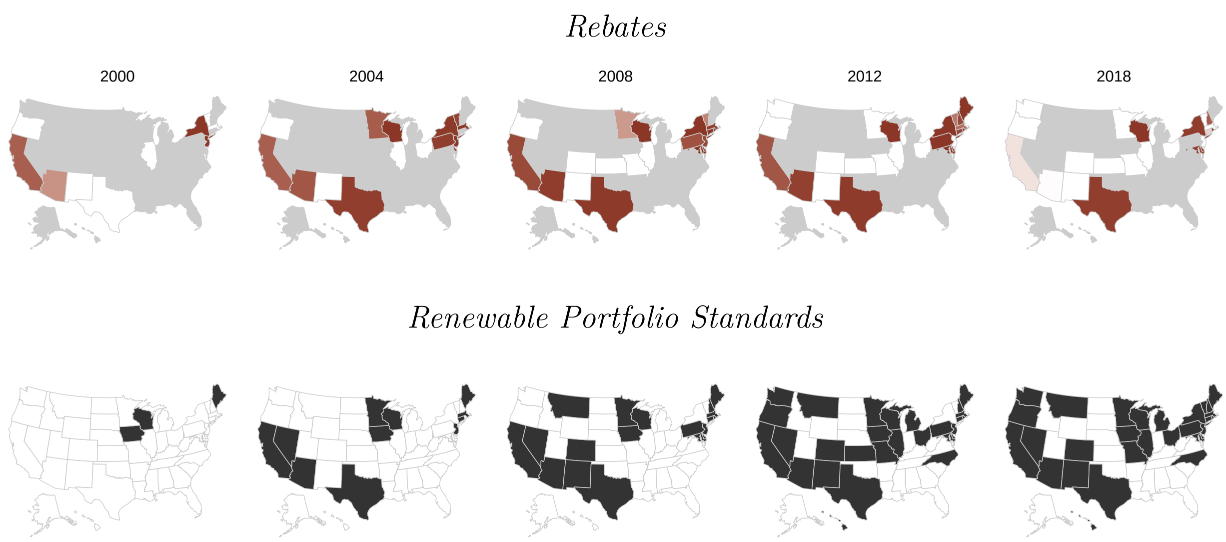

regulations that were repealed or rolled back by the executive branch or Congress from 2017 to 2020. These include the repeal of the Clean Power Plan and rollback of greenhouse gas standards for light-duty vehicles.

This phenomenon is not limited to the most recent American administration. Other states have passed and then repealed energy production policies. Major electricity market reforms enacted by California in 1998 were repealed three years later, leading to the recall of the governor. In an empirical study of the Renewable Fuel Standards, which provides incentives for biofuel production, Lade, Lawell, and Smith (2018) find that announcements of planned policy changes had large effects on traded performance payments (Renewable Identification Numbers). They estimate that the effect of this change in future revenue expectation was to devalue the industry by 6 billion USD.

I model the socially optimal choice of subsidy margin given a stochastic political state—the policy may either be "alive" or "dead." The setting I consider is one where the marginal benefit, maintenance costs, and political risks are known by all parties. Therefore, unlike much of the literature that compares policies for pollution reduction (e.g. Weitzman 1974), there is no cost or benefit uncertainty. For convenience, I call all subsidies on capital "rebates," and all subsidies on output "performance payments" (or simply, "payments"). I show that the choice of policy can be thought of as a choice of the extensive margin (more investments) versus the intensive margin (higher output over the life of a given investment). Rebates act exclusively on the extensive margin; they are not exposed to policy risk, but do not encourage additional maintenance or improvements to increase output. Payments, on the other hand, increase both the intensive and extensive margins, but the effectiveness of payments to encourage either of these margins is eroded by political uncertainty.

In this essay, I first write down a continuous-time model of potential investors in renewable energy projects and a regulator that maximizes social benefits. In Section 1, I solve the model assuming no political risk. In Section 2, I develop the mathematics of an alternating Poisson process to model stochastic political risk. In Section 3, I incorporate a stochastic policy state variable into the model and find three main results: 1) the optimal payment with political uncertainty is larger than marginal benefits; 2) for political risk that is not too big, optimal payments will be lower when paired with optimal rebates compared to a payment exclusive policy; and, 3) Pigouvian payments are preferred to Pigouvian rebates when the price elasticity of per-project benefits is larger than the price elasticity of the number of projects. Hochman and Zilberman (2021) also show that optimal payments under political risk are higher than the Pigouvian level; however, they do not consider rebates or mixed policies. In Section 4, I illustrate the use of the model by estimating model parameters using utility-scale solar data from a sample of American solar projects. I then use these estimated parameters to simulate the model over a range of political risk, and find that Pigouvian rebates are likely to dominate Pigouvian payments in this market. In the final section, I discuss the similarity of the model to the real world target.

1.2 Continuous time problem with a credible regulator

First, let us consider the optimal regulation where there is no risk of the regulator reversing themselves and repealing the regulation.

In the initial period, $t = 0$, producers appraise the expected profit of investment

$$\pi_i = V_i(0) + r - k_i \tag{1.1}$$

where $V_i(0)$, the value function at the initial period, is the expected value of the investment after installation given their optimal behavior in responce to all future conditions; r is the rebate offered by the regulator; and, k is the capital cost. Assume that producers will invest if expected profits are non-negative or $k_i \geq V_i(0) + r_i$. The value function for any time $t \in [0, T]$ is

$$V_i(t) = \max_{c(t)\geq 0} \int_0^{T-t} e^{-\rho t}(pq - c)\ dt \tag{1.2}$$

given the equation of motion

$$\dot{q} = -\alpha + g(c(t)) \tag{1.3}$$

where T is the life of the investment after which there is no bequest value or disposal cost; ρ is the discount rate; p is the payment offered by the regulator for units of production q; and c is the ongoing cost of maintenance, which is chosen by the producer. Changes to output are governed by the equation of motion where α is the rate of depreciation of q and $g(c)$ translates money and effort spent on maintenance to increases in q. This implies that maintenance incurred at any time period increases q in all future time periods, which is eroded at a constant rate α. $q(0)$ is assumed to be known in the initial time period and $\frac{\partial g}{\partial c} \geq 0$ and $\frac{\partial^2 g}{\partial c^2} < 0$, making the problem convex, and $g(c = 0) = 0$.

The regulator chooses rebates (r) and payments (p) to maximize the total social benefit net of costs from all producers

$$T.B. = \max_{p,r} \int_0^{V(p,0)+r} D(k)\Big(B(p,0) - k\Big)dk \tag{1.4}$$

where $D(k)$ is the density of potential producers with capital cost k. The per-project net benefit function is given by

$$B_i(0) = \int_0^T e^{-\rho t}(bq - c(p))\ dt \ \ \text{s.t. } \dot{q} = -\alpha + g(c(p,t)) \tag{1.5}$$

where b the social benefit per unit of output.

The first-order condition of Equation 1.4 with respect to the rebate is

$$[r] : D(k)\Big(B(p,0) - V(p,0) - r\Big) = 0. \tag{1.6}$$

Solving the first-order condition gives the optimal rebate as a function of payments

$$r^*(p) = B(p,0) - V(p,0). \tag{1.7}$$

Setting $p^* = b$, the Pigouvian level of p, makes the producer's problem equal to the regulators (and is therefore the optimal level). This implies the social and private benefit are equal, or $B(p^*,0) = V(p^*,0)$, and therefore $r^* = 0$.

1.3 Policy uncertainty as a two-state Poisson process

To incorporate policy shocks into the model, consider a political process that has two states indexed j: either a regulation is "alive" ($j =$ ☼) or "dead" ($j =$ ☠). The states alternate so that if the regulation is alive it may become dead and if it is dead it may be renewed and become alive again. The duration of each instance of a given state is governed by a stochastic process that is different for each state. These "sojourns" are random variables that form the sequence $\{S_{jn}\} = \{S_{\text{☼}1}, S_{\text{☠}1}, S_{\text{☼}2}, S_{\text{☠}2}, ...\}$, which is a continuous-time Markov process.

Modeling this switching behavior as an alternating (or two-state) Poisson process,[2] given a state j, the arrival rate of the other state $-j$ is λ_{-j}. $\lambda_{\text{☠}} dt$ is then the instantaneous probability of repeal given that the regulation is alive and $\lambda_{\text{☼}} dt$ is the instantaneous probability of enactment given that the regulation is dead. The expected value of each sojourn is $\mathbf{E}S_j = 1/\lambda_{-j}$.

Some additional properties of the alternating Poisson process will be helpful. The probability of being in state ☼ (alive) and ☠ (dead) a time t periods in the future conditional on being in state ☼ now (at $t = 0$) is

$$f_{\text{☼}|\text{☼}}(t) = \frac{\lambda_{\text{☼}} + \lambda_{\text{☠}} e^{-t(\lambda_{\text{☼}}+\lambda_{\text{☠}})}}{\lambda_{\text{☼}} + \lambda_{\text{☠}}} \quad \text{and} \quad f_{\text{☠}|\text{☼}}(t) = \frac{\lambda_{\text{☠}} - \lambda_{\text{☠}} e^{-t(\lambda_{\text{☼}}+\lambda_{\text{☠}})}}{\lambda_{\text{☼}} + \lambda_{\text{☠}}}. \tag{1.8}$$

Similarly, the probability of being in state ☠ and ☼ starting from state ☠ is

$$f_{\text{☠}|\text{☠}}(t) = \frac{\lambda_{\text{☠}} + \lambda_{\text{☼}} e^{-t(\lambda_{\text{☼}}+\lambda_{\text{☠}})}}{\lambda_{\text{☼}} + \lambda_{\text{☠}}} \quad \text{and} \quad f_{\text{☼}|\text{☠}}(t) = \frac{\lambda_{\text{☼}} - \lambda_{\text{☼}} e^{-t(\lambda_{\text{☼}}+\lambda_{\text{☠}})}}{\lambda_{\text{☼}} + \lambda_{\text{☠}}} \tag{1.9}$$

(see Lancaster 1990 or Di Crescenzo and Meoli 2016). In the long run as t becomes large, the probability of being in state j at time t converges, regardless of the initial state, to $\frac{ES_j}{ES_j+ES_{-j}} = \frac{1/\lambda_{-j}}{1/\lambda_{-j}+1/\lambda_j}$ (Zacks 2012).

Four additional relations will be useful for derivations later on:

$$f_{\text{☼}|\text{☠}}(t) - f_{\text{☼}|\text{☼}}(t) = -\, e^{-t(\lambda_{\text{☼}}+\lambda_{\text{☠}})} \tag{1.10}$$

$$\lambda_{\text{☠}} \int_0^x f_{\text{☼}|\text{☠}}(t)\ dt - f_{\text{☼}|\text{☼}}(t) = -\, f_{\text{☠}|\text{☼}}(x) \tag{1.11}$$

$$f_{\text{☼}|\text{☼}}(t) - f_{\text{☼}|\text{☠}}(t) = e^{-t(\lambda_{\text{☼}}+\lambda_{\text{☠}})} \tag{1.12}$$

$$\lambda_{\text{☼}} \int_0^x f_{\text{☼}|\text{☼}}(t)\ dt - f_{\text{☼}|\text{☠}}(t) = f_{\text{☼}|\text{☠}}(x). \tag{1.13}$$

1.4 The continuous-time problem with political uncertainty

Now, consider the same problem as before but with a regulator may repeal and reinstate their regulation at stochastically determined intervals. Since rebates are issued in the initial

[2] Or, more precisely, a process that alternates between two homogeneous Poisson death processes.

period, they are immune from this uncertainty. However, once a producer invests, future payments are only given if the policy is alive and are therefore stochastic. The producer's value function at $t \geq 0$ is now

$$V_{ij}(t) = \max_{c(t)\geq 0} E_t \int_0^{T-t} e^{-\rho t}(\tilde{p}q - c)\ dt \ \text{ s.t. } \dot{q} = -\alpha + g(c(t)) \tag{1.14}$$

where the tilde indicates that payments are stochastic. If the regulation state follows the alternating poison process described above, realized payments at time t are $\tilde{p} = p1(j = \text{☼})$ where 1 is the indicator function. The Hamiltonian-Jacobi-Bellman equation is then

$$\rho V_{ij}(t) = \max_{c(j,t)\geq 0} pq - c(t) + p\int_0^{T-t} f_{\text{☼}|j}(t)\ dt\big(-\alpha + g(c)\big) + \lambda_{-j}\big(V_{-j}(t) - V_j(t)\big), \tag{1.15}$$

subject to the terminal condition

$$V_j(T) = 0. \tag{1.16}$$

The Hamilton-Jacobi-Bellman equation is derived for a generic value function with an alternating Poisson process in Achdou, Han, Lasry, Lions, and Moll 2017. The first-order condition which solves the value function is then

$$[c(j,t)] : -1 + \frac{\partial g}{\partial c} p\Big(\int_0^{T-t} f_{\text{☼}|j}(t)\ dt + \lambda_{-j}\int_0^{T-t} f_{\text{☼}|-j}(t') - f_{\text{☼}|j}(t')\ dt'\Big) = 0 \tag{1.17}$$

as long as $\frac{\partial g}{\partial c}$ is not a function of t and where g_c^{-1} is the inverse function of g_c. The difference in partial derivatives of V with respect to c is the change in $\dot{q}$ for a change in c multiplied by the expected value of $\tilde{p}$. Recall that $\tilde{p}$ is equal to p if the policy is alive and zero otherwise.

For tractability, define the maintenance cost function as

$$g(c) = \beta \ln(c+1). \tag{1.18}$$

The optimal cost response function is then[3]

$$c^*(\text{☼}, t) = p\beta\Big(\int_0^{T-t} f_{\text{☼}|\text{☼}}(t)\ dt - f_{\text{☠}|\text{☼}}(T-t)\Big) - 1 \tag{1.19}$$

[3]The generic form is $c^*(j,t) = g_c^{-1}\Big(\Big[p \times \Big(\int_0^{T-t} f_{\text{☼}|j}(t)\ dt + \lambda_{-j}\int_0^{T-t} f_{\text{☼}|-j}(t') - f_{\text{☼}|j}(t')\ dt'\Big)\Big]^{-1}\Big)$.

and

$$c^*(\text{☠}, t) = p\beta\Big(\int_0^{T-t} f_{\text{☼}|\text{☠}}(t)\ dt + f_{\text{☼}|\text{☠}}(T-t)\Big) - 1 \qquad \text{s.t. } c \geq 0. \tag{1.20}$$

The optimal cost is an increasing function of p and β and is scaled by the cumulative probability that the policy is alive for the remaining life of the investment. Cost is also decreasing in the risk of repeal such that as the risk of repeal ($\lambda_{\text{☠}}$) goes to zero c^* goes to $p\beta - 1$. Increasing or decreasing c has a direct effect on future q which implies that a higher risk of repeal reduces the intensive margin. The positivity condition for c^* creates a kink in the function. I will assume an interior solution for the sake of tractability.

To simplify the regulator's problem, assume that all producers are alike except in their capital cost k_i. This is not an burdensome simplification if we normalize k_i to be the capital cost for a normalized $q(0)$, but it does imply that the cost response function $g(c)$ is alike across producers given this normalization.

The regulator must also now account for uncertainty when choosing a policy. They face producers that respond differently given the state of the stochastic regulation which implies stochastic costs and output, while social benefits of output remain unchanged given the state. The function that determines the socially optimal rebate (Function 1.7) is unchanged by political risk. I show in the Appendix that if the regulator is able to choose $r^*(p^*)$, then the regulator's problem is to maximize the per-project social benefit. The regulator's Hamilton-Jacobi-Bellman equation for individual projects is then

$$\rho B_{ij}(t) = \max_p bq - c^*(p, j, t) + \int_0^{T-t} b\ dt\big(-\alpha + g(c^*(p, j, t))\big) + \lambda_{-j}\big(B_{-j}(t) - B_j(t)\big). \tag{1.21}$$

Taking the derivative with respect to p the first-order condition is solved by

$$p^* = \Big(bT + \lambda_{\text{☠}}\Delta\tilde{c}/\beta\Big)\Big/\Big(\int_0^T f_{\text{☼}|\text{☼}}(t)\ dt - f_{\text{☠}|\text{☼}}(T)\Big). \tag{1.22}$$

which is derived in the Appendix for the proof of the following proposition. Notice that as $\lambda_{\text{☠}}$ goes to zero, the denominator converges to T and $p^* = b$ as in the case without political uncertainty.

I now state and discuss three propositions resulting from the model. Proofs for each of these are provided in the Appendix.

Proposition 1. *If the risk of repeal is positive, then $p^* > b$—Pigouvian payments are not optimal.*

As political risk increases, it decreases the probability that the producer will receive a payment. This increases the payment required to induce producers to engage in maintenance. Further, the producer chosen level of c^*, as shown above diverges depending on state. As the regulator has only one choice of payment, there is no way to induce the first best level of maintenance in all periods for all states. If rebates are optimally chosen, then the optimal payment under political uncertainty is

$$p^* = \Big(bT + \lambda_{\skull} \Delta\tilde{c}/\beta\Big) \Big/ \Big(\int_0^T f_{\sun|\sun}(t)\ dt - f_{\skull|\sun}(T) \Big). \tag{1.23}$$

A corollary is that p^* is not "Pigouvian in expectation." We can see this by recognizing that the expected value of payments in any period from the perspective of the initial period is equal to ${}_{\sun|\sun}(t)p$. Clearly, if we plug in p^*, terms do not cancel to yield b as would be required for a payment that is Pigouvian in expectation. Instead, for reasonable choices of T (not too small) and $\lambda_{\skull}$ (not too big), the denominator is larger than the probability that the policy will be alive in any period. This implies that the expected payment is less than b in any period. This further implies that the optimal, jointly chosen, rebate is positive in most cases.

Proposition 2. *If the regulator is constrained to only set payment levels while rebates are zero, then the optimal second-best payment, p^o, is larger than the optimal payment p^* under the first-best policy that includes both payments and rebates.*

The intuition for this result is straightforward. Without the use of rebates, prices must now control both the intensive and extensive margins. This leads producers to invest in maintenance over the optimal level induced by p^* in order to increase the number of projects towards the first best no-uncertainty level.

Let us turn our attention to comparing Pigouvian mechanisms. What if the regulator may only choose either a price or a rebate and the maximum subsidy level that they may set is the subsidy that is equal to expected social benefit without political uncertainty? There are many instances where the regulator could either be explicitly directed to set marginal subsidies equal to marginal benefit without considering political risk or the political dynamics that drive political risk also prevent the regulator setting prices at the first-best level. The Pigouvian price has the standard definition of setting price equal to marginal benefits—$p^p = b$. The equivalent Pigouvian rebate is to set the rebate equal to the expected benefit conditional on a subsidy price of zero—$r^p = B(p = 0)$.

As producer response on the intensive and extensive margin does not change for a fixed rebate and zero price for changing levels of political risks, r^p is constant for all $\lambda_{\skull}$. It is also unaffected by the maintenance effectiveness parameter β. Because of this, the extensive margin is also unaffected by changes to these parameters. Instead, all of the action occurs only on the performance payment mechanism as risk and maintenance parameters shift. Changes in these parameters affect both the intensive and extensive margins under a price mechanism. It is simple to show that the difference in total social benefit between a price mechanism and a rebate is given by the difference in area of two rectangles

$$TB^p - TB^r = N^p \Delta B - \Delta N B^r \tag{1.24}$$

where N is the total number of adopters, ΔB is the additional benefits per project from prices compared to rebates, and ΔN is the additional number of projects from rebates compared to prices. This is derived in the Appendix. While there is no simple closed form solution to this difference in terms of project level parameters, it can be expanded to an intuitive comparison.

Proposition 3. *If the price elasticity of the intensive margin is larger than the price elasticity of the extensive margin, then the Pigouvian price mechanism is preferable to the Pigouvian rebate. Also, if the price elasticity of the intensive margin is smaller than the price elasticity of the extensive margin, then the Pigouvian price mechanism is not preferable to the Pigouvian rebate.*

Importantly the price elasticity of the intensive margin is determined by the maintenance parameter β as well as the risk parameters. β increases the effectiveness of cost incurred towards maintenance. Therefore higher values of the parameter imply a higher price elasticity, while a zero value implies no effect of price on the intensive margin. Increasing political risk on the other hand reduces the expected benefit for producers which in turn reduces the optimal cost and output levels.

1.5 Estimation and simulation

In this Section I estimate the parameters α and β for utility scale solar and use these parameters to simulate the ratio of total social benefit under Pigouvian taxes to Pigouvian rebates over a prater space that varies the risk of repeal $\lambda_{☠}$. The data I use is the Lawrence Berkeley National Laboratory Utility-Solar Technology trends deta set (LBNL 2020). This data set is a panel of annual electricity generation for a sample of 157 utility solar projects in United States 38 states. In addition to electricity generated, these data include the revenue per unit of energy for each solar project from power the electricity markets, which varies by utility, and from Power Purchase Agreements (PPA), which varies by project. PPAs are contracted price for electricity generated for a fixed term. No projects in the data have expired PPA terms.

Normalizing $b = q(0) = 1$, the remaining parameters of the model are T, $\lambda_{☼}$ and $\lambda_{☠}$, α, and β. When $q(0) = 1$, the maintenance shift parameters β is the percentage point increase on q resulting from a percent increase in c. It is simple to see that if the density function $D(k)$ is a constant, which corresponds to a linear supply curve, the constant drops out of a comparison between mechanisms. Therefore, the density is also normalized to 1.

It can be shown that $\frac{\partial \dot{q}}{\partial p} \times \frac{1}{p} = \beta$. Therefore, I estimate the equation

$$q_t = \delta_{s,t} + (1-\alpha)q_{t-1} + \beta \ln(p) + \varepsilon \tag{1.25}$$

where q_t is normalized by dividing output in all time periods by the output in the first year and p is per megawatt-hour revenue for each project in 2019; $\delta_{s,t}$ are state-year fixed effects. The fixed effects capture other state wide shifts including years with more sunshine and shifts in Renewable Energy Certificate prices. Since power purchase agreements are contracts that are determined before construction, there is no risk of endogeneity for that component of price. So, these estimates may be interpreted as causal as long as no project is large enough to affect equilibrium electricity prices on their own. Parameter point estimates from both models are similar in magnitude.

Table 1.1: Estimation results

	Dependent variable: q_t	
	OLS	*fixed effects*
	(1)	(2)
$\ln p$	0.005^{*}	0.008^{**}
	(0.003)	(0.003)
q_{t-1}	0.660^{***}	0.697^{***}
	(0.026)	(0.026)
Constant	0.297^{***}	
	(0.027)	
Observations	1,538	1,538
R^2	0.302	0.350
Adjusted R^2	0.301	0.275
Residual Std. Error	0.083 (df = 1535)	
F Statistic	332.500^{***} (df = 2; 1535)	370.973^{***} (df = 2; 1379)

Note: *p<0.1; **p<0.05; ***p<0.01

Results for two models are shown in Table 1.1. Model 1 omits state-year fixed effects while Model 2 includes them. I use the point estimate of $\hat{\alpha}$ from the fixed effects estimation to simulate total social benefit for a Pigouvian performance payment and a Pigouvian rebate over a parameter space that varies β and the repeal arrival rate $\lambda_{☠}$. Other parameters of the model are assumed and fixed for the whole parameter space. The reinstatement arrival rate $\lambda_{☼}$ is assumed to be 0.2, which corresponds to an mean time of 5 years to re-enact. The total duration of projects is assumed to be 20 years. I also perform the simulations assuming 40 years as a sensitivity check.

Figure 1.3: Simulated ratio of total social benefits—20 year project life

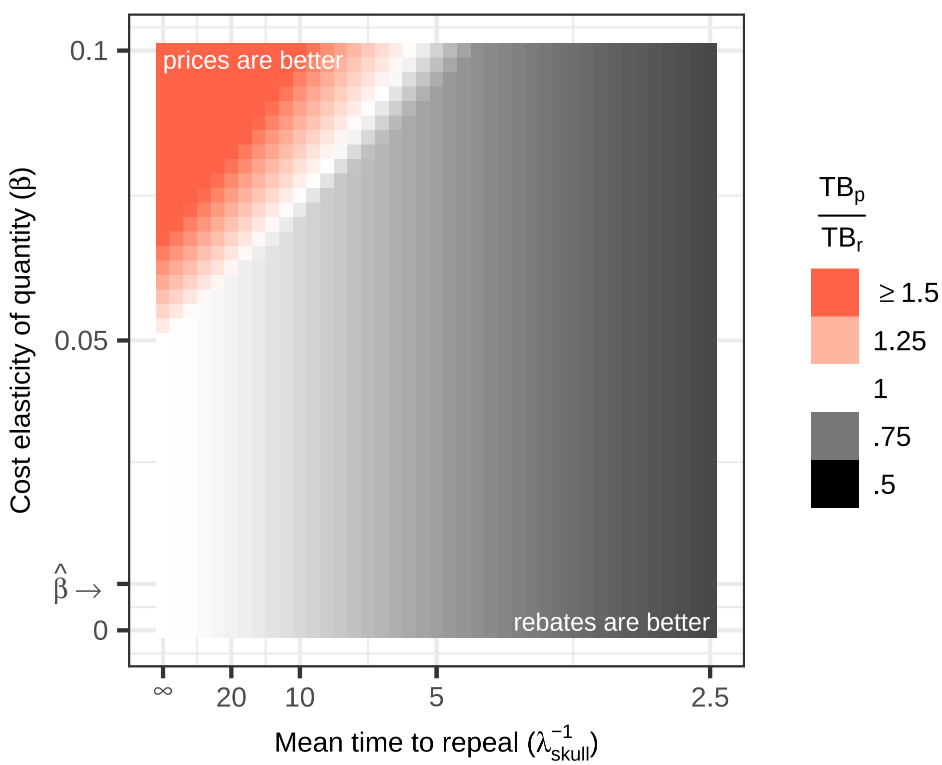

The results of the simulations are shown in Figure 1.3, which assumes a 20 year project life, and 1.4, which assumes a 40 year project life. The estimated cost-quantity elasticity ($\hat{\beta}$) is indicated on the vertical axis. For both sets of simulations, Pigouvian rebates outperform the Pigouvian payments if there is any political uncertainty and are equal to performance

Figure 1.4: Simulated ratio of total social benefits—40 year project life

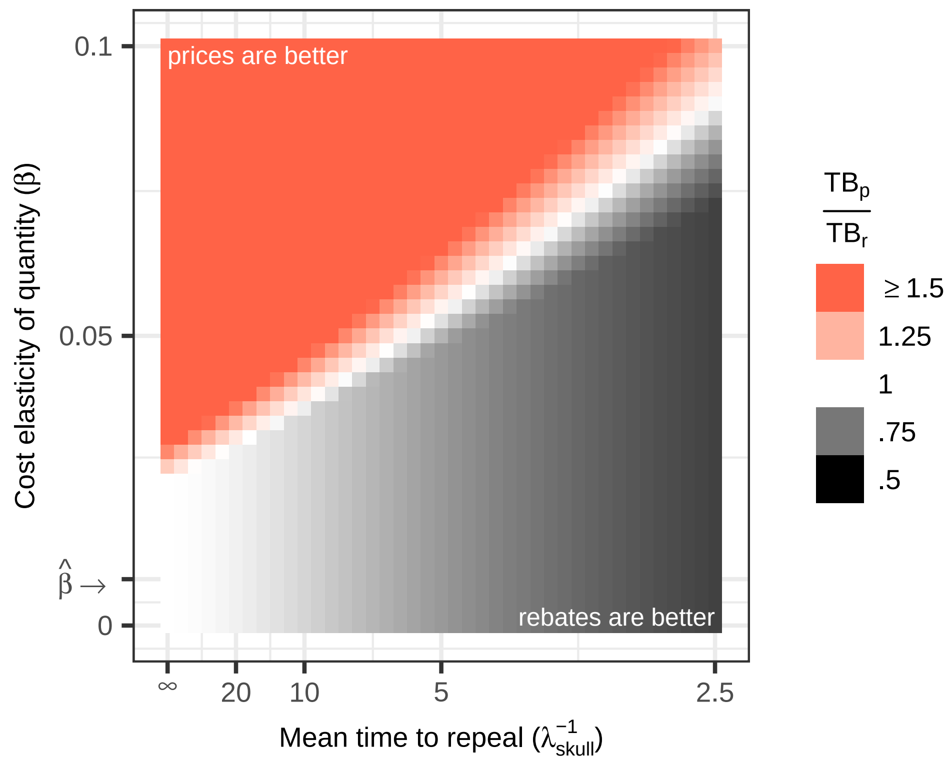

payments if there is no political uncertainty.

1.6 Conclusion and discussion

In this essay I show: 1) the optimal policy deviates from the standard model when there is political risk; 2) the optimal policy depends on the particulars of technology—what level of continued investment is required to maintain the technology; 3) while rebates are never preferred in a setting without political risk, they are increasingly preferable to payments as political risk increases. This effect is intensified as the effectiveness of maintenance expenses decrease.

The model yields two empirically-testable predictions. First, given a positive response in output to maintenance activity, increasing payments should increase output over the life of a project. I test this result for the utility-scale solar market and find small and statistically-significant average effect. It is likely, however, that this effect is spatially heterogeneous as local environmental factors contribute to the rate of solar PV fouling. As such regulators should not take the implications to be the case for all settings. Further, the result is specific to this particular technology. It is often assumed in rooftop solar cost estimation tools that there is nearly zero maintenance cost incurred by homeowners with solar panels. If this is so, then the model would predict that rebates are always preferable to payments to induce rooftop solar adoption. On the other hand, power generated from wind farms requires more maintenance to prevent decline, since wind turbines have large moving parts that can break and are prone to lightning strikes.

The second testable prediction is that political risk reduces adoption rate under a payment regime. This is the subject of the following essay.

Is the model similar enough to the relevant mechanisms in the real world? I discuss four threats to similarity to the target market: heterogeneity of producers, endogenous risks, electricity markets, and the regulator's familiarity with the target technology.

First, producers are assumed to be alike, while in reality they are likely heterogeneous. This concern can be mitigated by normalizing output and capital costs per unit of expected output in the first period. However, this implies that the maintenance response function is also alike across producers for this normalization, which may or may not be the case.

Second, there are valid complaints to make about the use of an alternating Poisson process to model political uncertainty. If it is the case that increasing the size of industry for the targeted technology increases the industry's political influence, then the probability of repeal ought to decline as the policy grows older and has not been repealed. Similarly, a swift repeal may discourage regulators with limited political capital to put effort into reinstating the policy. Or, if the probability of repeal is in fact endogenous to industry size, a collective action problem may emerge among producers, where entry is more attractive given the entry of others. Given these endogeneity concerns, the model is best interpreted as applying to a new industry over the short-to-medium run.

Third, I leave out the market for electricity. In most cases, renewable energy generators will also receive revenue from selling electricity to the grid, or, in the case of rooftop solar, reduce purchased electricity from the grid. By leaving out this source of revenue for investors, the model should be interpreted as excluding inframarginal behavior on either investment margin; it instead captures those projects that would not be built and maintenance that would not be incurred without public incentives.

Fourth, I assume that regulators are good at predicting the social benefit of projects. In many cases, this is a good assumption, provided that administrative and political hurdles can be cleared to assess expected output. As experience with and data on renewable energy technologies increases, so does the ability to predict future output with tools such as Google's Project Sunroof or the National Renewable Energy Laboratory's WIND Tool. A notable exception may be predicted energy savings from building weatherization, which Fowlie, Greenstone, and Wolfram (2018) find to be 2.5 times actual energy savings in a weatherization assistance program in Michigan.

These threats to similarity mainly have the effect of adding nuance to interpreting the model and tempering expectations of precise predictions in a messy world. What the model does do is to capture the effect of large policy shocks on the direction of adoption and maintenance decisions. In the specific case of utility-scale solar, I estimate a price elasticity of output to be very small, which agrees with our understanding of maintenance of these systems, and makes rebates preferable to payments for lower levels of political risk.

CHAPTER 2

POLITICAL RISK REDUCES ADOPTION UNDER RENEWABLE PORTFOLIO STANDARDS

I study the effect of political risk on subsidies for renewable electricity generation. I find robust evidence that political risk reduces the effectiveness of Renewable Energy Certificate (REC) prices using a national opinion survey as a proxy measure of political risk. Renewable Energy Certificates are state-level subsidies on renewable electricity generation that are traded in markets for states with Renewable Portfolio Standards. I use the variation in political risk while controlling for REC prices in linear and Poisson fixed effects models to identify the effect of risk on rooftop solar adoption. The estimated effect is large. Adoption of solar panels among states with Renewable Portfolio Standards could have been 70 to 85% higher if the effect of political risk was reduced.

2.1 Introduction

In the previous essay, I wrote down a model of a regulator designing a policy to induce renewable energy investment and who is at risk of losing power and having their regulation repealed. A prediction of this model was that as the risk of repeal increases, the effectiveness of payments for energy on inducing adoption decreases. In this essay, I test this prediction for rooftop solar adoption in the United States. I use opposition to Renewable Portfolio Standards (RPS) as measured by the Yale Climate Survey (YPCC and Mason 4C 2020) as a proxy of political risk, and use the variation between states to identify a change in the efficacy of Renewable Energy Credits (RECs).

RECs are a state-level subsidy on electricity generated by renewable sources and are traded on a market between suppliers of renewable energy and obligated parties, typically utility companies. As of 2021, 30 states have RPS on their books, which account for 58% of all electricity sales (Barbose 2021). Because RECs are traded, the per kilowatt-hour subsidy varies on this market. Theory predicts that this variation comes from changing expectations about the future marginal suppliers of electricity to the RPS. These expectations can shift due to changing expectation of short- and long-term energy demand, supply, and expected future policy affecting either of those. Therefore, a home or commercial building owner weighing the decision to install solar PV faces two uncertainties arising from RECs: the REC price over the lifetime of the investment, and the risk of repeal of RPS, ending all future REC payments.

Using an unbalanced panel of nearly 20,000 zip codes with 1.7 million PV installations, I test if higher political risk reduces how effective RECs are at inducing adoption, and find a negative and statistically significant effect across a range of specifications. I then simulate counterfactual scenarios for 2018 and the 8 RPS states in my data for that year. I find that the total effect is potentially substantial, as shown in Figure 2.1 and discussed in the Results section.

Figure 2.1: What if opposition was reduced?

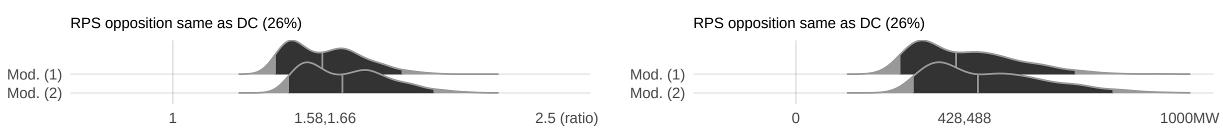

Note: Distributions of 500 bootstrap simulations using baseline specifications. Results are for the following states with RPS: CT, DE, MA, NH, NJ, OH, TX . The figure on the left shows the ratio of predicted adoptions compared to the baseline. The right column shows the additional capacity in megawatts. Inner 95% interval shown dark grey. Medians shown on the density by vertical lines and means are reported underneath. The District of Columbia's level of opposition to Renewable Portfolio Standards is used as the counterfactual because it is the territory with the lowest level of opposition among those with a standard.

In 2018, of non-hydro renewable, 23% of capacity energy was added through RPS (Barbose 2019). If the mean estimate I find were to hold across all states and renewable energy sources, there would have been an additional 15% to 20% of renewable energy added that year.

How likely is it for RPS to be repealed? The Texas legislature in 2007 was the first to try to undermine its RPS by allowing large industrial electricity consumers to opt out. There was an active push in the 2010s to repeal or greatly reduce RPS in all 20 states with RPS at the time (Cullenward and Victor 2020). The political scientist Leah Stokes (2020) has written of the period: "By 2011, a movement to repeal renewable energy laws was underway across the country, led by groups with ties to fossil fuel companiesâmost prominently the American Legislative Exchange Council (ALEC)." This led to repeal of RPS in Kansas, and rollbacks in Maryland and Ohio. Maryland reduced the cost of alternative compliance payments while Ohio first put a two-year freeze on the state's RPS in 2014, and in 2019 Ohio reduced its RPS target from 12.5% in 2026 to 8.5% and eliminated the standard afterwards.

In recent years, many other significant renewable energy policies have been repealed in the United States and abroad. In the United States, the Clean Power Plan was rolled back before it came into force. Australia reversed its carbon tax in 2014 after only two years. Multiple bills have been put forward to repeal Canada's federal carbon tax. The Yellow Vest protest movement in France, in part responding to fuel and carbon taxes, was successful in rolling back taxes on petroleum.

Matisoff and Johnson (2017) also study the effectiveness on non-utility solar incentives in the United States. They use data on the net present value of incentives to compare their effect on adoption, and find that RPS (among other mechanisms) has statistically insignificant effects on adoption. They argue that this result is consistent with research that shows that incentive salience is important. De Groote and Verboven (2016) study household myopia in the face of solar subsidies. Using variation in pre-announced changes in electricity generation subsidies in Belgium, they find that consumers undervalue future payments and that upfront rebates would have cost the government less for the same rate of adoption. While Masitoff and Johnson (2017) show that up-front subsidies are more effective and speculate that the result is driven by salience, and De Groote and Verboven (2016) show that potential adopters discount future subsidy payments and speculate that the result is driven by myopia, this essay contributes to the literature by showing that political risk could be the driving mechanism.

Salience and myopia could be ignored as only affecting consumers with limited attention. Corporations in the business of renewable energy may be less vulnerable or even largely immune to these miscalculations. However, to the extent that corporations are paying more attention to calculating expected future returns, they ought to be *more* sensitive to political risk and volatility, not less. If this is so, then my results have implications not only for solar subsidy programs, but for cap and trade policies and Pigouvian taxes when they are exposed to political risk.

2.2 Data and summary statistics

I use data provided by Lawrence Berkeley National Laboratory (Barbose et al. 2019) of non-utility-scale solar PV installations in the United States. These data are collected by utilities and state and local government agencies that provide incentives. These data include all PV systems that received an incentive for the areas and years in the data. Lawrence Berkeley National Laboratory estimates that the data for 2019 covers 82% of all distributed PV systems in the United States. I use 25 of the 27 states in the data set. I drop Maryland and Utah, which are only in the data for one year. These data include total system cost before incentives, rebates, system capacity, and zip code. A summary of the panel at the state level is provided in Table B.1 in the Appendix.

The population-weighted number of rooftop solar projects is shown in Figure 2.2 for states with RPS and without. California is included as a state without a RPS even though a standard is law in that state. This is because rooftop solar installers are given an option to either receive a rebate or a contracted payment price over a fixed term. This is discussed further in the next section.

I use REC data from two sources. Some states have solar carve-outs in their RPS. This creates a solar REC market exclusively for the purchase of solar PV certificates—so-called

Figure 2.2: PV adoption per 1,000 people

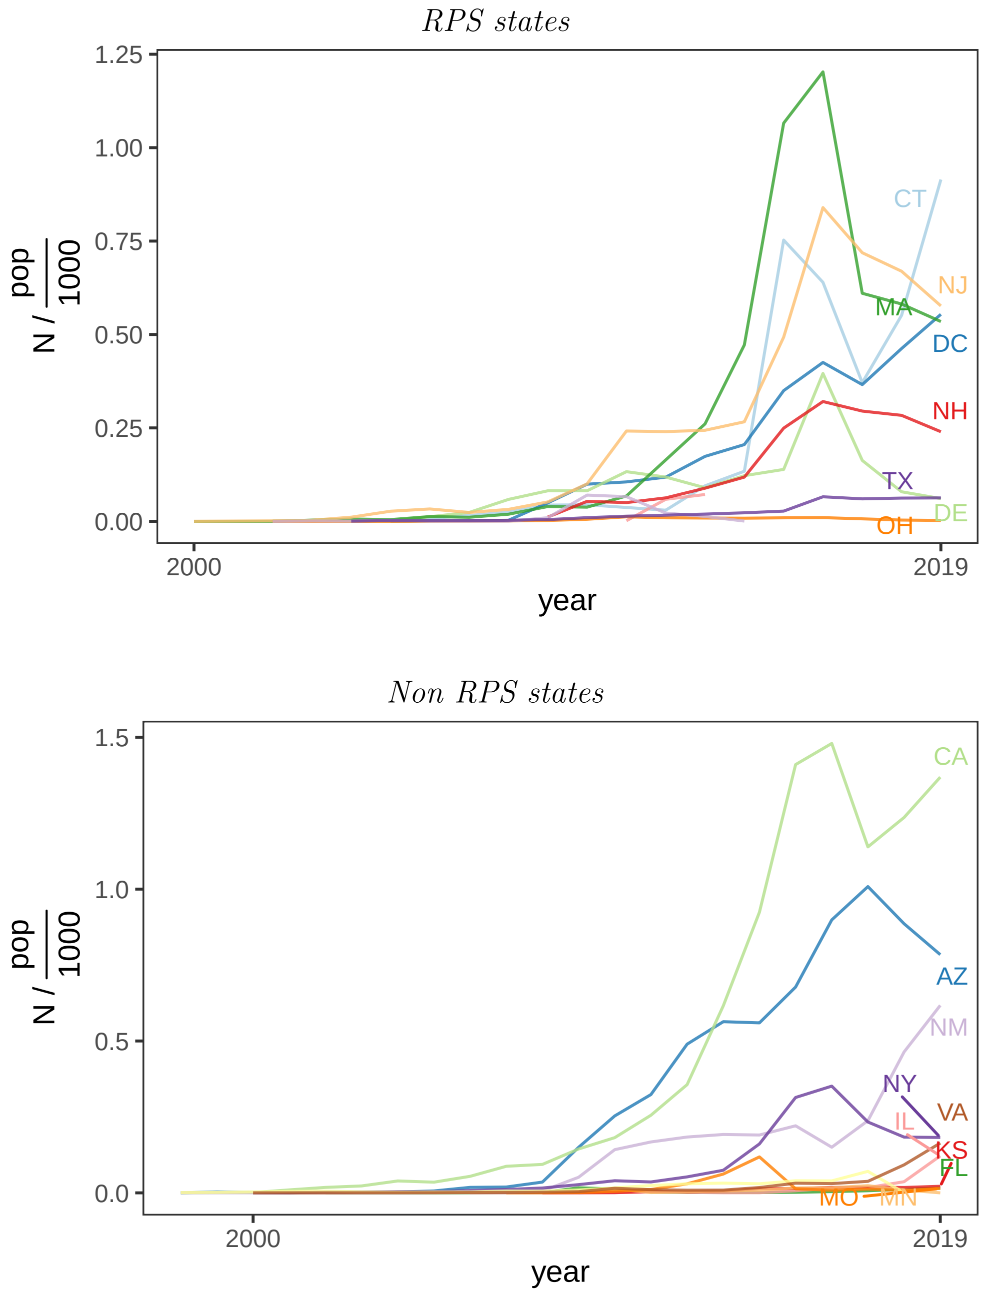

SRECs. I use purchased data from Sun Trade for all states with SRECs. Sun Trade provides daily traded SREC prices. For generic RECs, I scraped prices reported by Department of Energy reports. When there are both generic RECs and SRECs for a state-year, I use

SRECs. Unlike the Renewable Identification Number market where total credit value per gallon of biofuel is the sum of all qualifying credits, electricity providers only sell into the highest tier that they qualify for. After merging the PV data with REC data and dropping state-years where I am missing REC data, I have 25 states that are in the unbalanced panel ranging from 3 to 22 years. Of these, 10 states have a RPS and therefore REC data.

Figure 2.3: REC prices

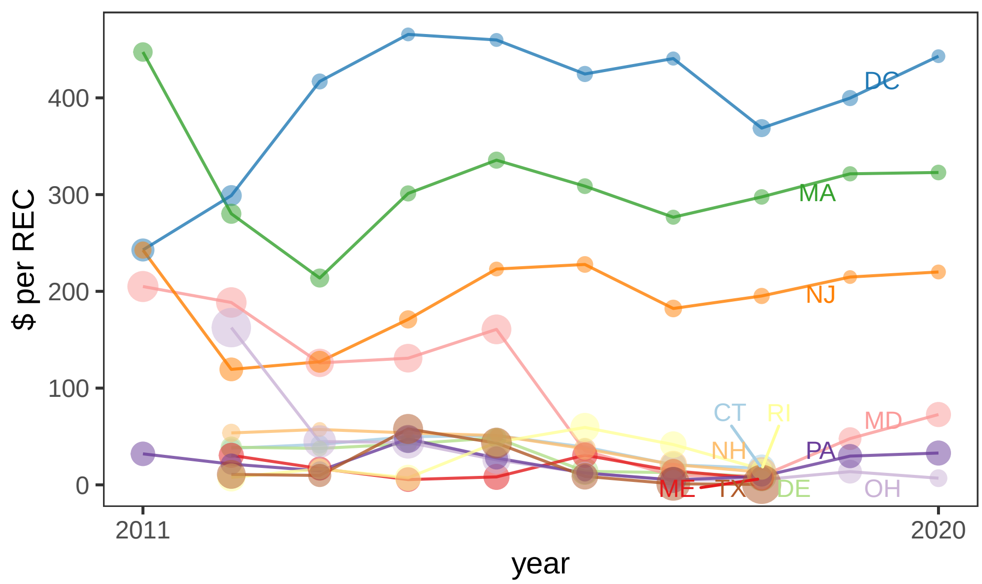

Note: Lines show average Renewable Energy Certificate prices for the highest value qualifying REC for each state in the data; size of dots shows volatility estimated as a GARCH(1,1) process.

For the empirical analysis, I use yearly averaged REC prices by state, and I use the daily traded averages to estimate yearly average volatility for each state using a GARCH(1,1) process. Figure 2.3 shows variation in REC prices and volatility over time. Prices vary substantially between and within state REC markets.

I use data from the Yale Program on Climate Change Communication to measure state and local (county) level opinions on climate policy (YPCCC and Mason 4C 2020, Howe et al 2019). YPCCC conducts an ongoing national survey (n>25,000) of public opinion about climate change and climate change policy. Survey data is reported at both the state and county level. Local opinion data is estimated using multilevel regression with post-stratification that includes demographic and geographic characteristics (Howe et al 2015). The model is validated using multiple cross-validation, comparisons to independent local surveys, and comparisons to data collected by other researchers.

Figure 2.4: Opposition to RPS among states with RPS in data

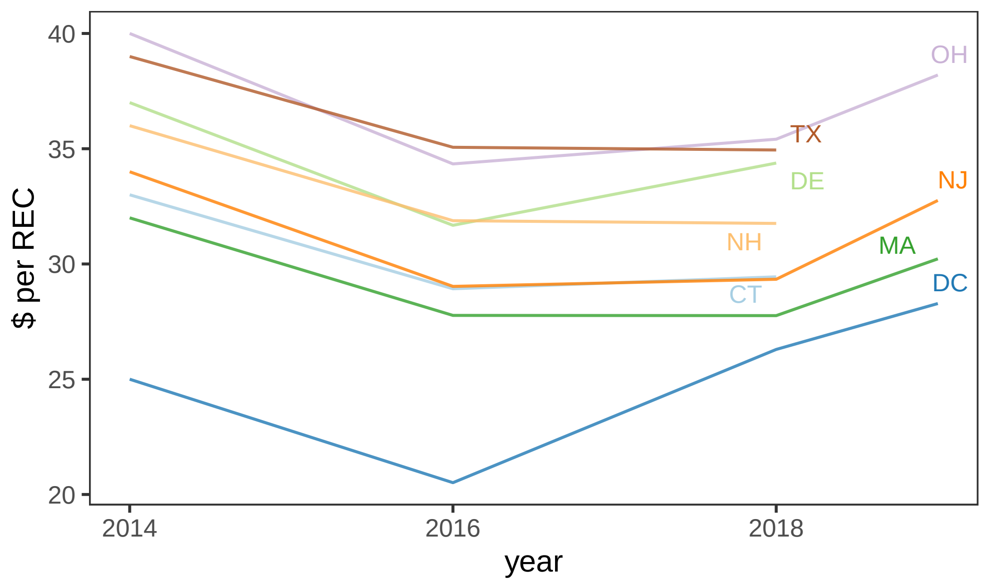

Note: Shown are percent of survey respondents that oppose Renewable Energy Standards for states with Renewable Portfolio Standards for the years they are in the data set; survey data from Yale Climate Survey (YPCC and Mason 4C 2020).

I use all survey years in the range of the rooftop solar data, which are: 2014, 2016, 2018, 2019. For years without surveys, I average across neighboring years with surveys. For years before 2014, I use the 2014 survey. Survey methodology changed somewhat after 2014; the primary change was to include adults ages 18 and over, while only respondents 25 and older were used in 2014. There were also changes to the weighting at the state and county level of respondents after 2014. Assuming that the difference between observations are comparable across years, these changes will not affect the estimated effect of the uninteracted risk term, but could effect the interacted risk term. The survey includes questions on support and opposition for several climate policies such as RPS, government investment in renewables,

Figure 2.5: Support for RPS

States with RPS (left) and without RPS (right)

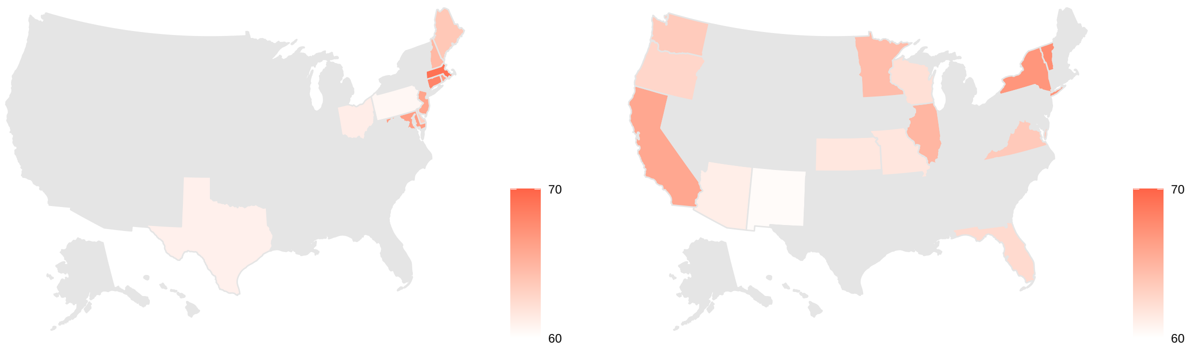

Note: Shown are percent of survey respondents that support Renewable Energy Standards for states in the complete data set; survey data from the Yale Climate Survey (YPCC and Mason 4C 2020).

and a carbon tax. It also asks questions about beliefs (e.g. "global warming is happening"), risk perceptions (e.g. "worried about global warming"), and behaviors (e.g. "citizens should do more to address global warming"). County level data is not reported for 2014. So, I use local data for the most recent year only in the estimation.

Table 2.1: Correlation between RPS support and REC level

Dependent Variable:	$/REC
(Intercept)	-231.4***
	(2.243)
supportRPS	3.843***
	(0.0351)
Observations	264,613
r^2	0.04337

Standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Figure 2.4 shows percent of respondents who opposed RPS over time for states with RPS in my data. The ordering of states is relatively consistent over the time period. While opposition trends upward after 2016 in general, the level of opposition appears to fall from the 2014 to 2016 survey years. It is unknown if this is due to methodological changes or reflects true changes in the population. Figure 2.5 maps state average support for RPS for states with and without RPS for the 2019 survey. Support for RPS varies between about 60 and 70% both for states with and without RPS. Visually, there is an association between RPS states with high support and with higher REC prices. These states tend to have larger standards, which would explain this correlation. I estimate the correlation between RPS support and REC level, which is shown in Table 2.1. On average, a one percentage point increase in support for RPS is associated with a 3.8 USD increase in REC price.

2.3 Empirical specification

I wish to estimate the effect of political risk on the effectiveness of payments (RECs) to induce adoption. As I do not observe the probability of repeal or a large change in the level of payments, I use survey measures instead, and compare states with populations that have higher or lower opposition to the payment policy as measured by the Yale Climate Survey (YPCC and Mason 4C 2020). Using the survey measure to proxy for political risk assumes that there is some function that translates support in the general population to the risk of repeal by the state government. Additionally, it assumes that this function is approximately the same across states. We could imagine several mechanisms that translate opposition or support into changes in law, including elections, lobbying, and financial support of politicians. Support in the general population itself may be a reflection of lobbying and advocacy done by industries and other political groups. Table 2.2 shows the correlation between opposition interacted with RECs and the population-weighted average of rooftop solar projects controlling for REC levels. On average, there are 5.3 rooftop solar projects per 1,000 people in a year without any opposition and for the average REC price (280 USD). If opposition is at 35% then there are only .4 solar projects per 1,000 people per year. However, there are likely other confounders that are not included in this simple correlation that determine both opposition and adoption. This section describes my strategy to control for these confounders.

Table 2.2: Correlation with treatment variables

Dependent Variable:	$N \Big/ \frac{pop.}{1000}$
Variables	
(Intercept)	0.1515***
	(0.0046)
\$/REC × % oppose RPS (state)	-0.05***
	(.0025)
\$/REC	0.0190***
	(0.0008)
Observations	263,202
r^2	0.00558

Standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

To estimate the effect of political risk, I specify two estimators: a linear fixed effects model, and an exponential mean fixed effects model. I estimate these with OLS, and Poisson quasi-maximum likelihood estimation, respectively. Poisson regression is often used in models with dependent variables that are count data; it has some advantages over a log-linear model, such as accepting outcomes that are zeros. The Poisson Fixed Effects (PFE) estimator is consistent under simple assumptions regardless of the distribution of the error term. Unfortunately, the Poisson fixed effects estimator is not suitable for simulation. This is because all zip codes with no installations in the data are perfectly predicted by the multiplicative Poisson fixed-effect, and the estimator cannot predict changes for alternative scenarios.

The general form of the linear model that I use is:

$$\frac{N_{zt}}{\text{pop}_z/1000} = \alpha_z + \gamma_1 \text{risk}_s \text{rec}_{st} + \gamma_2 \text{rec}_{st} + \gamma_3 \text{volatility}_{st} + \mathbf{X}_{zt}\beta + \varepsilon_{zt} \tag{2.1}$$

where $\frac{N_{zt}}{\text{pop}_z/1000}$ is the PV adoptions per 1,000 people in zip code z and year t; α_z is the time invariant zip code level fixed effect; risk_s is the probability of RPS being repealed for state s; rec_{st} is the annual mean of traded REC prices and volatility_{st} for state s in year t. I estimate REC volatility for each state using a GARCH(1,1) process. $\mathbf{X}_{zt}$ is a vector of additional time-varying controls for each zip code.

In this model, zip code fixed effects control for all time-invariant characteristics of a given zip code. Of particular concern are characteristics that influence solar adoption, such as solar potential, characteristics of the housing stock such as roof type, and the tastes and environmental concern of the local population. Lags of the dependent variable are not included as they are known to cause bias and are not consistent as the number observational units becomes large so long as the number of time periods remains small (Nickell 1981).

Similarly, I specify the exponential mean model as:

$$\frac{N_{zt}}{\text{pop}_z/1000} = \phi_z \exp\left(\gamma_1 \text{risk}_s \text{rec}_{st} + \gamma_2 \text{rec}_{st} + \gamma_3 \text{volatility}_{st} + \mathbf{X}_{zt}\beta + \rho \frac{N_{zt-1}}{\text{pop}_z/1000}\right) u_{it} \tag{2.2}$$

where ϕ_z is the unobserved (fixed effect) parameter for zip code z and u_{it} is mean 1 conditional on covariates. Wooldrige (1997) shows that the PFE is consistent as long as the expected value of the multiplicative error u_{it} is 1, conditional on covariates for time periods preceding and including t (the so-called Conditional Moment Condition). This allows the researcher to include lags of the dependent variable unlike in the linear fixed effects case. I include lags in the Poisson model to improve the efficiency of the estimator. It is likely that exposure to rooftop solar and marketing from solar panel installers contributes to installations in future periods. I choose the number of lags that minimizes the Schwarz-Bayesian Information Criterion (Scott and Hatemi-J 2008), which I find to be one lag. Figure B.5 in the Appendix shows the results of the test. Lags in the model control for aspects of the local solar industry such as marketing or the effect of seeing neighbors with solar panels.

In both models, the main coefficient of interest is γ_1, which is the coefficient on the interaction between the level of REC prices and a state's political risk. The interaction term takes the value zero for all non-RPS states. A negative value γ_1 indicates that the value of the RECs is reduced by risk. In the linear model, $\gamma_1 \times \text{risk}_s \times \text{rec}_{st}$ has the interpretation of being the total effect of political risk in state s in year t on PV installations per thousand people while in the Poisson model it is the total percent change in installations.

The cost of PV installations per unit of capacity has fallen substantially. These trends may vary regionally due to regional demand for labor. However, including local rooftop solar prices runs the risk of introducing simultaneous equations bias. Instead, I use utility-scale solar without storage averaged at the state level as a proxy for rooftop solar costs to control for regional supply shocks. This assumes that demand at the zip code level for solar installations has a negligible effect on state average utility-scale solar costs.

Political change could also affect potential adopters in states without RPS. If adopters in non-RPS states think it is highly likely that their state will adopt an RPS in the near future, they may be more likely to adopt, conditional on the cost of adoption. This anticipation may also induce them to wait to see if policy changes. To control for this, I include the survey response "support for RPS" for all states. I also include the same response interacted with the cost proxy and an indicator if the state is a non-RPS state.

While trends that are common across all observations are controlled for with year fixed effects, we may be concerned that there are differing trends that are correlated with adoption and the zip code fixed effect. I interact county level concern for climate change with time to control for differing time trends for locations with higher or lower environmental awareness.

There are a number of state and local policies that affect PV adoption that are not captured by rebate or REC data and vary over the panel. All of these time-varying policies that I am aware of and that are not otherwise captured by other controls are in California. In 2011, California began offering PV adopters a choice between a five-year fixed contract for electricity and an upfront rebate. There is no way to tell from the data who took the upfront cost and who took the contract for electricity, so this policy is controlled for with a dummy variable. I include additional dummy variables to indicate changes in California statewide zoning requirements related to PVs for years 2007 to 2014, 2014 to 2019, and 2020. The most recent of these changes requires all new homes built in California to have rooftop solar. I also include city-specific policies for Los Angeles (2013 to 2015) and San Francisco (2015 onward) which offered feed-in tariffs and new construction requirements respectively, as these changes are not captured by zip code fixed effects.

2.4 Results

Table 2.3 shows results for the main specifications for OLS and Poisson regressions. Models 1 and 3 use levels of the risk variable for the explanatory variable, while Models 2 and 4 use the natural log of risk.

The coefficient of interest is estimated to be negative and significant at the 5% level in Model 1 and at the 1% level for all four models, indicating a reduction in REC efficacy with political risk. As we would expect, the coefficient on the REC level is positive in all four models, indicating that higher payments induce more adoption. The coefficients on the cost proxy and the cost proxy interacted with state support for RPS are small and not significant across models. The coefficient on rebates is also small and not significant. The coefficient on rebates is perhaps not surprising. The bulk of rooftop solar rebates are provided at the federal level and therefore do not vary across observations. Changes to local incentives may be less salient or require additional administrative hurdles. Counter to what theory would predict, the coefficient on REC volatility is positive and significant across all models.

The coefficient on county-level environmental concern interacted with time is negative. This could be explained by locations with high environmental awareness adopting rooftop solar sooner than those with lower environmental awareness.

The fit of the Poisson estimator is substantially better than the fit of the linear model (0.92 compared to 0.71). This is also illustrated in Figures B.4 and B.3 in the Appendix, which shows fitted values versus predicted values as well as residuals for Models 1 and 3. While the Poisson estimator does a better job at fitting the model, it appears to under-predict observations with very high numbers of installations.

Table 2.3: Estimates for main specifications

Dependent Variable:	$N\Big/\frac{pop}{1000}$			
Model:	(1)—OLS	(2)—OLS	(3)—Poisson	(4)—Poisson
\$/REC × % oppose RPS (state)	-0.11**		-0.04***	
	(0.04)		(0.01)	
\$/REC × ln % oppose RPS (state)		-0.0330***		-0.0123***
		(0.0126)		(0.0030)
\$/REC	0.0384***	0.1189***	0.0155***	0.0452***
	(0.0144)	(0.0442)	(0.0034)	(0.0104)
REC volatility	0.2500**	0.2415**	0.4533***	0.4525***
	(0.1070)	(0.0974)	(0.0513)	(0.0522)
Rebate per MW cap.	0.0003	0.0003	0.0008	0.0008
	(0.0003)	(0.0003)	(0.0005)	(0.0005)
Cost proxy	0.0006	0.0006	0.0019	0.0019
supportRPS	-0.1112	-0.1061	0.0491	0.0517
Cost × support	1.51×10^{-6}	2.46×10^{-6}	-5.57×10^{-6}	-5.2×10^{-6}
CA policy 1	-0.3980***	-0.3927***	0.0532	0.0516
CA policy 2	0.4436***	0.4526***	0.4943***	0.4949***
LA policy	0.0349	0.0340	0.0071	0.0068
SF policy	-0.7649***	-0.7638***	-0.6877***	-0.6876***
t × % worried about global warming (county)	$-1.7 \times 10^{-7***}$	$-1.7 \times 10^{-7***}$	-7.15×10^{-8}	-7.07×10^{-8}
Lag dependent var.			0.0684***	0.0683***
			(0.0135)	(0.0136)
Observations	80,790	80,790	49,842	49,842
R^2	0.71120	0.71126	0.92007	0.92009

*Note: Standard-errors clustered by state code in parentheses. Signif. codes: ***: 0.01, **: 0.05, *: 0.1. Some standard errors suppressed for brevity.*

I conduct a series of robustness checks with alternative models. Figures B and B show estimates for the coefficient of interest, γ_1, in alternative specifications of the model, and 95% confidence intervals for these alternative models. The coefficients and confidence intervals are scaled by the average REC price in the data (280 USD) and by 10%. With this scaling, the estimates from the linear estimator can be interpreted as the effect adoptions per 1,000 people for a 10 percent increase in opposition for the average REC level. The scaled coefficients for the Poisson estimator can be interpreted as a percent change arising from the same.

Changes to the model shown are the same for both estimation strategies, and are divided into three groups. The first group is the baseline; the second group uses residential rooftop solar installations only—dropping commercial solar installations; while the third group drops all observations in California. This is done to test if the results are driven by a small number of commercial installations, or by California, which is by far the largest adopter of solar panels both per capita and in total. Within each group, the treatment and controls are varied. The treatment is varied using the level of the risk measure and the log. The estimated coefficient of interest is remarkably consistent across groups.

In one specification, I use the survey response "oppose funding of renewables" instead of "oppose Renewable Portfolio Standards." The coefficient is negative but much smaller. This is possible evidence that potential solar adopters are responding to perceived risk for the specific policy, not general attitudes towards green investments. I also vary the local green preference time trend. While in the main specification I interact years with the response "worried about climate change," in others I use "citizens should do more to fight climate change." I also use the share of the county level vote for the democratic presidential election for the closest presidential election year. Changing these variables have negligible effect on the estimated coefficient of interest.

Using the linear Models 1 and 2, I bootstrap (n=500) counterfactual scenarios for the eight states in the data with RPS in 2018—the most recent year in my data with many RPS states. The results are shown in Figure 2.1. I use opposition to RPS for the District of

Columbia (26%) as the counterfactual level of low opposition. As shown in Figure 2.4, DC has the lowest level of opposition to RPS among sub-national governments with RPS. These bootstraps estimate that with opposition similar to that of DC, we would have seen 73% more solar panel installations for states with RPS in 2018 using Model 1 and 86% using Model 2. Using the average capacity of solar panels in each state, I compute the counterfactual addition in capacity for these states, which is estimated to be 171 megawatts in Model 1 and 201 megawatts in Model 2.

2.5 Conclusion

This essay estimates the effect of political risk on an incentive that pays out over the working life of a capital investment. In doing so, it tests the prediction of the theoretical model in the previous essay, and finds strong support for the prediction that political risk leads to reduced renewable energy adoption conditional on the size of payments for renewable energy.

This study has important implications for policy makers. If policy makers wish to induce renewable energy generation, they ought to consider the risk that they may no longer be in power at some point in the future, and how this risk affects the mechanisms they choose.

CHAPTER 3

ATTRIBUTE PRODUCTION AND TECHNICAL CHANGE: RETHINKING THE PERFORMANCE AND FUEL ECONOMY TRADE-OFF FOR LIGHT-DUTY VEHICLES

With Gloria Helfand and Soren T. Anderson

We develop a theory of attribute production for multi-attribute (differentiated) goods and apply our theory to the production of performance and fuel economy in gasoline-powered cars and light-duty trucks. We investigate the heterogeneous effects across the marketed vehicle fleet for three types of technology changeââattribute-neutral, biased, and discreteââas well as a change in regulatory standards. We find that new technologies are not applied evenly across the fleet and that fuel economy standards tend to induce the largest improvements in the worst performing vehicles. We coin the term Marginal Rate of Attribute Substitution (MRAS), a measure of attribute trade-offs, and demonstrate how to estimate it as the equilibrium of heterogeneous consumer preferences and the cost of attribute production. Using attribute data from marketed vehicles we estimate a reduction in MRAS for power and fuel consumption of about 20% to 50%. This is done by adapting an Expectations-Maximization algorithm to our theory model. This result is robust to parameter assumptions and implies that improving fuel economy by reducing power has become less attractive relative to adding fuel saving technology.

Often the goods we consume are made up of many attributes. Many of these goods use energy, which regulators may want to limit. Houses, appliances, power plants, and cars are all examples of multi-attribute (or differentiated) goods that include the attribute of energy-efficiency or energy consumption in addition to other attributes such as natural light, storage capacity, reliability, or speed.[1] However, economics has not yet formalized a theory of the production of attributes in a competitive setting nor the role of technical change in attribute production. Such a theory is crucial to understanding the effect of regulation on marketed attributes, welfare, and incidence.

We borrow the term âdifferentiated goodsâ to describe our setting. However, unlike the two literatures that have used this term—monopolistic competition (e.g., Chamberlin 1949) and the hedonics literature (e.g. Rosen, 1974)—our interest lies in how producers, in a competitive market, respond to heterogeneous consumer preferences. We also take inspiration from the stochastic frontier estimation literature (Ainger, Lovell, and Schmidt 1977), which formulates the production frontier of multiple goods within a firm. However, our model departs substantially from this literature in order to capture the physical interaction of attributes within a single product.

In the 1980s and 90s mandated fuel economy increases at first came at the expense of performance and later size at the fleet-wide level (EPA 2019). This led to a belief that making vehicles more efficient would necessarily reduce their performance and to the conclusion that regulation to increase the efficiency of vehicles would force consumers into cars that are less enjoyable to drive. Figure 3.1 shows how attribute levels changed in the fleet from 1988 to 2019. Horsepower and fuel economy moved in opposite directions in the 1990s to the early 2000s, but both increased in the second half of the 2000s.[2] We ask: how has the relationship between performance and fuel economy changed, what technologies drive this change, and how can changes in technology and regulation affect the composition of the vehicle fleet?

We define an attribute trade-off as the minimum reduction of one attribute required to increase another by one unit without affecting the per-good production cost of the differen-

[1]Further examples: buildings have attributes such as natural light, interior volume, and energy consumption, each which may affect the production of the other. Desirable qualities in power plants include longevity, flexibility, and capacity. Energy consuming processes of a computer such as high definition graphics conflict with battery life. Plant breeding or production may be directed towards disease resistance, drought hardiness, flavor, or size. City property may be designated for different public uses.

[2]The apparent trade-off between power and fuel economy predates our data by at least a decade (EPA 2019).

Figure 3.1: Percent change in mean attributes (sales weighted)

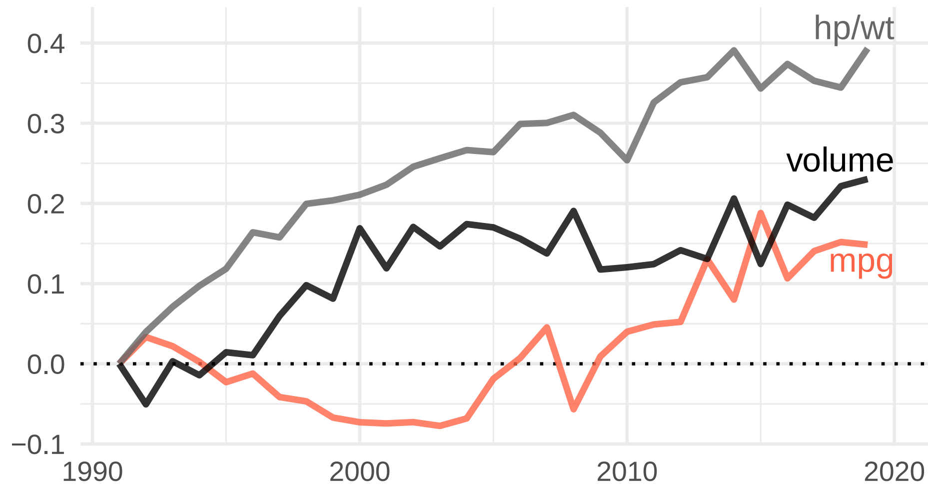

Note: Data are sub-model attributes as reported by Wards Automotive and are sales weighted using EPA administrative data.

tiated good. We coin the term Marginal Rate of Attribute Substitution (MRAS) to refer to this trade-off. MRAS is the slope, at a given combination of attribute levels, along the set of all attribute level combinations for a fixed costââthe slope along an iso-cost curve plotted in attribute space. Each iso-cost curve is a slice of a special type of Production Possibilities Frontier (PPF) for a single unit of a differentiated good. This PPF takes a single input, cost, which aggregates all factors of production, and gives the set of all combinations of attributes such that no attribute may be improved without worsening another.

MRAS is the same concept as the marginal rate of transformation but for a differentiated good. It is meant to emphasise the differences of the PPF at the level of a single unit of a good with attributes as the output, compared to a PPF at the level of a firm or of an economy with multiple factors of production and many products. While a PPF is typically understood to describe the relation of separate inputs and outputs (goods), in our setting, these outputs are not separable from the differentiated good which is purchased by consumers as a bundle. An additional difference is that firms are not constrained to a single attribute production vector. They can make a single model of a good, each with the same attribute production vector, or they may make many models, each with a distinct attribute production vector. Firms produce different models of a good with different attribute levels in response to consumers who differ in their enjoyment of marginal increases in the attributes. As this work is focused on the trade-off between power and fuel economy, future discussions of trade-offs or MRAS will refer to this particular trade-off in the design of light-duty vehicles.

MRAS is pivotal for understanding how differentiated goods—in our case the fleet of marketed light-duty vehicles—change due to technical change or constraints on one of the attributes such as fuel economy regulation. We define three types of technical change and analyze their implications in the theory section, which include: 1) an overall reduction of marginal (per-vehicle) production costs by a scalar; 2) a reduction in production costs that favors one attribute over another; and, 3) the invention of a new technology that is either adopted or not. For the last type, we use turbochargers as an example and show that vehicle weight is correlated with turbocharger adoption despite no direct relationship between turbochargers and weight, as our model would predict.

For regulatory analysis, MRAS is necessary, but not sufficient, for understanding how the distribution of attributes in a fleet will change in response to attribute regulation. If the distribution of consumer preferences is known, then MRAS may be used to compute an upper bound on the cost of regulation. Alternatively we may fix attribute levels and vehicle-specific demand and only allow the addition of costly technologies to meet a standard. This is the approach commonly used in regulatory analysis—EPA (2009), for example. These two approaches can be thought of as estimating corner solutions for manufacturers. Manufactures in practice can both change a vehicle's attribute mix and cost. So, the cost of regulation is bounded above by whichever is the smaller of the two constrained approaches. Whitefoot, Fowlie, and Skerlos (**?**) also makes this point: â[..] analyses of the regulations that do not include [fuel-economy-acceleration] trade-offs may significantly overestimate compliance costs and underestimate GHG emission reductions.â

There are other methods for determining the cost of attribute regulation. Anderson and Sallee (2011) use the observed cost of an alternative compliance mechanism to estimate the shadow cost of regulation. Shadow cost has the benefit of incorporating the cost both of production and the opportunity cost to consumers. Whitefoot, Fowlie, and Skerlos (2017) instead estimate a complete demand system using consumer and engineering cost software for consumers and producers to simulate counterfactuals. We do not make use of this strategy as we are interested in the cost function over time, and it may not always be possible to reconstruct cost simulations of vehicles from one or two decades in the past. Further, it is useful to develop a methodology that relies on market data alone.

In our application to the vehicle market, we are interested in how MRAS and costs have changed over time to learn about the trajectory of technology. This interest is inspired by observations of engineers that suggest that the MRAS of fuel-economy and acceleration has declined substantially. Others have been interested in measuring how technology in the form of production costs have changed over time (Knittel 2011) including how regulation induces innovation (Whitefoot, Fowlie, and Skerlos 2017; Newell, Jaffe, and Stavins 1999).

Economists are familiar with several common production functions such as the Constant Elasticity of Substitution set of functions including the special case of the Cobb-Douglass production function. These are functions that take factors of production (e.g. labor and capital) and map them to the output of a single good. Some research attempts to use these functions, the Cobb-Douglass production function in particular, and re-write them with attributes modeled as inputs and costs as outputs. Unfortunately this approach fails for several reasons. It violates the "no free lunch" principle of PPFs; assuming free-disposal, the PPF is a non-convex set and so a unique solution for consumers maximizing utility or firms maximizing profits is not guaranteed; further, it violates a common sense understanding of attribute cost. Knittel (2011), MacKenzie and Heywood (2012), and Klier and Linn (2012) interpret estimates from a log-log regression on attributes as tracing the mean iso-cost curve with the slope estimating the elasticity of the attribute trade-off. As we discuss in Section 2, the micro-economic foundations of the log-log relationship modeled by many of studies may not be consistent with utility-maximizing consumers. In estimating MRAS with attribute data for marketed vehicles alone, the researcher must make assumptions about the joint distribution of consumer preferences, or equivalently must assume that the portion of the cost function that is determined by a collection of related attributes is constant on average for the entire range of any one of the attributes.

The paper proceeds in four sections. In Section 1, we develop an augmented Cobb-Douglas-like functional form for a single-input, multiple-attribute-output PPF. We find an analytic solution for attribute levels when paired with linear consumer utility functions. Augmenting the Cobb-Douglass function is necessary to satisfy second order conditions. We develop a parsimonious model of attribute production with competitive firms and heterogeneous consumers. We then probe the model for economic intuition and conduct comparative static analysis that is informed by our understanding of how vehicle technology has changed.

In Section 2, we map the physical determinants of horsepower and fuel economy to illustrate the engineering design decisions that produce performance and fuel economy and therefore the MRAS. This exercise demonstrates the importance of engine displacement—the total volume of an engine's combustion chambers typically measured in liters—to the MRAS in vehicles. Engine displacement is the primary design choice that reduces one of our two attributes of interest (acceleration and fuel economy) at the expense of the other.

We then investigate how the relationships between displacement and performance and displacement and fuel economy have changed over time. In 6 simulated vehicles, each with drive-trains representative of technology from a particular year, we show that, given a fixed engine size, vehicles have become much faster and more efficient, and that changes in engine displacement have a substantially larger proportional effect on performance than on fuel economy. If engine displacement were the only source of a attribute trade-off, these simulations provide evidence for show about a 30% decrease in MRAS since the 1980s.

In Section 3, we use market data from Wards Automotive to estimate our cost function.

In lieu of estimating a demand model for consumers, we assume mean preferences for a multi-variate distribution of consumer preferences. We then fit our model and the covariance matrix of the preference parameters using an Expectation Maximization (EM) algorithm for an early time period (1990-94) and a later one (2015-19). We find a consistent estimate of 25 to 28% reduction in the MRAS elasticity. This reduction suggests that, holding costs fixed, increasing performance has a much smaller effect on reducing fuel economy than 20 years earlier. Or, equivalently, holding cost fixed, increasing fuel economy decreases performance much more than it did 20 years ago. This suggests that adding fuel saving technology has become more attractive for meeting regulation relative to making slower vehicles. Section 4 concludes.

3.1 Theoretic framework

We model the relationship between performance, fuel economy, costs and consumer preferences for vehicle attributes that results in an individual vehicle model attribute level. Our model includes profit-maximizing producers who balance per-vehicle manufacturing cost with marginal attribute benefits for distinct consumer segments, modeled as individual consumers. The primary innovation of our model is a Cobb-Douglas-like function that maps attribute levels for physically interacting attributes to per-vehicle production cost. The functional form maintains helpful aspects of a Cobb-Douglas functional form, such as constant elasticity of MRAS, as well as analytic solution when consumer preferences are locally linear. After solving our model for manufacturers and consumers, we compare the effect of different types of technical change and of regulation on attributes for each consumer group.

Model

Manufacturers

The types of attributes we have in mind are those where it is always preferable to increase the attribute (or decrease the attribute depending on the measure of the attribute) within a differentiated good. In contrast, our model does not take into account attributes such as color which are only a matter of taste—some may get utility from a red car while others dis-utility—and which seem unlikely to physically interact in the production of other attributes.

We make two simplifying assumptions. First, we assume that manufacturers are price takers. Second, we assume that there is free entry in the vehicle market implying that the fixed cost of developing a new vehicle is not substantial. Both of these assumptions are likely not correct but are useful for our parsimonious model and for comparison to other work.[3]

Assume that manufactures have locally constant returns to scale so that the per-vehicle cost of production is constant with respect to quantity. The manufacturer's problem is to maximize the profits across all vehicle models by choosing the attributes of all individual vehicle models indexed by l.

$$\max_{q_l, g_l, X_l, Z_l} \pi = \sum_l q_l(p_l - c(g_l, X_l, Z_l)), \tag{3.1}$$

$$c(g_l, X_l, Z_l) = c^1(g_l, X_l) + c^2(Z_l) \tag{3.2}$$

where q is quantity sold, p_l is the price of the model and c is per-vehicle costs. Per-vehicle costs are represented as additively separable functions: a function, c^1, of attributes (X_l) that affect fuel consumption (g_l)—e.g. volume, weight, and acceleration—and a function, c^2, of attributes that are unrelated to fuel consumption Z_l—e.g. the vehicle's sound system, interior design, etc.[4]

If the manufacturer is subject to a fuel economy or greenhouse gas emissions standard, the firm's maximization problem in Equation 3.1 is subject to the constraint

$$\frac{1}{Q}\sum_l q_l \cdot g_l \leq \sigma \tag{3.3}$$

where g_l is fuel consumption in gallons per mile for vehicle model l, σ is the level of the standard, and Q is the sum of all vehicle sales by the manufacturer.[5] If the standard is

[3] The presence of market power and markups should change the analysis little so long as markups are either constant across all vehicles or increasing in attribute levels. In this case, markups are simply captured by the cost function discussed later. A problem may arise if there is some other pattern of markups.

[4] This representation of vehicle costs also appears in Knittel (2011).

[5] This is identical to writing the constraint as requiring the quantity-weighted harmonic mean of fuel economy (miles per gallon) to be greater than a fixed level, which is a typical representation of the constraint.

an attribute-based standard, then the level of the standard can be written as a function of the attribute, $\sigma(x_{j'l})$.[6] Note that a CO_2 per mile standard is easily translated into a fuel economy standard, and so for our purposes, we will only discuss a fuel economy standard. The manufacturer's constrained maximization problem is then

$$\mathcal{L} = \sum_l q_l(p_l - c(X_l, Z_l)) + Q\lambda \left(\sigma - \frac{1}{Q} \sum_l q_l \cdot g_l \right). \tag{3.4}$$

We have multiplied λ by Q so that λ may be interpreted as the per vehicle shadow cost of the standard. If the constraint is binding, the first-order necessary condition is

$$[q_l] : p_l = c^1(g_l, X_l) + c^2(Z_l) + \lambda(g_l - \sigma) \quad s.t. \lambda \geq 0. \tag{3.5}$$

If the standard is not binding, then $\lambda = 0$ and price is equal to the cost of producing a vehicle.

We want to solve for equilibrium attributes given consumer preferences. To do that we assume a functional form for the c^1 cost function. In keeping with Knittel (2011), although substantially altered as discussed below, we assume a Cobb-Douglas-like functional form for cost function of physically related attributes:

$$c^1(g_l, X_l) = T \left(c^{\max} - g^{\alpha_g} \prod_j \underline{x}_j^{\alpha_j} \right). \tag{3.6}$$

We define x_j as the measure of an attribute as a "good," and can then define $\underline{x}_j = x_j^{-1}$ as the same attribute measured as a "bad."[7] We define x_j as a "good" in the sense that increasing the attribute level, else equal, is desired by all consumers; consumer utility is monotonically non-decreasing in x_j, which means that utility is monotonically non-increasing in $\underline{x}_j$. T corresponds to technology. The effect of decreasing T makes the cost of all bundles of attributes decrease.

[6] However, a reformulation of the problem as an attribute standard does not change the first-order condition with respect to quantity in Equation 3.5, which determines the price function. An attribute based standard would impact our problem later on if c^1 is a function of the attribute-basing attribute.c^1

[7] The function would work just as well with other transformation of "goods" into "bads" such as $\underline{x}_j = x_j^{\max} - x + j$, but the inverse is commonly measured (e.g. fuel consumption) and has convenient properties that we will see later on.

This functional form breaks with Knittel (2011) as well the implicit assumptions of all other studies that use a log-log form to estimate the relationship between attributes. While Cobb-Douglas models are standard for utility and production functions, they are not appropriate for the production of multiple attributes. The functional form fails for several reasons. First, it fails the no free lunch principles of PPFs. As we discussed in the introduction, the set of all iso-cost curves for a particular differentiated goods describes a special type of PPF. The no free lunch principle states that there can be no production for free or without inputs.

However, the Cobb-Douglass model allows any level of attribute production at no cost so long as another attribute is set to zero. Second, assuming free disposal, the set of production possibilities is not convex. This means that, in a model that includes consumers, the second order conditions are not satisfied unless the shape of consumer utility curves are constrained to be more convex than iso-cost curves. In particular, using linear utility functions fails the second order conditions and has no unique or finite maxima bundle of attributes. We show this in the Appendix. Third the shape of the cost function has non-intuitive implications. It implies that for a given cost extremely large values for some attributes if other attributes approach zero. As we do not observe tiny cars that break the sound-barrier or slow rolling cars that are the size of wales, this functional form suggests the availability of implausible products.

The two changes we have made to the Cobb-Douglas functional form, subtraction of the second term from a maximum cost and the transformation of X so that $\underline{X}$ increasing is less preferred by consumers, is critical for our model to produce a unique solution. This change creates convex production sets, satisfies the no free lunch principle, and implies a maximum attribute level for each attribute (in goods space) for each cost level, which is intuitive.[8]. Using the functional form in Equation 3.6 makes iso-cost curves convex to the origin, when graphed in goods-space, for positive level of attributes j less than $x_j^{\max}$. A proof sketch of

[8]To see this, rearrange the cost function so it becomes $x = \left(c^{\max} - \underline{x}_2^{\alpha_2} - c^1\right)^{-1/\alpha_g}$

the second order conditions and drawing is provided in the Appendix. Subtracting from a theoretic maximum level of output is common to stochastic frontier models (for example Aigner, Lovell, and Schmidt 1977). Figure 3.1 illustrates the diffence in our Cobb-Douglas like cost function and the Cobb-Douglas cost function.

Figure 3.2: Illustration of functional form considerations

Our Cobb-Douglas-like cost function

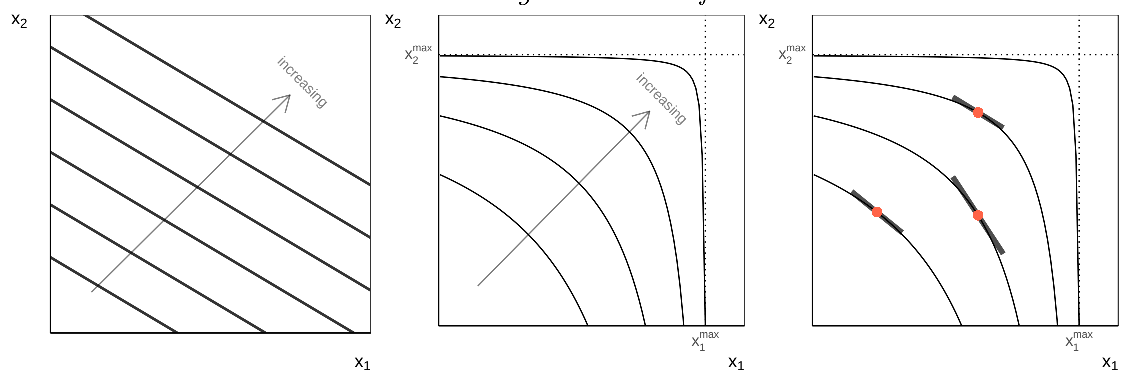

Left: an individuals linear iso-utility curves; middle: our Cobb-Douglas-like iso-cost curves trading off between two attributes; right: bundles of attributes (vehicles) are points of tangency for individual iso-utility curves and iso-cost curves.

A Cobb-Douglas cost function

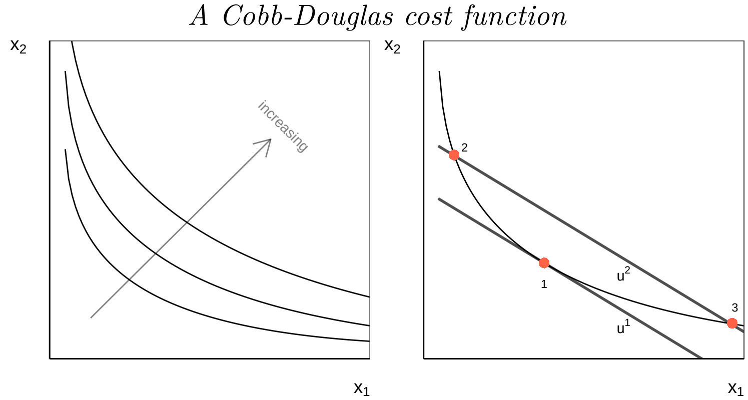

Left: A Cobb-Douglas iso-cost curve trades off two attributes. Right: The unique solution, point 1, is a point of tangency of the iso-utility curve, u^1 *on a given iso-cost curve; this consumer can be made better off by moving to points 2 or 3, which are are on a higher utility curve,* u^2*.*

A possible downside of this functional form is that it has a cross-partial derivative that is negative. This implies that the marginal cost of a given attribute is decreasing as other attributes increase. While this may seem unintuitive at first, the relationship is consistent with our understanding of the engineering in this context. Both fuel consumption and acceleration reflect the transformation of energy into propulsion. Some engineering inputs that improve one attribute improve the other or make the other easier to achieve. For example, light-weighting a vehicle with a fixed drive-train will improve both acceleration and fuel consumption. Adding a turbocharger to a fixed drive-train will improve acceleration, while affecting fuel consumption negligibly; turbochargers have allowed OEMs to reduce the engine size to improve fuel consumption while keeping acceleration constant.[9]

Heterogeneous consumers

Assume that there are a mass of consumers, indexed by i, with heterogeneous preference for the attributes. We then can specify a separable, locally-linear utility maximization problem

$$\max_{l} u_i \approx v_i(g_l, X_l, Z_l) + y_i - p_l \tag{3.7}$$

where yi is income for consumer i. Assume further that v_i is linear in g_l and X_l. To conform to the transformation used in Equation 3.6, we write

$$v_i(g_l, X_l, Z_l) = v^{\max} - \sum_j \beta_{ij} \underline{x}_{lj} - g_l p_g \phi_i \tag{3.8}$$

where p_g is the price of fuel, ϕ_i is an individual's preference for (dislike of) fuel consumption, and β_i is a vector of preference parameters on other fuel consumption-related attributes.

Assuming either perfect competition or constant markups, we replace vehicle price p_l with the vehicle cost function in Equation 3.2. We then find equilibrium outcome by choosing attribute levels that maximize consumer utility accounting for the effect of attribute levels on price. The first-order necessary conditions with respect to a given attribute x_{lj} on chosen vehicle l are

[9] Alternatively, we may use a quadratic functional form, which has the advantage of a flexible cross-partial derivative but has many more parameters—too many for us to estimate with our data.

$$\left[\frac{\partial u_i}{\partial x_j}\right] : \beta_{ij} = -c^1_{x_j}(g_l, X_l)x_j^2, \tag{3.9}$$

$$\left[\frac{\partial u_i}{\partial g}\right] : \phi_i p_g + \lambda = c^1_g(g_l, X_l). \tag{3.10}$$

Instead of Equation 3.10, we may instead take the partial deriviative with respect to fuel economy, g^{-1}, which yields $\left[\frac{\partial u_i}{\partial g^{-1}}\right] : \phi_i p_g + \lambda = -c^1_g(g_l, X_l)(g^{-1})^2$.

To simplify notation define $\beta_{ig} \equiv \phi_i p_g + \lambda$. Making use of Equations 3.10 and 3.9, we can solve for the optimal attribute level as a function of preferences and cost para maters (see Appendix for derivation):

$$g^{-1*} = \left(\frac{\alpha_g}{\beta_{ig}} T \prod_{j \in J} \left(\frac{\beta_{ig}}{\beta_{ij}} \frac{\alpha_j}{\alpha_g}\right)^{\alpha_j}\right)^{\frac{1}{\psi}}, \tag{3.11}$$

where $\psi \equiv \alpha_g + \sum_j \alpha_j - 1$ is negative due to constraint given by the second order conditions, $\alpha_g + \sum_j \alpha_j < 1$, and shown in the Appendix. We further assume that all cost parameters are strictly positive, which is required for the cost function to be increasing in all "good" attribute levels. The optimal level of attributes x_j is similarly:

$$x_j^* = \left(\frac{\alpha_j}{\beta_{ij}} T \prod_{k \in K} \left(\frac{\beta_{ij}}{\beta_{ik}} \frac{\alpha_k}{\alpha_j}\right)^{\alpha_k}\right)^{\frac{1}{\psi}} \tag{3.12}$$

where $K \equiv \{\forall j' \neq j, g\}$ââall attribute indexes including fuel consumption but not including the current attribute that is being solved for. Note that Equations 3.11 and 3.12 assume $\underline{x}_j \equiv x^{-1}$. To solve for attributes as "bads," $\underline{x}_j$ and g, simply take the inverse of both sides.

Attribute trade-offs

Our object of interest, the attribute trade-off or MRAS, is the slope of the isocost curve in attribute space, or

$$\frac{\partial C}{\partial \underline{x}_j} / \frac{\partial C}{\partial g_j} = \frac{\alpha_j}{\alpha_g} \frac{g}{\underline{x}_j} \tag{3.13}$$

when it is written in terms of our bads.

The MRAS can be written in a very similar manner for more familiar variables where x_j is a good and g^{-1} is fuel economy (e.g. miles per gallon)

$$\frac{\partial C}{\partial x_j} / \frac{\partial C}{\partial g_j^{-1}} = -\frac{\alpha_j}{\alpha_g} \frac{g^{-1}}{x_j}, \tag{3.14}$$

which is derived in the Appendix. The elasticity of the MRAS—the percent change of one attribute for a percent change in the other holding costs fixed—is then simply α_j/α_g when the elasticity is either in "goods" or "bads" space in the way that we have defined $\underline{x}_j$. However, if the transformation of the attribute measured as a "good" into a measure that is a "bad" is not the inverse, then this is not the case and the elasticity is not constant for all levels of the good.

Comparative statics

We investigate four changes to the cost function to characterize the effect of each on the distribution of vehicle attributes. The first three changes we investigate represent changes in technology while the third captures a change in the stringency of the regulatory standard.

We describe four effects corresponding to the four changes in model parameters: 1. "More of everything." When technology improves and is attribute-neutral, all consumer want more of all attributes by the same proportion regardless of preferences. 2. "More of what I like." When technology improves that is biased towards the production of one attribute, consumers will increase their optimal level of consumption for the attribute favored by the technology, proportional to their prior optimal level of consumption. Depending on their taste for attributes that do not benefit from the biased technological improvement, consumers may increase or decrease their consumption of the attribute not favored. 3) "Rich get richer." Discrete technologies (e.g. turbochargers) are first adopted by consumers who consume the most of all attributes. 4) "Leave no cars behind." When regulatory standards become more stringent, consumers who least preferred fuel economy respond by increasing their fuel economy more than those who most prefer fuel economy.

There are two senses of the word technology in our analysis. We use it when discussing discrete technology in the common way: a physical thing created by innovative engineering. The second use of technology abstractly indexes the innumerable production and engineering innovations that allow a manufacturer to produce one ore more attributes using fewer resources. This is how technology typically appears in the economics literature, including the economics literature on attribute trade-offs.

Attribute-neutral technical change

Technical change that is not biased towards one attribute or another is the sort of technical change that has been typically modeled in the literature. The technical restriction for unbiasedness is that is that the ratios of marginal costs do not change ($\Delta \frac{c_{x_k}}{c_{x_j}} = 0$). This can be accomplished with a decrease in T, which reduces the cost of producing of all attribute combinations. Visually, this is the same as shifting isocost curves outwards in goods space so that, for any cost level, more of all attributes may be produced. However, as we do not constrain the exponent parameters in our cost function to sum to 1, attribute neutral technical change may also be accomplished by reducing all α_j by the same proportion.[10]

Consider a marginal decrease in the technology index term T. We can easily see that a change in T changes the marginal cost of each attribute by the same proportion, making the change not attribute-biased. By taking logs of Equation 3.12 we get a simple representation of the effect of a percent change in T on the attributes

$$\frac{\partial \ln x_j^*}{\partial \ln T} = \frac{1}{\psi} < 0. \tag{3.15}$$

A reduction in T leads to a reduction increase in all attributes. With our chosen functional form, this increase is a fixed proportion of the baseline level of the good. Alternatively, if we scale all α_j by ω, then the partial effect of ω is the same as Equation 3.15 and reducing ω

[10] Recall that $\frac{c_{x_k}}{c_{x_j}} = \frac{\alpha_k}{\alpha_j}\frac{x_j}{x_k}$.

increases all attributes by the same proportion. However, when it comes to estimation, we will need to account for both pathways for attribute neutral technical change.

Attribute-biased technical change

To capture different types of technology change we investigate both changes that are biased towards a particular attribute as technical changes that are not biased. A technical change that is biased towards attribute x_k is defined as a change in the cost function where

$$\Delta \frac{c_{x_k}}{c_{x_j}} < 0$$

where j and k index different attributes. This implies $\Delta \frac{\alpha_k}{\alpha_j} < 0$ using our specification in Equation 3.6.[11]

To our knowledge, attribute-biased technical change has not been considered in the differentiated goods literature. However, there is a considerable literature of factor-biased or skill-biased technical change in macroeconomic aggregate production (for example Acemoglou 2002). We believe attribute-biased technical change is relevant in this context because of evidence from engineering data discussed in the next section. These data provide some evidence that iso-cost curves have not only shifted outwards (in goods space) but have changed in slope.

To interrogate this, we simply analyse the effect of reducing α_k. The own-attribute effect on optimally chosen attribute levels from the attribute-biased technical change is

$$\frac{\partial \ln \underline{x}_k^*}{\partial \ln \alpha_k} = 1 + \frac{\alpha_k}{\psi}\left(1 + \ln \underline{x}_k^*\right) > 0, \tag{3.16}$$

which is always positive. The cross-attribute effect on optimally chosen attribute levels from the attribute-biased technical change is

$$\frac{\partial \ln \underline{x}_j^*}{\partial \ln \alpha_k} = \frac{\alpha_k}{\psi}\left(\ln \underline{x}_j^* + \ln \frac{\alpha_k}{\alpha_j}\right) \lesseqgtr 0, \tag{3.17}$$

[11] $c_{x_j} = T\alpha_j \frac{c_1}{\underline{x}_j}$. So, $\Delta \frac{c_{x_k}}{c_{x_j}} = \frac{\underline{x}_j}{\underline{x}_k}\left(\frac{\alpha_{k,t=1}}{\alpha_{j,t=1}} - \frac{\alpha_{k,t=0}}{\alpha_{j,t=0}}\right)$, where t is time.

which may be positive or negative depending on sign and magnitude of $\ln \frac{\alpha_k}{\alpha_j}$ compared to the baseline attribute level. Therefore, while attribute k-biased technical change will always increase consumer's consumption of attribute k, it may increase or decrease attribute j. This type of technical change may even increase attribute j more than attribute k. The cross-attribute effect could be negative if α_j is large relative α_k and the consumer has high baseline levels of j, which implies that they have a high value of the attribute. This gives the non-intuitive result that under certain conditions, those with the highest preferences for attribute j will consume less of it if there is k-biased technical change, while those who prefer j less will consume more of both attributes.

Later we show evidence that attribute-biased technical change is important for understanding changes in attribute trade-offs for light duty vehicles since the 1980s.

Discrete technology adoption

Periodically, manufactures invent new technologies that may be added to vehicles to improve one ore more attributes. For this discussion, turbochargers exemplify the sort of discrete technology that we have in mind. We adapt our model to understand how discrete technologies of this type would be added onto vehicles given each vehicle's consumer segment and baseline attribute levels. There are other examples of discrete technologies that also fit this analysis, such as continuously variable transmission, direct injection (replacing carburetor), hybrid or electric hybrid drive-train.

In simple terms, a turbocharger works by temporarily increasing the performance output of an internal combustion engine by forcing additional compressed air into the combustion chamber. Adding a turbocharger to a vehicle's drive train allows an otherwise identical vehicle to accelerate faster. Manufactures have frequently used turbochargers while decreasing engine displacement ("downsizing"), therefore improving fuel economy and maintaining acceleration at similar levels (Shahed 2009).

Our discrete technology has two important features for our analysis. First, the cost, t, of installing the technology in some year is approximately the same for all vehicles regardless of baseline attributes. Second, the technology has a proportional effect on the attribute that the technology targets. Vehicles with a higher baseline of the targeted attribute (power) see higher gains.

The simplest way to model the attribute-targeted effect of the technology is as an individual consumer choice to adopt the technology, incur cost t, scale the benefit of the targeted attribute by $\omega > 1$. Assume that the technology targets x_j.

Assume that a given consumer i chooses to adopt the technology. This consumer's utility will then increase by how much more of all attributes they will consume when adopting the technology, scaled by how much they enjoy the technology, and decrease by the cost of the technology, or:

$$\Delta u_i^A = -\sum_j \beta_{ij} \left(x_j^{*^{-1}} \frac{\partial \ln x_j^{-1^A}}{\partial \omega} \Delta\omega \right) - t. \tag{3.18}$$

In equilibrium, the adopter will increase their level of the targeted attribute x_j by the proportion

$$\frac{\partial \ln x_j^{-1^A}}{\partial \omega} \Delta\omega = -\frac{\psi - \alpha_j}{\psi\omega} \Delta\omega < 0. \tag{3.19}$$

Both numerator and denominator are negative and so the effect of increasing ω is to increase x_j. The adopter will also increase attributes that are not directly affected by the discreet technology as the increase in x_j makes marginal increases in other attributes, x_k, less costly. They will increase these non-targeted attributes by

$$\frac{\partial \ln x_k^{-1^A}}{\partial \omega} \Delta\omega = -\frac{-\alpha_j}{\psi\omega} \Delta\omega < 0 \tag{3.20}$$

As the change in utility if the consumers do not adopt is zero, consumers will adopt if 3.18 is positive. As the equilibrium level of all attributes is increasing in all consumer preference parameter—$\partial x_j^* / \partial \beta_j, \partial x_k^* / \partial \beta_j > 0$— Equation 3.18 is increasing in both preferences for the attribute that is directly improved by the technology as well as the attributes that are not directly improved. The result is that technologies are first adopted by those who prefer higher attributes in general. As the cost, t of the technology drops over time, more consumers with lower attribute levels adopt the vehicles with those consumers with the lowest overall preference adopting the technology last.

We then expect that new technology is not uniformly adopted in the fleet. Instead, attributes that are not directly affected by the technology are predictive of adoption. This result has implications for estimation. To estimate MRAS by fitting a line through attribute data, we must assume that the mean level of c^1 is constant across all levels of a given attribute. However, this result implies that discrete technologies are adopted more, and therefore average cost is higher, for vehicles with higher attribute levels.

Change in standards

Finally, we consider a change in the level of the fuel economy standard. As the level of the standard, σ, does not appear in the the optimal attribute choice for the consumer, we instead examine an increase in λ, the shadow cost of the standard, as an increase in the level of the standard must increase λ at least in the short to medium run. An increase in the shadow cost of the standard has the following effect

$$\frac{\partial \ln g^{*-1}}{\partial \lambda} = \frac{\psi - \alpha_g}{\psi(\phi_i p_g + \lambda)} > 0 \tag{3.21}$$

As $\psi < 0$, both the numerator and denominator are negative and the effect is positive. Standards improve fuel economy. As the effect is decreasing in the value of fuel economy,ϕ_i, consumers with the highest preference for fuel economy are affected the least by increased standards while those with the lowest preference for fuel economy are improve their fuel economy the most. Fuel prices also appear in the equation interacting with fuel economy preferences. This implies that higher fuel prices have a dampening effect on an increase in standards in the short term. This dampening effect is largest for consumers who are the most sensitive to the cost of driving.

We can also look at how standards change other attributes.

$$\frac{\partial \ln x_j^*}{\partial \lambda} = \frac{-\alpha_j}{\psi(\phi_i p_g + \lambda)} > 0 \tag{3.22}$$

Again the effect is positive, implying that standards increase other non-regulated attributes. This effect is the result of the model's structure. As any one attribute increases, it the marginal cost of others decrease. So as consumers are encouraged to purchase more efficient vehicles, the marginal cost of other attributes falls and they buy more until marginal costs meet marginal benefits. Similar to the previous result, this effect is dampened by higher preferences for fuel economy and by higher gas prices.

3.2 Empirical analysis

The empirical analysis takes two approaches, both using data of attribute levels from all produced sub-vehicle models as reported by Wards Automotive. These data are matched to sales volumes and indicator variables for vehicles with diesel engines and turbochargers from EPA administrative data as reported in EPA (2019).[12]

The first is a structural approach. We estimate the theoretical model for two time periods: model years 1990 to 1995 comprise an older period, and model years 2014-2019 a contemporary period. We use this approach to estimate the α parameters and T from our cost function for the two time periods and report the MRAS for power and fuel economy and measure unbiased and biased technical change. We find evidence that the MRAS has declined substantially, indicating biased technical change. Our second approach is to deploy reduced-form models to test the predictions of our model for discrete technologies and an increase in the standards. In this section we first discuss the mechanics of the two methodological approaches. We then present results in the same order as the theory section.

To supplement the structural estimation of MRAS we take a reduced form approach using the National Household Travel Survey. This survey asks households about vehicle ownership

[12]The Wards data has more sub-model vehicles than EPA administrative data as well as variables not reported to EPA. Sales volumes are split evenly across all matched vehicles in Wards.

and use. We use the 2001 and 2017 surveys. We use only vehicles for the five years before the survey was conducted to make it likely that vehicles were purchased new. We match these data to EPA and Wards attribute data using model names and fuel economy.

Methods

A structural model to estimate cost parameters

Using our model to estimate cost parameters, we either need to know the total cost of each vehicle[13] or the value of each bundle of attributes to the consumer who buys them, which in our model is equal to the per-vehicle cost of production. Unfortunately neither manufacturing costs nor consumer value are directly observable or at least, in the case of manufacturing costs, they are not shared with researchers.

Consumer preferences have been estimated using discrete choice models of observed purchase decisions or elicited in stated preference surveys. However this literature has so far yielded estimates of consumer preference values that do not agree with each other. This broad range of estimates also appears sensitive to assumptions and the inclusion or exclusion of controls (Greene et al. 2018). Table 3.2 is reproduced from Greene et al. (2018) and summarizes willingness to pay estimates from the literature. We are therefore reticent to choose one particular set of consumer preference values from the literature or to attempt our own estimate.

Rather than estimate heterogeneous consumer preference parameters in a discrete choice model, we estimate our model assuming we know the set of mean attribute preference parameter values from a multivariate normal distribution of logged preferences. We then estimate our model for each combination of mean attribute preference parameter values in evenly spaced increments along a range. For each estimation we fit both the covariance matrix of the preference parameter distribution and the cost function parameters (treated

[13] Then we could simply write a linear approximation of our cost function and regress with observed attributes as right hand side variables and cost on the left.

Table 3.1: Willingness to pay estimates in the literature

	N	Units	Median	SD
Aggregate fuel cost per mile	117	\$/cpm	991	6,875
Gallons per mile	24	\$/0.01 gpm	1,027	1,484
horsepower/ weight	29	\$/0.01hp/lb	198	1,449

Table reproduced from data in Greene et al. (2018).

as homogeneous across all vehicles and firms) to the data using a modified Expectation-Maximization (E-M) algorithm (Dempster, Laird, and Rubin 1977; Dempster, Rubin, and Tsutakawa 1981). Repeating this procedure for each combination of candidate mean preference values, we produce a grid of estimated parameters, MRAS elasticities, and differences between time periods, which are shown in Figures 3.3 and 3.5.

Why not simply control for vehicle price as a proxy for cost? After all a proxy for cost would allow us to preform a simple linear regression on our attribute data. Here we run into several problems. The first is that list prices, which are readily available, are not the real transacted price which is negotiated between dealers and consumers. Second, the price of vehicles is affected by the shadow cost of regulation as a function of attribute levels so that we cannot recover a valid proxy for production cost from list price. Third, prices reflect additional attributes, such as leather seats, that have not physical relationship with the production cost of the drive-train but we would expect to be correlated with drive-train related attributes such as performance.

Parameters of the model are estimated using a modified E-M algorithm, which consists of an iterative loop between an expectations step (E-step) and a maximization step (M-step). Define $\mathbf{A} \equiv \{\alpha_j\}$, the set of cost parameters to be estimated, and $\tilde{\mu} \equiv \{\mathbb{E}[\beta_j]\}$ as the set of assumed mean consumer preference parameters for a given run of the estimation procedure. Assume that preferences are distributed as multivariate log-normal, written $N(\ln \mu, \Sigma)$.[14]

[14] We use the log-normal distribution as it is 1) quickly computed for the multivariate case, 2) requires no other parameters of the distribution to be estimated other than the covariance matrix, 3) constrains preferences to be positive (disallows negative marginal value of fuel power).

Let p index the current iteration of the algorithm. The E-M algorithm proceeds by first assuming initial values of $\hat{\mathbf{A}}^{(o)}$ and then alternating between an E-step and an M-step until the value of the likelihood function converges.

The E-step updates $\hat{\Sigma}^{(n)}$ given $\hat{\mathbf{A}}^{(n)}$. To do this, Equations 3.11 and 3.12 is re-written to solve for preferences $\hat{\beta}^{(n)}(X, \mathbf{A}, p_g)$ and applied to the data, which gives a matrix of estimated preferences, where each row of the vector corresponds to a row of the vehicle data. $\hat{\Sigma}^{(n)}$ is updated using the sample covariances of $\hat{\beta}^{(n)}$.

The M-step uses maximum likelihood to update $\hat{\mathbf{A}}^{(n+1)}$ and $\hat{\beta}^{(n+1)}$ with respect to $\hat{\mathbf{A}}$ assuming β is distributed $N(\ln \tilde{\mu}, \hat{\Sigma}^{(n)})$, or

$$\hat{\mathbf{A}}^{(n)} = \arg\max_{\mathbf{A}} \sum_{it} \ln \Pr\Big(\hat{\beta}_{it}(X_{it}, \mathbf{A}, p_{gt})\Big|\beta \sim N(\tilde{\mu}, \hat{\Sigma}^{(n)})\Big) w_{it}, \tag{3.23}$$

where w_{it} are sales volumes for vehicle i in year t. This maximum likelihood is found numerically using stochastic gradient ascent. Each E-step gives a likelihood value, $L^{(n)}$, given the current estimate of $\hat{\Sigma}^{(n)}$.

Dempster, Laird, and Rubin (1977) show how, by alternating between the E and M-steps, the log-likelihood must increase in each iteration. We verify that the log-likelihood does increase with each iteration and allow the algorithm to run until the difference in log-likelihoods, $L^{(n)} - L^{(n-1)}$, is less than 10. We compute each likelihood value , $L^{(n)}$, using sales volumes for each observed sub-model vehicle. Using sales volumes in estimation is important as certain vehicles have many more sub-model versions per sale than others and without sales weighting have an out sized effect on $L^{(n)}$.

For simplicity, we estimate a model using only acceleration and fuel consumption. We use the net horsepower-to-weight ratio[15] for each vehicle as a proxy for acceleration. While this may not be an ideal proxy—MacKenzie and Haywood (2012) have shown that some measures of acceleration such as 0 to 60 miles-per-hour time have increased for a given horsepower-to-weight ratio over time—it is consistently measured for all time periods.

[15]"Net" refers to an engine's output net of drive-train accessories such as the alternator.

We assume a range of consumer preference parameter values, $\tilde{\mu}$. For the mean marginal benefit of power, our range extends from 100 to 300 $US for a 0.01 increase in net horsepower-to-weight ratio. For the mean marginal benefit of fuel consumption, our range extends from 500 to 1,500 $US per 0.01 gallons per mile decrease. Each range is divided into ten equal intervals, making a 100 cell grid of candidate $\tilde{\mu}$ combinations. These chosen ranges are approximately centered on the mean value in the literature as summarized by Greene et al. (2018). We choose a range that is half to 1.5 times this center point, as Greene et al. (2018) find implausibly large ranges that include negative values within one standard deviation, as summarized in Table 3.2. It is worth noting that the middle value of mean marginal benefit for a reduction in fuel consumption, $1,000, is to a first approximation the full monetary value of the discounted value of a marginal change in fuel consumption over 10 years of driving for an average driver.[16]

Finally, we assume a value of λ, the shadow cost of the regulatory constraint, to input into the estimation of the model for each time period. For the earlier period we use an the shadow of approximately $15 per MPG estimated by Anderson and Sallee (2011). This shadow cost is then transformed to 0.01 GPM by evaluating it for a one MPG difference at the standard.[17] For the later period we use $81 per regulatory credit (in Mg CO_2 units) as reported in Leard (**(year?)**). The credit cost is then transformed to 0.01 GPM using the formula and lifetime mileage assumption given by the EPA.[18](CITATION). As it turns out, we need not worry about the accuracy of our assumed shadow costs. Figure C.2 in the Appendix shows the range of cost parameter estimates using all preference parameters assumptions. As the table summarises, the value of λ has nearly no effect on these estimates, except for a handful of outlier estimates.

[16] $10 years \times 10^4 \frac{miles}{year} \times 10^{-2} \frac{gallons}{mile} \times \frac{\$1}{gallon} = \$10^3$.

[17] $(\frac{1}{26} - \frac{1}{27}) \times 100 \times 15 \approx 2$

[18] Credits (Mg CO_2) $\times 10^6$ / (Lifetime mileage assumption $\times 8887$ g CO_2/gallon ≈ 4.6, where lifetime mileage is 195,264 for cars.)

Structural model results

We compute the MRAS elasticity for each set of cost parameters. Recall that the MRAS elasticity is simply α_j/α_k, the same formulation as a Cobb-Douglass elasticity, and is constant for all levels of attributes. This is the elasticity in either "goods" or "bads" space so long as the "bads" that are used in the model are the inverse of the "goods" measure. We evaluate this scalar for the mean power level of sub-model vehicles in the contemporary time period.

Figure 3.3: MRAS elasticity

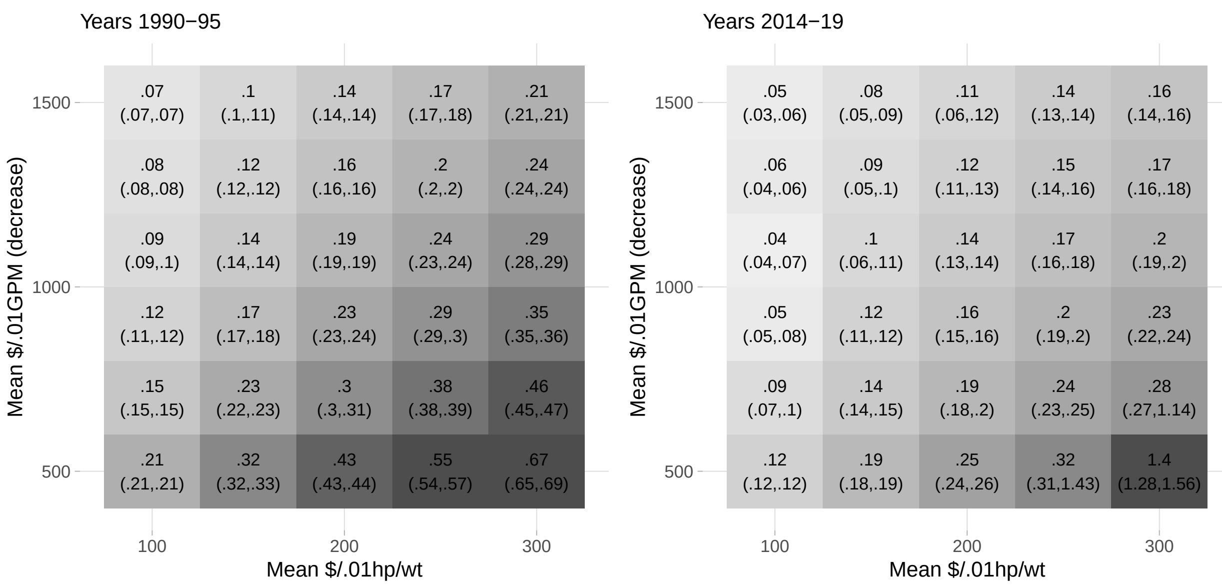

Note: The inner 95% interval of 500 bootstrapped estimates is shown in parenthesis.

Figure 3.3 shows the scaled MRAS elasticity for each time period. The earlier period elasticity point estimates range from 0.07 to 0.67, while the later period range from 0.05 to 1.4 These elasticises are interpreted as the percent improvement in efficiency[19] caused by a percent reduction in power, holding vehicle costs fixed. Each cell of the tables corresponds to a mean preference parameter combination, as described above. The elasticity monotonically

[19]Fuel consumption and fuel economy are equivalent for the analysis.

increases in the assumed mean marginal value of power and monotonically decreases in the assumed mean marginal value of reduced fuel consumption.

Figure C.1, in the Appendix, shows the change of the parameter T. As we discuss in the comparative statics section, T alone is not a sufficient measure of overall cost reduction as it is possible for T to increase and for costs to decrease as we are allowing the other cost parameters to change overtime as well. Figure 3.4 reports the percent change of producing each attribute for the sales weighted average level of each attribute in 1995. Nearly all of the point estimates shows decline in the marginal cost of producing each attribute while the marginal cost of performance shows larger decreases than fuel economy.

Figure 3.4: Percent change in marginal cost of attributes

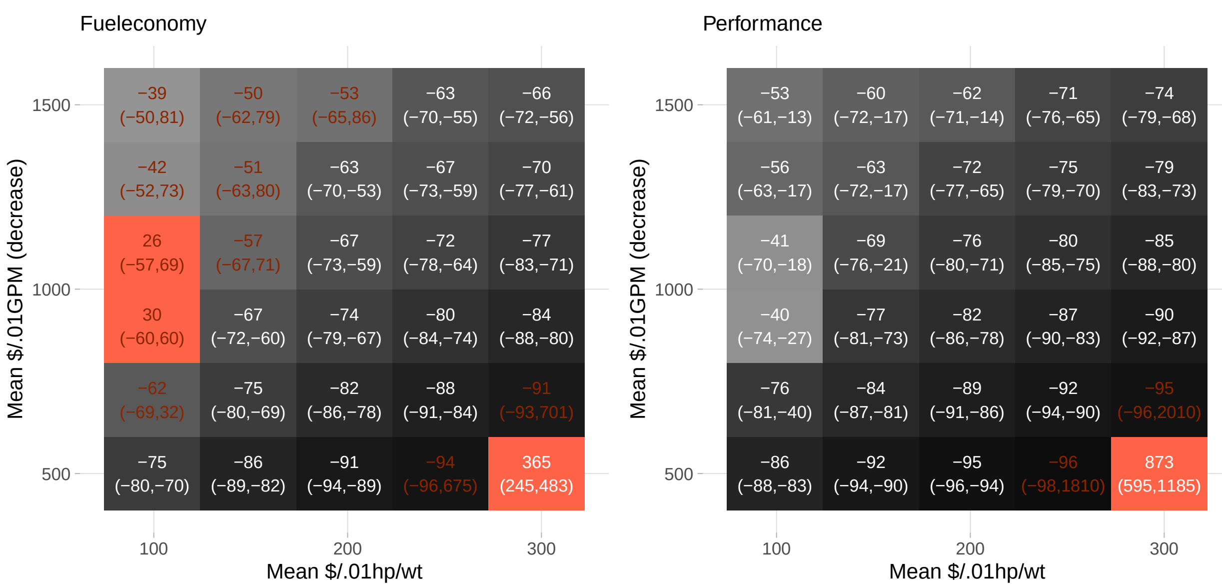

Note: The inner 95% interval of 500 bootstrapped estimates is shown in parenthesis; red text indicates estimates with an interval that crosses zero; negative values are reductions in the marginal cost of attributes.

We also find robust evidence to support substantial biased technological change between the two periods that is remarkably consistent for all preference assumption. Taking the percent difference of corresponding elasticities yields Figure 3.5 that shows the percent change in MRAS elasticity between the two time periods. With the exception of one outlier, these point estimates range from a decrease from 25 to 54%. Parenthesis show the inner 95% range of estimates from 500 bootstraps[20]

Figure 3.5: Percent change in MRAS elasticity

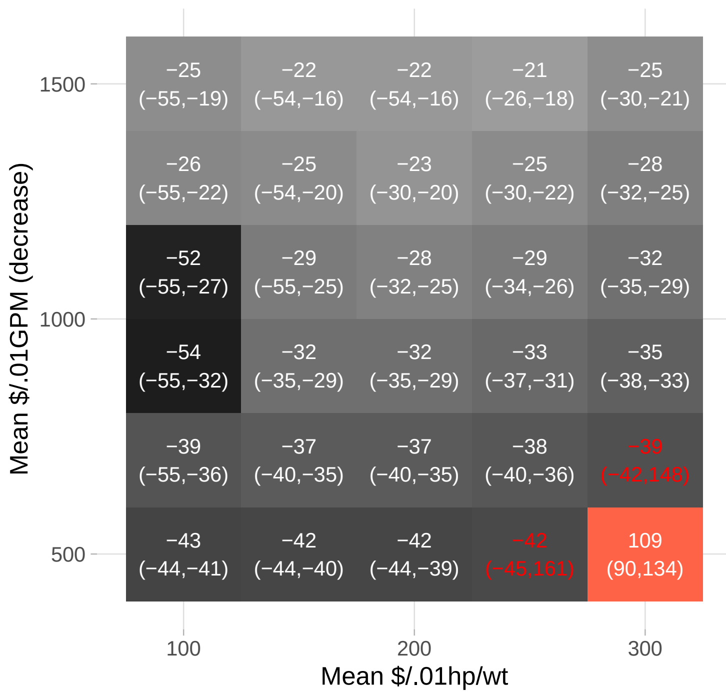

Note: The inner 95% interval of 500 bootstrapped estimates is shown in parenthesis; red text indicates estimates with an interval that crosses zero; negative values are improvements in biased technical change.

This difference assumes that mean consumer preferences have not changed between the two periods (while the covariance matrix in each period can change) when vehicle attribute data are sales-weighted. There is also no way to know if the mean of consumer preference values has changed substantially from our estimation. While it is clear that there are different types of vehicles on the road in recent years than there were in the early 1990s, these changes could be caused by changes in technology and fuel prices, which are captured in our model. Taking the central elasticity value from the earlier period, the mean marginal benefit of

[20]Estimates for each time period were bootstrapped 500 times. Therefore the interval is the percent difference of the 250,000 combinations of boot straps.

reduced fuel consumption would need to fall by about half, or the mean marginal benefit of power would need to increase by nearly double, for there to be in change in the elasticity between the two time periods. The former would imply either a substantial decrease in the preference for driving or future fuel savings. This seems unlikely. More plausible is that the mean marginal benefit of a power increase is lower than it was in the earlier period. We have constrained our consumers to have linear preferences for attributes (for tractability purposes), but it is likely that marginal benefit from power decreases as vehicles become more powerful. A marginal increase in 0 to 60 time from 13 seconds, as was typical in the 1990s, may have been more valuable for the average consumer than a marginal improvement from the 8 seconds more common currently.

Reduced form estimation of cost parameters

A shortcoming of the structural model is that we must make assumptions about the mean preferences in the population over time. If there are large and unknown changes in consumer preferences then we cannot match estimated parameters from two time periods. We can relax the mean preference assumption by assuming that preferences are a fixed function over time of demographic measures and using data of consumer purchases. If we then observe a similar person on one period buying a car with very different attributes we can take this as evidence that attribute production technology has changed.

Dividing and then rearranging first-order conditions from our theoretic model, we can write

$$\ln(g^{-1}) = \Big(\ln \alpha_g T - \ln(\phi_i p_g + \lambda) + \sum \alpha_j \ln \underline{x}_j \Big) \Big/ \Big(\alpha_g - 1 \Big) \tag{3.24}$$

If the shadow cost of regulation is small, then we can write a reduced form model that relates to Equation 3.24

$$\ln(g_i^{-1}) = \gamma_0 + \gamma_1 \ln \frac{hp_i}{wt_i} + \gamma_2 \ln p_g + D_i' \delta + \varepsilon_i \tag{3.25}$$

where D is a vector of demographic variables. If $D_i'\delta$ equals ϕ_i, then

$$\gamma_0 = \frac{\alpha_{power} + \ln T\alpha_g}{\alpha_g - 1} \tag{3.26}$$

$$\gamma_1 = -\frac{\alpha_{power}}{\alpha_g - 1} \tag{3.27}$$

$$\gamma_2 = -\frac{1}{\alpha_g - 1}. \tag{3.28}$$

We can then solve for the elasticity

$$\frac{\alpha_g}{\alpha_{power}} = \frac{\gamma_2}{\gamma_1}\Big(1 - \gamma_2^{-1}\Big). \tag{3.29}$$

Reduced form results

Table 3.2 shows the results. Model 1 includes annual vehicle mileage to control for consumer preferences for fuel savings while Model 2 included other demographic data. These variables are: household size, number of vehicles, state, life cycle stage (which includes age of youngest child, retirement status, partnered status, etc.), and urban category (urban, suburban, rural). All coefficients are significant at the 1% level. The calculated percent change in elasticity is -90% for Model 1 and -94% for Model 2. This is a substantially larger estimate than from the structural model.

A reduced form model of turbocharger adoption

Using our vehicle attribute data from 1996 to 2019, the years for which we have reliable data on turbochargers, we test the predictions of our model for discrete technology adoption. Figure C.3 in the Appendix summarises adoption of turbochargers throughout the fleet and shows a steady increase in the proportion of vehicles with adoption over time. Our model predicts that turbocharger adoption will be associated with higher attribute levels (in goods space). We test this by regressing log attribute levels (in goods space) on the a binary turbocharger variable that equals one if a vehicle has a turbocharger and zero otherwise.

Table 3.2: Consumer vehicle choice

	Dependent variable:	
	ln MPG	
	(1)	(2)
ln(hp/wt)	-0.139^{***}	-0.136^{***}
	(0.001)	(0.001)
$\ln(p_g)$	-0.138^{***}	-0.125^{***}
	(0.001)	(0.001)
ln(hp/wt)×t	-0.034^{***}	-0.032^{***}
	(0.0001)	(0.0001)
$\ln(p_g)\times t$	0.049^{***}	0.045^{***}
	(0.0001)	(0.0001)
t	-0.126^{***}	-0.117^{***}
	(0.0004)	(0.0004)
ln(annual miles)	-0.011^{***}	-0.014^{***}
	(0.00004)	(0.0001)
Constant	2.799^{***}	2.757^{***}
	(0.002)	(0.011)
Additional controls	NO	YES
$el\hat{as}t._{2000}$	8.196	8.264
$el\hat{as}t._{2017}$	0.781	0.529
$\%\Delta el\hat{as}t.$	-90%	-94%
Observations	3,884,720	3,884,720
R^2	0.092	0.173
Adjusted R^2	0.092	0.172
Note:	*p<0.1; **p<0.05; ***p<0.01	

The attributes included in the regression are miles per gallon, horsepower-per-weight, and volume measured as wheelbase × width × height. The effect of the decreasing cost of turbochargers is controlled with a time trend. Additional validation of our prediction is in the Appendix, where we show that performance and fuel economy are predictive of more transmission gears.

Results are shown in Table 3.3 for four models. Models 1 and 2 allow trucks to have a shift in the intercept compared to cars while models 3 and 4 allow a change in slope as well. Models 2 and 4 control for the 2015 inflation adjusted dollar US annual average price of gasoline per gallon as reported by the US Energy Information Administration.

Turbochargers have a direct effect on performance and fuel economy (through engine downsizing) and the estimation in Table 3.3 shows a large and statistically significant association between turbocharger adoption and these two attributes. However, turbochargers do not directly affect vehicle size and yet are associated with turbocharger adoption as our theoretical model predicts.

An event study for the 2011 increase in fuel economy standards

Next we test our model predictions on an increase in fuel economy standards. Our model predicts that standards will have the largest effect on vehicles with the worst fuel economy in what we call the "leave no car behind" effect. We use 2012, the year that new standards were implemented, for an event study. We use sub-model names for each manufacture to match 275 vehicle model lines for models that were made in 2005. These model lines have a range from 1988 to 2019 in our data, though most do not span all years. For each vehicle in the data we compute the average fuel consumption for 2005 model year vehicles within each vehicle model line.

Four models of an event study are shown in Table 3.4. Models 1 and 2 regress 2015 average fuel consumption, a time trend, a treatment period dummy, and the interaction of baseline fuel consumption and the treatment dummy. Model 2 additionally controls for log

Table 3.3: Turbocharger adoption regression results

	Dependent variable:			
	Turbo (yes/no)			
	(1)	(2)	(3)	(4)
ln mpg	0.254***	0.207***	0.194***	0.152***
	(0.015)	(0.015)	(0.023)	(0.023)
ln hp/wt	0.179***	0.169***	0.157***	0.143***
	(0.010)	(0.010)	(0.015)	(0.015)
ln volume	0.083***	0.074***	0.160***	0.149***
	(0.011)	(0.011)	(0.025)	(0.025)
time trend	0.013***	0.017***	0.020***	0.023***
	(0.0003)	(0.0004)	(0.001)	(0.001)
ln mpg × truck			0.214***	0.200***
			(0.031)	(0.031)
ln hp/wt × truck			0.038*	0.051**
			(0.021)	(0.021)
ln vol. × truck			-0.038	-0.041
			(0.028)	(0.028)
time trend × truck			-0.013***	-0.013***
			(0.001)	(0.001)
truck indicator	-0.083***	-0.090***	-0.420***	-0.350***
	(0.006)	(0.006)	(0.105)	(0.104)
2015$/gal		-0.062***		-0.058***
		(0.003)		(0.003)
constant	-0.382***	-0.134***	-0.447***	-0.226***
	(0.047)	(0.048)	(0.072)	(0.072)
Observations	32,003	32,003	32,003	32,003
R^2	0.166	0.178	0.182	0.192
Adjusted R^2	0.166	0.178	0.181	0.192
Note:				*p<0.1; **p<0.05; ***p<0.01

gasoline price in 2015 dollars. Models 3 and 4 replace the 2005 level of fuel consumption variable with model-line fixed effects. The coefficient of interest in all models is the coefficient on the interaction term.

All models show that fuel consumption was reduced on average after 2012 by about 40%. Models 1 and 2 estimate coefficients on baseline fuel consumption near 1, indicating, as expected, that vehicles with higher fuel consumption continue to have higher fuel consumption. All four models give a negative value for our coefficient of interest, indicating those vehicle model lines with lower baseline fuel economy improved fuel economy more than others. These effects are in line with our theoretical model and are significantly different from no effect at the 0.01 level.

Table 3.4: Event study regression results

	Dependent variable:			
	log fuel consumption			
	OLS		*vehicle model fixed effects*	
	(1)	(2)	(3)	(4)
ln mpg in 2005	0.937***	0.937***		
	(0.007)	(0.007)		
year>2011	0.407***	0.396***	0.396***	0.387***
	(0.046)	(0.046)	(0.042)	(0.042)
year count	-0.001*	0.002***	-0.0001	0.002***
	(0.0003)	(0.0004)	(0.0003)	(0.0004)
2015$/gal		-0.029***		-0.026***
		(0.003)		(0.003)
ln mpg in 2005 × year>2011	-0.106***	-0.107***	-0.104***	-0.105***
	(0.015)	(0.015)	(0.014)	(0.013)
constant	0.179***	0.255***		
	(0.022)	(0.023)		
Observations	3,348	3,348	3,348	3,348
R^2	0.868	0.871	0.148	0.174
Adjusted R^2	0.868	0.871	0.072	0.100
Note:		*p<0.1; **p<0.05; ***p<0.01		

In summary, we find robust evidence for biased-technological change. We also are able to validate the predictions of our model using reduced form estimation of technology diffusion ("rich get richer" effect) and the response of vehicle models to standards ("no car left behind" effect).

3.3 Engineering and design

Underlying the production of the attributes are are design choices. Manufactures designing vehicles to have a certain power and fuel economy have many options available to them to achieve the desired outcome. These options include discrete technologies, choice of materials, placement and interaction of components, and shape.

Figure 3.6 illustrates how design choices, shown in the center column, determine acceleration and fuel economy. It illustrates three categories of design choices: those that affect only acceleration (turbo- and super-chargers); those that increase both attributes (engine efficiency, light-weighting, aerodynamics, advanced transmissions); and those that have an opposite effect on the attributes (engine displacement and tuning). The figure was created in discussion with automotive engineers at the US Environmental Protection Agency who have extensive experience testing and modeling design choices and technology options on vehicles.

Not shown on the figure are energy consuming technologies unrelated to propulsion that can reduce fuel economy, but little or no effect on power. Some examples are stereo systems, air conditioning, and computing power for self-driving technologies. While this figure is not comprehensive, engine displacement and tuning are the only design choices that we are aware that increase acceleration at the expense of fuel economy.[21]

Small changes in engine displacement can be considered to have nearly no effect on the marginal cost of producing vehicles.[22] This makes engine displacement the most obvious

[21]Transmission upgrades may be another design choice where a manufacturer can increase one of these attributes at the expense of the other. However, manufactures typically choose to use transmissions to improve both.

[22]The fixed cost of design and manufacturing changes may be substantial. However, in this analysis we do not take into account fixed costs.

Figure 3.6: Directed causal graph of performance and fuel economy

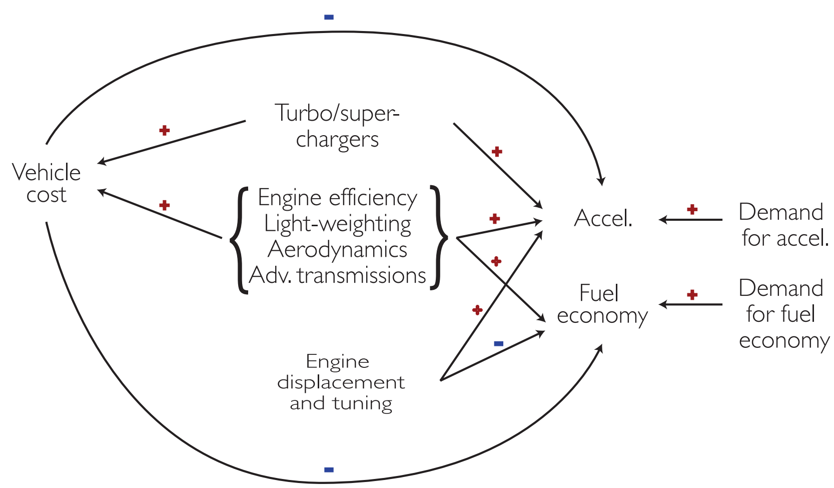

Note: "+" and "−" indicate the sign of the effect of increasing or adding the given vertex.

driver of MRAS, though not the only possible contributor. To the extent that the other engineering levers favor one attribute relative to other levers, substituting one for the other (or dialing one up and the other down) may allow for a trade-off. For example, a manufacture may choose between turbo-charging the engine, which only improves acceleration, or aerodynamics technology of a similar cost. If the aerodynamic improvements of the vehicle aid acceleration less than the turbocharger, then there is an opportunity for a trade-off.

These engineering substitutions are known to take place. Adding a turbocharger and reducing engine displacement, known as "engine downsizing" in the engineering literature, has been a common method used by manufactures to increase fuel economy with a smaller engine while maintaining acceleration using a turbo charger (Shahed 2009). As the price of turbochargers has fallen and they have become widely deployed in the fleet, there are fewer opportunities to make this substitution.

Technology adoption also has important interactions with other engineering design choices. Turbochargers, along with other technologies, substantially dampen the effect that small changes in power have on fuel economy while maintaining displacement's effect on acceleration. Taken together, these two effects—foreclosed pathways for a trade-offs once technology becomes ubiquitous and technology's trade-off dampening effect—would imply a substantially decreased trade-off over time. Reducing engine displacement appears to have become a less effective way to improve fuel economy.

Market factors, which are included in Figure 3.6, also influence acceleration and fuel economy. Consumers want both attributes, while vehicle technologies increase their cost.[23] Manufacturers then are expected to balance engineering inputs that increase costs with consumer demand and add technologies so long as their costs are less than the additional benefit it provides consumers.

We use evidence from a physics-based full vehicle simulation model, the Environmental Protection Agency's Advanced Light-Duty Power-train and Hybrid Analysis (ALPHA) tool (Dekraker et al. **(year?)**), to investigate how certain technologies have changed the effect of fuel displacement. Figure 3.7, which is reproduced from Moskalik et al. (2018) data, summarizes this change in simulated vehicles with representative drive-trains.

Each of the simulated vehicles vary only by drive-train technology. Each of the representative drive-trains are placed in a simulated mid-sized sedan. This allows the simulation to maintain the same road-load[24] across all simulated vehicles.The five power-trains included a 1980s carbureted engine with a three-speed automatic transmission; a 2007 Toyota port fuel-injected (PFI) engine coupled to a five-speed transmission; a 2013 GM gasoline direct injection (GDI) engine coupled to a six-speed transmission; a 2017 Honda turbocharged engine coupled with an eight-speed transmission; and, a future Ricardo 24 bar turbocharged engine with cooled engine gas recirculation (EGR) coupled to an advanced eight-speed transmission. The model is calibrated to each drive-train technology package using data gathered from real-world vehicles in a laboratory setting. Data on individual components and tech-

[23]We have omitted an edge (causal arrow) from engine displacement to vehicle cost as small changes in engine displacement have negligible effect on the marginal cost of a vehicle. This would not be the case for large changes in engine size, which would require much more steel and aluminium in the engine but also changes in the chassis to support a larger motor.

[24]Road-load is the combined measure of a vehicle's weight and aerodynamic resistance.

nology are captured through extensive vehicle testing using a vehicle's on-board sensors and sensors added to the vehicle.

Figure 3.7: ALPHA tool simulations

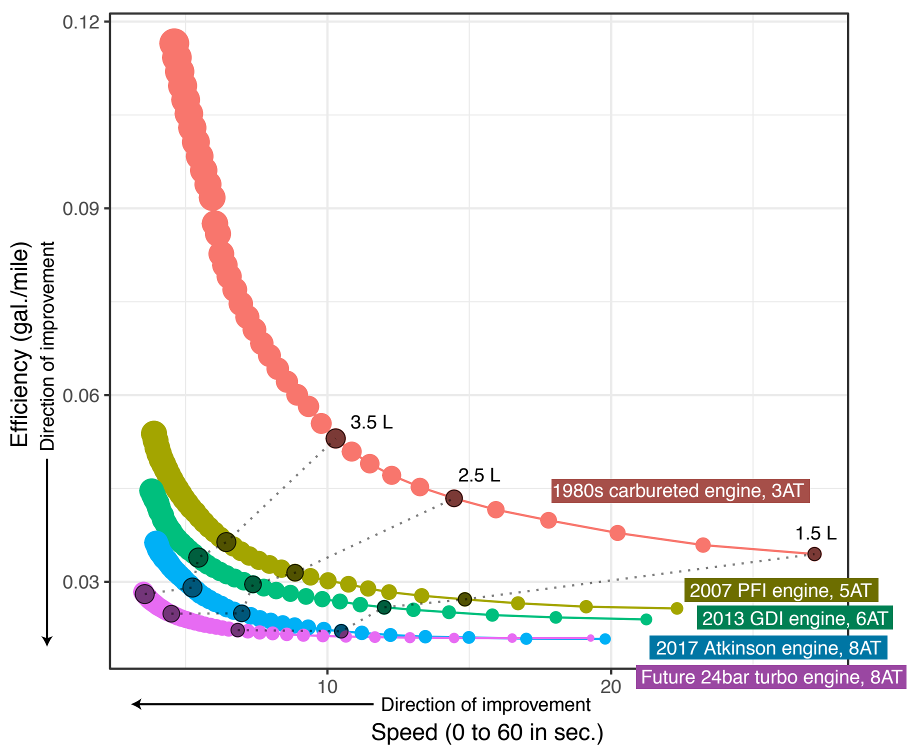

In the Figure each line plots the performance (0 to 60 miles-per-hour time in seconds) and fuel consumption (gallons per mile). Each line plots simulations for drive-trains that are representative of technology used at a given time, varied only by the displacement of the modeled engine, which is represented by the size of the bubble. Three displacement volumes—1.5, 2.5, and 3.5 liters—are marked with dashed lines connecting them for convenience. The lowest line represents a near-future drive-train using technology that is currently available taken from different production vehicles. Note that the direction of improvement is towards the origin on both axes.

Several observations are made plain by the Figure. 1) For any given displacement, engines are substantially faster in recent-year representative drive-trains compared to the 1980s drive-train. The smallest engines have sped up the most. In percent changes, time to 60 miles per hour has dropped by approximately half for most engine displacements comparing the 1980s simulated vehicle to the 2017 vehicle. In levels this means a 17 second improvement for 1.5 litre engines compared to a 5 second improvement for 3.5 litre engines. 2) For any given engine displacement, the drive-train is more fuel efficient. In levels, this change has been greatest in the largest engines. 3) The marginal change in fuel economy relative to performance caused by displacement is substantially lower in recent-year representative vehicles simulated, and the relative marginal effect is declining (or flattening) over time.

On average across the sweep of engine displacements, the 1980s simulated vehicle has a trade-off elasticity of 1.9 (a 1% increase in power decreases fuel economy by 1.9%) while the 2017 simulated vehicle has trade-off elasticity of 1.3—a 29% decrease.

This last point implies that it takes a much larger reduction in displacement volume to improve fuel economy by the same amount, and this reduction comes at a higher cost to performance. If displacement were the only pathway for trading off between the two attributes, we would then conclude that the trade-off has substantially flattened. As previously discussed, there are other potential pathways that could create a trade-off, but displacement is the most obvious cause of a large trade-off.

3.4 Conclusion and discussion

We have attempted to clarify the mechanics of attribute trade-offs through the physical interactions of design choices and how designs are determined by consumer preferences.

Using an engineering understanding of physical relationships between power and fuel economy as mediated through particular design choices, we find that design choices fall into one of three categories. A design choice may improve both attributes, improve one attribute and have no effect on the other, or improve one attribute at the expense of the other. We find that many of the design choices available fall into the first two categories, with the exception of engine displacement. We further find that engine displacement itself, while still important for determining power, has less of an impact on fuel economy than it once did.

However, when manufacturers design vehicles, all design options are on the table, which expands the possible achievable bundles for a given iso-cost curve. So, the changing effect of engine displacement on fuel economy relative to power is not definitive evidence for a changing MRAS overall.

This engineering foundation inspires our development of a theoretical model that can be taken to data. Our model of firms designing vehicles for consumers with heterogeneous preferences allows us to test if technology has changed the marginal rate of attribute substitution and predicts how the vehicle fleet will react to new technologies or more stringent standards.

Applying our model to market data, we find robust evidence that the Marginal Rate of Attribute Substitution (MRAS) has substantially fallen. The estimates across a wide range of consumer preference are remarkably consistent—a decrease of about a quarter. This implies that reducing power has a much smaller effect on fuel economy than it once did, keeping manufacturing costs equal. This reduces the desirability of using displacement reductions to improve fuel economy; instead, adding fuel-reducing technology becomes a more attractive way to improve fuel economy. The decline in MRAS is driven by technological change that is strongly biased towards the production of power, with all of the change in the relationship coming from the cost parameter on power. Our model anticipates that this change has an amplifying, "more of what I like," effect on vehicle attributes in the fleet.

When it comes to discrete technologies, we find that technology is not adopted randomly throughout the vehicle fleet. Instead, adoption is a function of consumer preferences and therefore the baseline level of attributes of vehicles. Further, baseline attributes that are not physically related to a particular technology are both predictive of adoption and are affected by it through the response of consumers—a "rich get richer" effect, where vehicles with the highest baseline attributes increase all attributes with adoption. Using the diffusion of turbochargers as an example, we find evidence in support of the theory. This result can inform our expectations of how new technologies will affect the composition of the vehicle fleet and has implications for the estimation of trade-offs using vehicle data. Without accounting for production costs or consumer preferences, a researcher cannot assume, as has been implicitly assumed in the past, that technology and costs are distributed evenly across attributes. Fitting a line through attribute data cannot measure an iso-cost curve and therefore is not able to measure attribute trade-offs.

Finally, standards that are applied over fleet averages have a greater impact on vehicles with the worst fuel economy. Intuitively, the most inefficient vehicles are the least costly to improve and therefore see the most improvement. We use the change in standards in 2012 to preform an event study. We find strong evidence to support this "leave no cars behind" effect.

This paper has brought both engineering and economic insights to the stylized fact that improvements in fuel economy must come at the expense of vehicle performance. The engineering assessment shows that innovation has changed the nature of the trade-off for the design choices where the trade-off is large. The economic assessment suggests that the diffusion of these technologies into the fleet is not uniform: take heterogeneous consumer preferences into account and the relative benefit of changing attribute mix or adding technologies when designing vehicles. Changes to vehicle technology and the fleet in recent years suggest that we are closer to a world where it is most profitable to add technologies to meet fuel economy standards than to do so by reducing power.

CHAPTER 4

A NOTE ON CROP ABANDONMENT AND REVENUE INSURANCE

With David Hennessy

Does crop revenue insurance affect crop abandonment? We develop a simple model of stochastic crop production and farmers who incur costs and update their expectations through the season. At any time the farmer can avoid further costs by abandoning their crop. We find that a simplified representation of revenue insurance in the US does not change the abandonment behavior of risk-neutral profit-maximizing farmers. We then discuss other sources of a wedge between behavior with and without insurance and argue that if farmers are risk averse crop insurance improves welfare by pushing abandonment decisions towards those made by a risk neutral farmer.

Crop insurance in the United States has grown substantially since the 1994 Farm Bill. In 2019, federally subsidized insurance policies covered roughly 90% of eligible acres (USDA 2019). Two concerns often follow the use of insurance: adverse selection and moral hazard. Adverse selection occurs due to information asymmetries in the riskiness of the insured. Moral hazard occurs when the insured change their behavior because of the risk reducing effects of insurance. Moral hazard could impact social welfare—reducing or increasing it—if insurance causes farmers to change their behavior either further from or closer to decisions the social planner would make with full information.

In this essay we explore the potential for moral hazard through crop abandonment. We develop a model of crop production and a farmer's decision to abandon their crop. As a crop grows, farmer's observe the progress of crop growth and update their expectations for end of season yield. Farmers also observe prices in the market and update their expected revenue per unit of output. Through the season farmers incur costs in labor and inputs of continuing to husband their crop. If updates are such that expected income is less than the expected cost of continuing husband their crop—conditional on optimal choices in the future–it is then optimal for farmers to abandon their crop.

Conceptually, we model crop production and the decision to abandon similar to an American put option, a financial derivative. Like an American put option a farmer makes an investment in an asset (planting a crop), observes the stochastic progression of their asset (crop growth), and may exercise the option (abandoning the crop) and receive the strike price (insurance indemnity payment or zero) at any point until the expiration of the option (harvest). If purchase of insurance increases the strike price of the crop-option, and has the potential to increase the probability that the farmer will exercise the option and abandon. If farmers abandon their crop with insurance when they would not abandon under without insurance, then this over abandonment creates an inefficiency in the market. Assuming commodities markets clear, the inefficiency would result in higher prices, lower quantities, and excessive crop production at farms that would be extra-marginal without the inefficiency. Additionally, over-abandonment would increase the actuarially fair price of insurance, which would increase the cost of insurance. However, our model shows no change in abandonment with insurance with the model of insurance that we specify.

We specify an indemnity payment based on our understanding of United States Department of Agriculture Risk Management Agency (RMA) rules for mid-season indemnities, which requires payments to be discounted by the expected value of the crop at the time of abandonment. Crop insurance products and rules for adjusters are developed by USDA RMA and sold by approved private providers. These rules are detailed in Loss Adjustment Manuals (LAM) that are revised periodically.[1] For our purposes, LAM provide guidance to adjusters for loss claims given before the end of the season. In recent years at least, LAM require farmers to receive written consent before a crop is replanted, put to another use, abandoned, or destroyed prior to harvest and making a claim on the loss of the crop. The

[1] A full history of these manuals may be found at https://www.rma.usda.gov/handbooks/25000/

Figure 4.1: Evidence of abandonment for winter wheat in Kansas

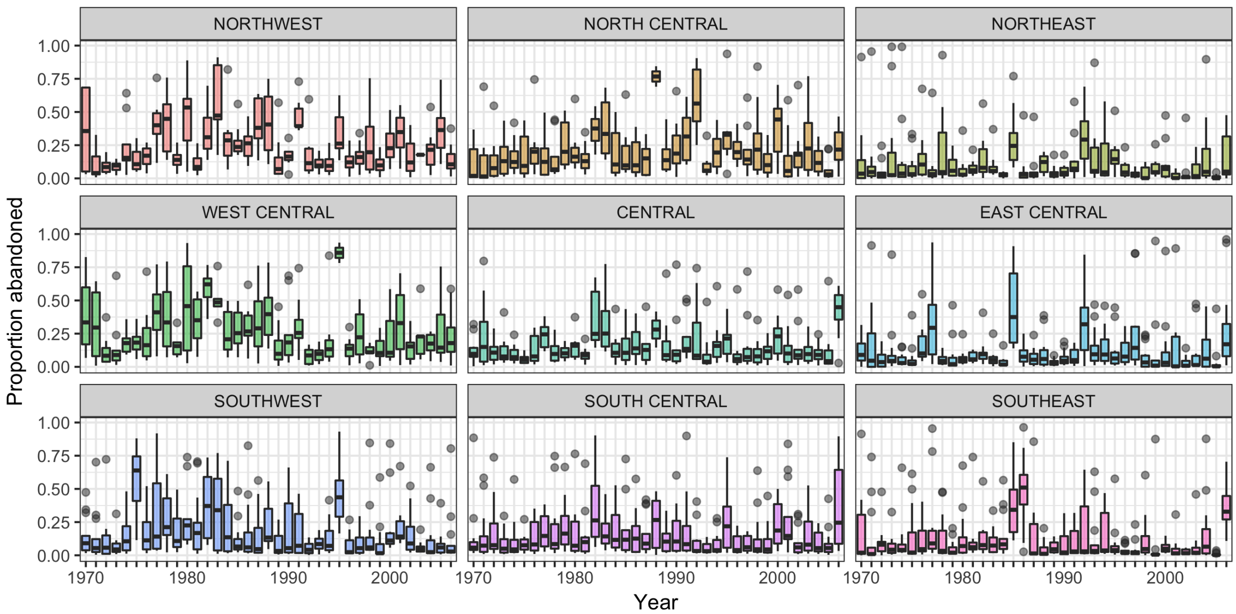

Note: This box-plot shows the distribution county-level proportion of unharvested acres to planted acres ($\frac{planted-harvested}{planted}$)*. These observations have been grouped by USDA defined Agricultural Districts and by year. These data are from the NASS Ag Statistics Database.*

loss adjuster's appraisal of the value of the crop is then given by the "the maximum potential production that could be expected with normal weather conditions and proper crop care from the time of the appraisal through the completion of harvest." This is this subtracted from the revenue or yield guarantee depending on the type of insurance purchased by the farmer. From our discussions with officials at USDA, "maximum potential production" is taken to mean the expected harvest if the farmer were to use best management practices until harvest. Evidence for this interpretation is a rule that allows the insurer to require the insured to maintain a small plot to determine the yield if the farmer had not abandoned in the event there is disagreement over the value of the crop to discount from the guarantee.

The literature on factors that contribute to crop abandonment is sparse. It has primarily focused on regression analysis of adverse intraseasonal events on yield and abandonment outcomes. For example Mulungu and Tembo (2015) find evidence that lack of moisture early in the season was associated with crop abandonment in Zambia.

Crop abandonment is a significant issue across the world especially for regions on the climatic boundary for a given crop. Some examples are cotton in West Texas and Oklahoma, wheat in western Kansas, corn in the Dakotas, and maize, ground-nuts, and cotton in the semi-arid regions of Africa. Figure (4.2) shows the proportion of panted less harvested acres to planted acres within a county, in each season from 1970 to 2007, and grouped by NASS Agricultural District. While it is likely that this ratio includes acres that were re-planted, and therefore double counted for our purposes, these data are evidence that crop abandonment is not uncommon: 8% of observations show a proportion that is greater than 50% and 46% has a ratio over 10%.

4.1 A simple model of a growing season

We develop a simple model of a growing season and a farmer's decision to abandon their crop within the season. By abandonment we mean: a farmer does not engage in the necessary activities to continue to husband the crop that lead to eventual harvest. The aim of the model is to capture the stochastic process of a crop's development through a growing season, the intraseason costs of continuing production, and the choice farmers face to continue to incur production costs or abandon after updating their expectations after observing crop development. After planting, weather and other conditions will effect the growth of the crop. If growing conditions are good, the crop may grow faster and become more resilient to poor conditions later on in the season. Observing the progress of the crop the farmer will update their beliefs of the yield they expect at harvest time and this updating allows them to adjust their production plan. If instead the crop has a poor start or crop prices fall after planting, then the farmer may choose to abandon instead of continue to incur costs. For simplicity the intraseason input decisions are assumed to be fixed and reflect well established "best practices" for the crop, geography, and climate of the farm.

Figure 4.2: A three period binomial tree model of crop production

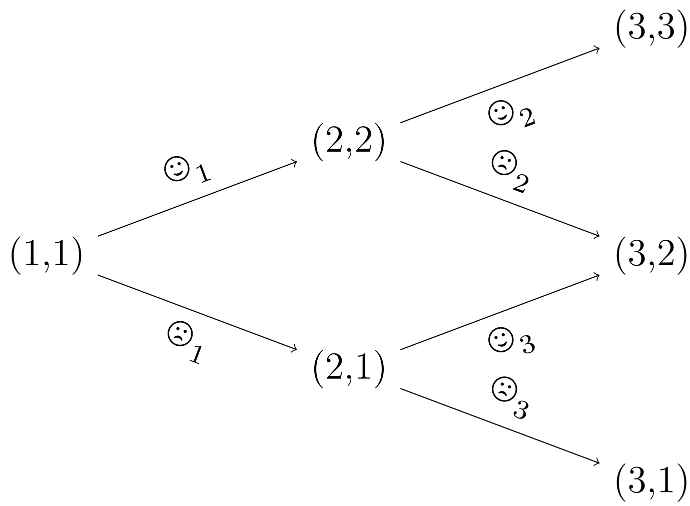

4.2 Setup and definitions

Consider a simple three period model of a single season of crop production illustrated in Figure (4.2). Each node is indexed for time t and order j in the lattice, (t, j). At the end of the season ($t = 3$) the value of the crop is either high (v_3), medium (v_2), or low (v_1) —corresponding to nodes (3,3), (3,2), and (3,1) respectively—depending on how well the crop has grown and how the market price for the crop has changed. Assume that each of the end of season values, v_n, are equal intervals apart.

Once planted ($t = 1$) the crop may do well or poorly by arriving at mid-season nodes $(2, 1)$ or $(2, 2)$ with probability $\smiley_1$ or $\frownie_1$, respectively, where $\frownie_n = 1 - \smiley_n$. If the farmer continues to care for the crop in period 2, the crop may improve or decline relative to the mid-season outcome with corresponding probabilities $\smiley_n$ and $\frownie_n$ shown in the figure. For simplicity, an improvement from node $(2, 2)$ or a decline from node $(2, 1)$ reach the same terminal node $(3, 2)$, which is a standard simplification from binomial tree analysis (see Dixit and Pindyk 1994).

At the beginning of the season, the farmer knows state transition probabilities and the probability of each final outcome (e.g. the probability of v_2 is $\smiley_1\frownie_2 + \frownie_1\smiley_3$). There is a cost at each period, $k_t \geq 0$, to continue to husband the crop and move forward a period. The final period (harvest) cost, k_3, must be incurred by the farmer before collecting the value of the crop in the final period, which is known with certainty. In any period including the last period, the farmer may choose to not continue production activities and avoid the cost k_t.

Denote a generic insurance indemnity payment function as $I(t, j)$. If the farmer has crop insurance, they may receive an indemnity payment, $I(t, j) \geq 0$, which is a function of the node where the farmer makes a claim. The cost of insurance is sunk at the beginning of the season and therefore not included in this model.

In the American context, Federal crop insurance indemnity payment rules are determined by the United States Department of Agriculture. While these rules are quite detailed (see RMA LAM), our simplified representation of revenue insurance available in the United States takes the following form:

$$I^r(t, j) = \max\{r - \mathbb{E}_{t,j}[v],\ 0\} \tag{4.1}$$

where r, the revenue guarantee, is a proportion of the expected value at the beginning of the season, or:

$$r = \lambda \mathbb{E}_{(1,1)}[v]. \tag{4.2}$$

$\lambda \in [0, 1]$ is the level of the insurance policy (e.g. 75%). This indemnity function simply says that if the expected value of the crop is less than the revenue guarantee, the payment to the farmer will be the guarantee less the expected value. For brevity, we refer to the indemnity function in Equation 4.1 as "revenue insurance" to distinguish it from any other possible indemnity function. The revenue insurance function at nodes $(2, 2)$ and terminal nodes are given in equations (4.3) and (4.4) respectively.

$$I^r(2, 2) = \max\{r - \smiley_3 v_2 - \frownie_3 v_3,\ 0\} \tag{4.3}$$

$$I^r(3, j) = \max\{r - v_j,\ 0\}. \tag{4.4}$$

Note that there can be no indemnity payment at node $(3,1)$ as it is the upper bound of possible crop value outcomes and the highest r could be, with a degenerate distribution that arrives at node $(3,3)$ with a probability of 1, would then receive an indemnity payment of zero. Notice that the mid-season indemnity payment is equal to the expected indemnity payment from the mid-season node. For example, if v_2 is such that there is an indemnity payment at the middle terminal node $(3,2)$, then the expected indemnity payment from $\mathbb{E}_{(2,2)}I = r - \mathbb{E}_{(2,2)}v$, which is the same as Equation 4.3.

Assume, conditional on arrival at node $(3,2)$, the middle terminal value, the farmer will not abandon but may receive an indemnity payment. Receiving an indemnity payment will occur at $(3,2)$ whenever $v_2 < \lambda\mathbb{E}_{(1,1)}[v]$, which implies negative skewness of the first period underlying distribution (see Appendix). Further assume that at the lowest terminal node, $(3,3)$, that there is always an indemnity payment under revenue insurance (or equivalently $v_3 < r$). These assumption reduce the number of cases to consider without a cost to generalizability. We assume that the farmer is profit maximizing and, for now, risk neutral.

4.3 Abandonment decisions with and without insurance

Result 1: *The incentive to abandon in the last period is the same with our without insurance.* Without insurance, the condition for abandoning the crop in the last period is that the cost of harvest, k_3, larger than the value of the crop, which is known with certainty. Using a generic insurance indemnity function, the payout is the same once yields are known with certainty. The condition for abandonment in the last period is:

$$\underbrace{I(3,j)}_{\text{rev. when abandon}} > \underbrace{I(3,j) + v_j - k_3}_{\text{rev. when harvest}} \implies k_3 > v_j. \tag{4.5}$$

Equation 4.5 shows that the decision to abandon in the last period does not depend on $I(3,j)$ and is identical to the abandonment condition without insurance so long as $I(3,j)$ is the same if abandoning immediately before harvest when yield is known or harvesting, which is the case for revenue insurance.

To determine if farmers are more likely to abandon mid-season, we must determine the conditions for abandoning at $(2,2)$. Consider the two possible cases for the relationship between k_3 and v_3 from Result 1. In Case 1, $k_3 < v_3$, implying that the farmer will not abandon at $(3,3)$ with or without insurance. In Case 2, $k_3 > v_3$, implying that the farmer will abandon at node $(3,3)$ with or without insurance.

Case 1: *When harvest costs are below worst-case harvest value ($k_3 \leq v_3$).*

In this first case, the farmer will *not* abandon in the last period regardless of insurance. Using using a generic indemnity function, the condition for abandonment with insurance at $(2,2)$ is:

$$\underbrace{I(2,2)}_{\text{profit when abandoning}} > \underbrace{\smiley_3\left(v_2 + I(3,2)\right) + (\frownie_3)\left(v_3 + I(3,3)\right) - k_2 - k_3}_{\text{expected profit when continuing}}, \tag{4.6}$$

and the condition for not abandoning with no insurance is:

$$\underbrace{0}_{\text{rev. when abandoning}} < \underbrace{\smiley_3 v_2 + (\frownie_3)\, v_3 - k_2 - k_3}_{\text{expected profit when continuing}} \implies k_2 + k_3 - (\frownie_3)v_3 < \smiley_3 v_2. \tag{4.7}$$

Plugging in the revenue insurance indemnity function—Equation 4.1—Equation 4.6 becomes

$$\begin{aligned} r - \smiley_3 v_2 - (\frownie_3)v_3 &> \smiley_3\left(v_2 + \max\{r - v_2,\ 0\}\right) + (\frownie_3)(v_3 + r - v_3) - k_2 - k_3 \\ &\implies \smiley_3 r + k_2 + k_3 - (\frownie_3)\, v_3 - \smiley_3 \max\{r - v_2,\ 0\} > 2\smiley_3 v_2 \end{aligned} \tag{4.8}$$

Case 2: *When harvest costs exceed worst-case harvest value ($k_3 > v_3$).*

In the second case, the farmer will abandon in the last period regardless of insurance. Again, using using a generic indemnity function, the condition for abandonment with insurance at $(2,2)$ is

$$\underbrace{I(2,2)}_{\text{profit when abandoning}} > \underbrace{\smiley_3\left(v_2 - k_3 + I(3,2)\right) + (\frownie_3)I(3,3) - k_2}_{\text{expected profit when continuing}}, \tag{4.9}$$

but not without insurance is

$$\underbrace{0}_{\text{profit when abandoning}} < \underbrace{\smiley_3\left(v_2 - k_3\right) - k_2}_{\text{expected profit when continuing}} \tag{4.10}$$

$$\implies \smiley_3 k_3 + k_2 < \smiley_3 v_2. \tag{4.11}$$

These are the same sign conditions for net revenue with and without insurance, however, expected costs and revenue have changed to reflect no abandonment at any node in the last period. Plugging in Equation 4.1 into Equation 4.9 we get

$$r - \smiley_3 v_2 - (\frownie_3)\, v_3 > \smiley_3 \left(v_2 - k_3 \smiley_3 \max\{r - v_2,\ 0\}\right) + (\frownie_3)\,(r - v_3) - k_2$$
$$\implies \smiley_3 r + \smiley_3 k_3 + k_2 - \smiley_3 \max\{r - v_2,\ 0\} > 2\smiley_3 v_2. \quad (4.12)$$

Note that to find the conditions for under-abandonment, we only need to reverse the inequalities in both of the cases.

Result 2: *Revenue insurance does not change the incentive to abandon mid season if farmers are profit maximizing and risk neutral.*

We provide a proof by contradiction. Result 1 shows that this is the case in the lat period. Therefore, we only need to find the conditions for over or under abandonment with insurance at $(2, 2)$ in each of the cases.

Case 1: If $\smiley_3 \max\{r - v_2,\ 0\} > 0$ then Equation 4.8 becomes

$$\smiley_3 r + k_2 + k_3 - (\frownie_3) v_3 - \smiley_3 (r - v_2) > 2\smiley_3 v_2 \implies k_2 + k_3 - (\frownie_3) v_3 > \smiley_3 v_2 \quad (4.13)$$

which contradicts Equation 4.7.

If $\smiley_3 \max\{r - v_2,\ 0\} = 0$, then $v_2 \leq r$, which when applied to Equation 4.7 implies $r \geq v_2 > \frac{k_2 + k_3 - (\frownie_3) v_3}{\smiley_3}$. Equation 4.8 becomes

$$\smiley_3 r + k_2 + k_3 - (\frownie_3) v_3 > 2\smiley_3 v_2 \implies r > \frac{1}{2} r + \frac{1}{2} \underbrace{\frac{k_2 + k_3 - (\frownie_3) v_3}{\smiley_3}}_{<r} > v_2 \quad (4.14)$$

which is a contradiction.

Case 2: If $\smiley_3 \max\{r - v_2,\ 0\} > 0$ then Equation 4.12 becomes

$$\smiley_3 r + k_2 + \smiley_3 k_3 - \smiley_3 (r - v_2) > 2\smiley_3 v_2 \implies k_2 + \smiley_3 k_3 > \smiley_3 v_2 \quad (4.15)$$

which contradicts Equation 4.11.

If $\smiley_3 \max\{r - v_2,\ 0\} = 0$, then $v_2 \leq r$, which when applied to Equation 4.11 implies $r \geq v_2 > \frac{k_2 + \smiley_3 k_3}{\smiley_3}$. Equation 4.12 becomes

$$\smiley_3 r + k_2 + \smiley_3 k_3 > 2\smiley_3 v_2 \implies r > \frac{1}{2} r + \frac{1}{2} \underbrace{\frac{k_2 + \smiley_3 k_3}{\smiley_3}}_{<r} > v_2 \tag{4.16}$$

which is a contradiction ■

The proof for no under-abandonment follows directly from the above by simply reversing all of the inequalities. This shows that under the risk neutrality assumption, the indemnity function is such that it exactly offsets any opportunity cost of forgone indemnity revenue mid-season with an increase in expected future revenues.

4.4 Can we find a wedge?

In the previous subsection we found that a risk neutral farmer will make the same abandonment decisions with or without revenue insurance. In this subsection we consider what we might be missing that could drive a wedge between abandonment decisions with and without insurance.

What if there are monitoring costs? A key component of the model is the ability of the insurer to verify the state of the crop. It is possible that the costs of collecting information change the behavior of either the insurer or the insured away from full information scenario that we have represented. The incidence of such a cost is not clear however. One possibility is that the insurer will reduce their efforts to verify the claims of the insured near the abandonment threshold because the gains to the insured are small. This argument would bear more weight if the insurer were fully privatized, which is not so in the United States. Coble and Barnett (2013) attribute the high administrative costs of crop insurance to the the cost of monitoring moral hazard and adverse selection at individual farms, suggesting that these costs are not avoided through lax monitoring.

The cost of verification could also fall on the farmer if the farmer is required to take actions to prove the state of their crop. This would increase the cost of abandoning under insurance, causing farmers to under-abandon. This is likely so to some degree. Under RMA rules the farmer must notify the insurer before abandoning their crop and the insurer can require the insured to continue to maintain a small area of the crop to verify how the crop would have grown if it were taken to harvest.

What if farmers are risk averse? By definition a risk adverse farmer values gambles with the same mean less and less as the variance of the gamble increases more and more. In the last period, this has no effect on the farmer's behavior as outcomes are certain regardless of insurance. In the middle period, however, there is a difference from the risk neutral farmer. Without insurance, the farmer in the middle period will value all gambles of the future value of the crop relative to abandonment less compared to the risk neutral version of herself. This means that the threshold for our risk averse farmer to abandon is lower and she will abandon more often.

If we now give this farmer revenue insurance, we will reduce the variance end of the season crop value while keeping the same expected value (as shown above). Insurance will increase the threshold for abandonment towards the risk neutral farmer's. Thus, insurance reduces abandonment towards the risk-neutral level if farmers are risk averse. If social welfare is maximized by maximizing expected value, insurance improves social welfare.

There is large literature that considers the risk aversion of farmers. Binswanger (1981) finds evidence in support of risk aversion among 330 farmers in India in a lab experiment. In the study farmers are increasingly likely to prefer gambles with the smaller spread compared to one with a larger spread but same mean, as the mean, and as the stakes of the gamble increase. This study was replicated in China in Kachelmeier and Shehata (1992) with similar results. Studies using surveys such as Friend and Blume (1975), Dohmen et al. (2005), Szpiro (1986), Cichetti and Dublin (1994), Riesgo, and Arriaza (2002) find evidence that US farmers exhibit a wide variety of risk preferences, but are genearlly risk averse.

Perhaps farmers respond to risk in some way that is not captured by expected utility theory. Du, Feng, and Hennessy (2016) and Doidge (**?**) investigate the risk preferences of

US farmers. Du, Feng, and Hennessy (2016) find that farmersâ crop insurance choices are inconsistent with expected utility maximization which would predict that farmers would either choose insurance plans with the largest subsidies or at the highest level while Doidge (2019), using a choice experiment over insurance coverage levels with farmers as subjects, estimates parameters that are consistent with a risk-aversion model that includes loss aversion. It is unclear, however, how loss aversion would affect farmer behavior in our setting.

What if there are interseasonal dynamics? The behavior of the insurer or the insured may incur costs in later seasons. For example, farmers may be reluctant to abandon if abandonment increases their insurance premium in future seasons. Evidence of this "pseudo-deductible" problem is presented in Braun et al. (2006), which finds evidences that homeowners are reluctant to make claims on their homeowners insurance because they are concerned that claims will increase their premium. There is no literature we are aware of regarding crop insurance that considers interseasonal dynamics. So, we leave this question to be answered by future research.

4.5 Conclusion and discussion

We show that revenue insurance is not likely to induce sub-optimal abandonment when farmers are profit maximizing and risk neutral. Why? 1) past history of yields is well documented and so the insurer and insured have common expectations; 2) actions by the farmer are assumed to be (and likely are in practice) observable by the insurer which removes moral-hazard regarding production input decisions; 3) the value of expected yield is subtracted from the revenue guarantee if the farmer abandons so that with insurance expected increased income if abandoning is exactly offset by increased expected revenue at harvest.

In this paper we isolated the decision to abandon, and treat other input and crop choice decisions as fixed. However, there is a large literature that estimates the change in input levels that are attributed to crop insurance. For example, Babcock and Hennessy (1996) find evidence that insurance leads to a decrease in fertilizer use for corn producers in Iowa. Seo, Mitchell, and Leatham (2005) use a mathematical programming model to solve for optimal fertilizer application for a cotton-sorghum farm in Texas. They find depending on a farmer's risk aversion that crop insurance changes fertilizer application from -6 to 3%. Importantly, in this work the authors assume that the farmer is unconstrained by the insurer regarding the optimal input level. Roberts, Key, and O'Donoghue (2006), use farm-level administrative data in Iowa, North Dakota, and Texas over a time period where they observe farmers switching from no-insurance to insurance. In a difference in difference framework that controls for farm level fixed effects, they find modest decreases in yields across three farm types associated with crop insurance indicating a modest effect of moral hazard (-.48 bushels per acre for Texas, for example).

However, insurers are able to monitor farmers production decisions for moral-hazard and adjust payments accordingly. RMA rules dictate that indemnity payments only be given if the farmer can demonstrate that they managed the crop using best management practices. To verify best management practices, the insurer may require receipts, compare the farm to others in the same region, or use bench-marking documents. Revenue insurance may not have always worked as well as what we have modeled here. Glauber (2004) identifies the 1980s and 90s as a period when crop production expanded to areas where there was little data and suggests that this lead to the high loss-ratios of the period. Further, Glauber attributes much improved loss-rates by the late 1990s to better data and the reform efforts enacted by Congress that increased monitoring.

Finally, our model highlights an interesting possibility: risk averse farmers may reduce abandonment with insurance. If society's goal is to maximize expected benefits, then insurance is welfare improving. This has the potential to change the integration of empirical results that find behavior farmer behavior changes with insurance and deserves further study.

APPENDICES

APPENDIX A

CHAPTER 1

Mathematical appendix

Proposition 1

Proof. Plugging the optimal r^* from Equation 1.7, the regulator's problem in Equation 1.4 becomes

$$\int_0^{B^*} D(k)(B-k)dt \tag{A.1}$$

The first-order condition is

$$\frac{\partial TB}{\partial p} = D(B^*)(B^*-B^*) + \int_0^{B^*} D(k)\frac{\partial B}{\partial p}\ dk = \frac{\partial B}{\partial p} = 0 \tag{A.2}$$

Taking the first-order condition of Equation 1.21 for $t=0$

$$\frac{\partial B(0)}{\partial p} = -\frac{\partial c(p,\text{☼},0)}{\partial p} + bT\frac{\partial g}{\partial c}\frac{\partial c}{\partial p} + \lambda_{\text{☠}}\frac{\partial B_{\text{☠}} - B_{\text{☼}}}{\partial p} = 0. \tag{A.3}$$

Notice that $\frac{\partial c_j^*}{\partial p} = \frac{c_j^*+1}{p}$ and $\frac{\partial \dot{q}_j}{\partial p} = \beta/p$.

Let

$$\Delta\tilde{c} = E[\int_0^T \tilde{c}\ dt|s_0 = \text{☼}] - E[\int_0^T \tilde{c}\ dt|s_0 = \text{☠}] \geq 0. \tag{A.4}$$

Then

$$\frac{\partial B(0)}{\partial p} = -\beta\Big(\int_0^T f_{\text{☼}|\text{☼}}(t)\ dt - f_{\text{☠}|\text{☼}}(T)\Big) + bT\beta/p + \lambda_{\text{☠}}\Delta\tilde{c}/p = 0 \tag{A.5}$$

implies

$$p^* = \Big(bT + \lambda_{\text{☠}}\Delta\tilde{c}/\beta\Big)\Big/\Big(\int_0^T f_{\text{☼}|\text{☼}}(t)\ dt - f_{\text{☠}|\text{☼}}(T)\Big). \tag{A.6}$$

As

$$\int_0^T f_{\text{☼}|\text{☼}}(t)\ dt \leq T \tag{A.7}$$

Therefore, $\lambda_{\text{☠}} > 0$ implies $p^* > b$ □

Proposition 2

Proof. Suppose not. Regulator's first-order condition is

$$\frac{\partial T.B.}{\partial p} = D(V(0))\Big(B(p^*,0)-V(p^*,0)\Big)\frac{\partial V(p^*,0)}{\partial p} + \frac{\partial B(p^*,0)}{\partial p}\int_0^{V(p^*,0)} D(k)\ dt = 0. \quad \text{(A.8)}$$

As shown in the proof of Proposition **??**, $B(p^*,0) > 0V(p^*,0)$. From the proof of Proposition 1 $\frac{\partial B(p^*,0)}{\partial p} = 0$. Further, $\frac{\partial V(0)}{\partial p} > 0$ for all p. Therefore,

$$\underbrace{D(V(0))\Big(B(p^*,0)-V(p^*,0)\Big)\frac{\partial V(p^*,0)}{\partial p}}_{>0} + \underbrace{\frac{\partial B(p^*,0)}{\partial p}\int_0^{V(p^*,0)} D(k)\ dt}_{=0} > 0, \quad \text{(A.9)}$$

which is a contradiction □

Proposition 3

Proof.

$$\begin{aligned} TSB^p - TSB^r =& N^p \times B^p - N^r \times B^r && \text{(A.10)}\\ =& N^p \times \big(B^p - B^r + B^r\big) - \big(N^r - N^p + N^p\big) \times B^r && \text{(A.11)}\\ =& N^p \times \big(\Delta B + B^r\big) - \big(\Delta N + N^p\big) \times B^r && \text{(A.12)}\\ =& N^p \Delta B - \Delta N B^r. && \text{(A.13)} \end{aligned}$$

So,

$$\begin{aligned} TSB^p &\lesseqgtr TSB^r && \text{(A.14)}\\ \Delta B/B^r &\lesseqgtr \Delta N/N^p && \text{(A.15)}\\ \frac{\Delta B/B^r}{\Delta p/p} &\lesseqgtr \frac{N/N^p}{\Delta p/p} && \text{(A.16)}\\ \varepsilon_{B,p} &\lesseqgtr \varepsilon_{N,p} && \text{(A.17)} \end{aligned}$$

□

APPENDIX B

CHAPTER 2

Table B.1: Panel data summary

	State	First year	Last year	N years	N installs	N res. installs	RPS	% oppose RPS
1	AR	2009	2011	3	416	341	No	40
2	AZ	1998	2019	22	354512	334489	No	36
3	CA	1998	2019	22	1058637	1034416	No	30
4	CT	2012	2018	7	27595	26631	Yes	29
5	DC	2011	2019	9	4972	4855	Yes	26
6	DE	2012	2018	7	2982	2830	Yes	35
7	FL	2002	2019	18	4391	3899	No	34
8	IL	2006	2019	14	7881	7469	No	31
9	KS	2008	2019	12	1016	865	No	35
10	MA	2011	2019	9	97393	93417	Yes	28
11	ME	2012	2013	2	521	521	Yes	34
12	MN	2001	2019	19	1691	1339	No	32
13	MO	2007	2019	13	5106	3181	No	34
14	NH	2012	2018	7	5601	5197	Yes	32
15	NJ	2011	2019	9	113107	107950	Yes	30
16	NM	1998	2019	22	16186	15924	No	37
17	NY	2000	2019	20	99102	93122	No	29
18	OH	2012	2019	8	2021	1987	Yes	35
19	OR	1998	2018	21	3816	3452	No	34
20	PA	2011	2014	4	3877	3311	Yes	35
21	TX	2012	2018	7	20672	19769	Yes	36
22	VA	2000	2019	20	10233	9681	No	33
23	VT	1999	2019	21	14938	13785	No	30
24	WA	2012	2019	8	6147	5941	No	33
25	WI	2002	2019	18	944	671	No	35

Figure B.1: Specification chart for linear models estimated with OLS

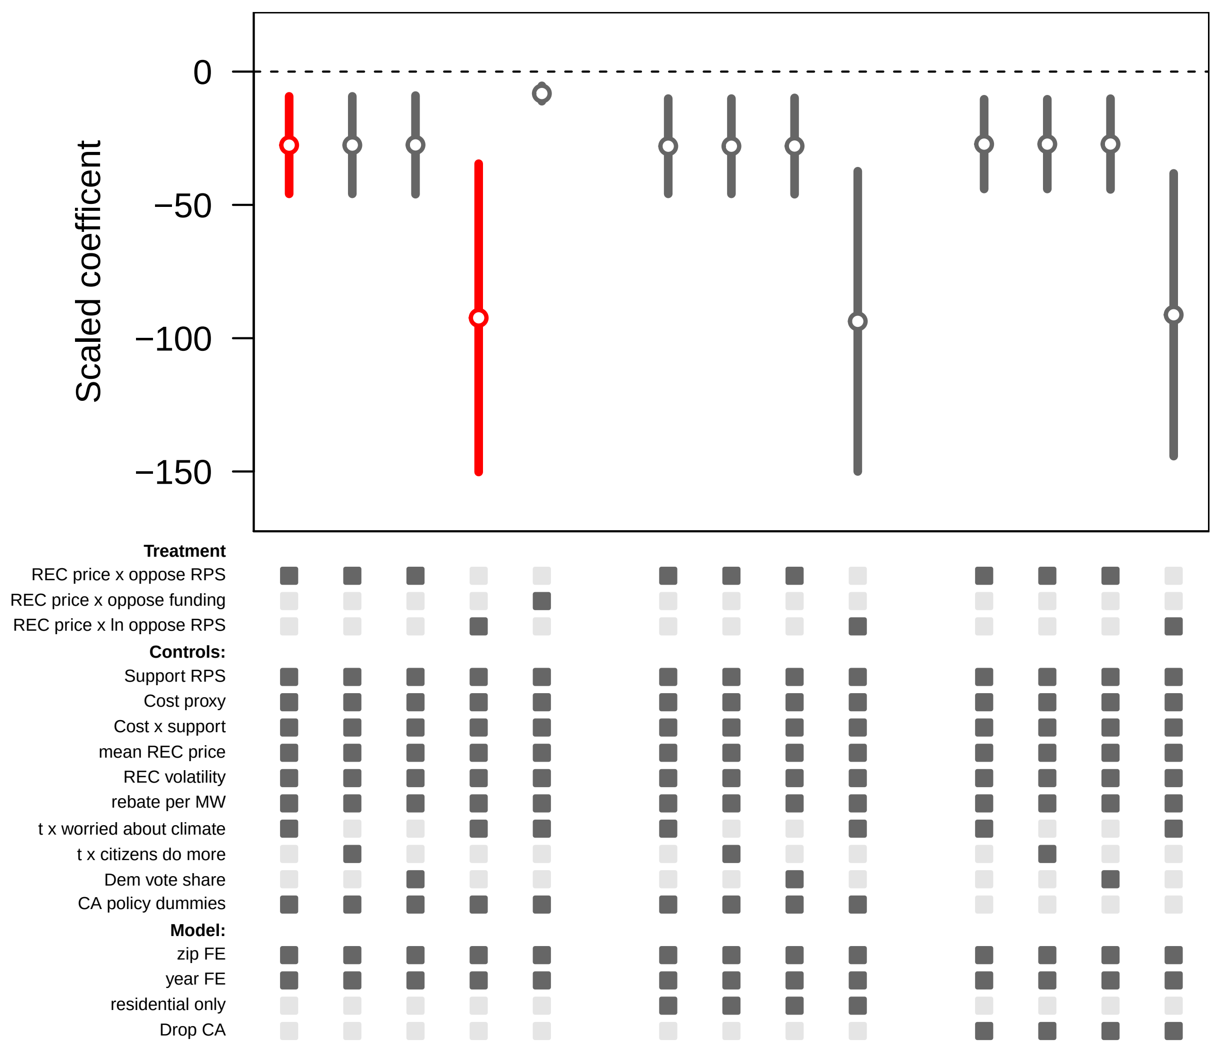

Figure B.2: Specification chart for exponential mean models estimated with Poisson QMLE

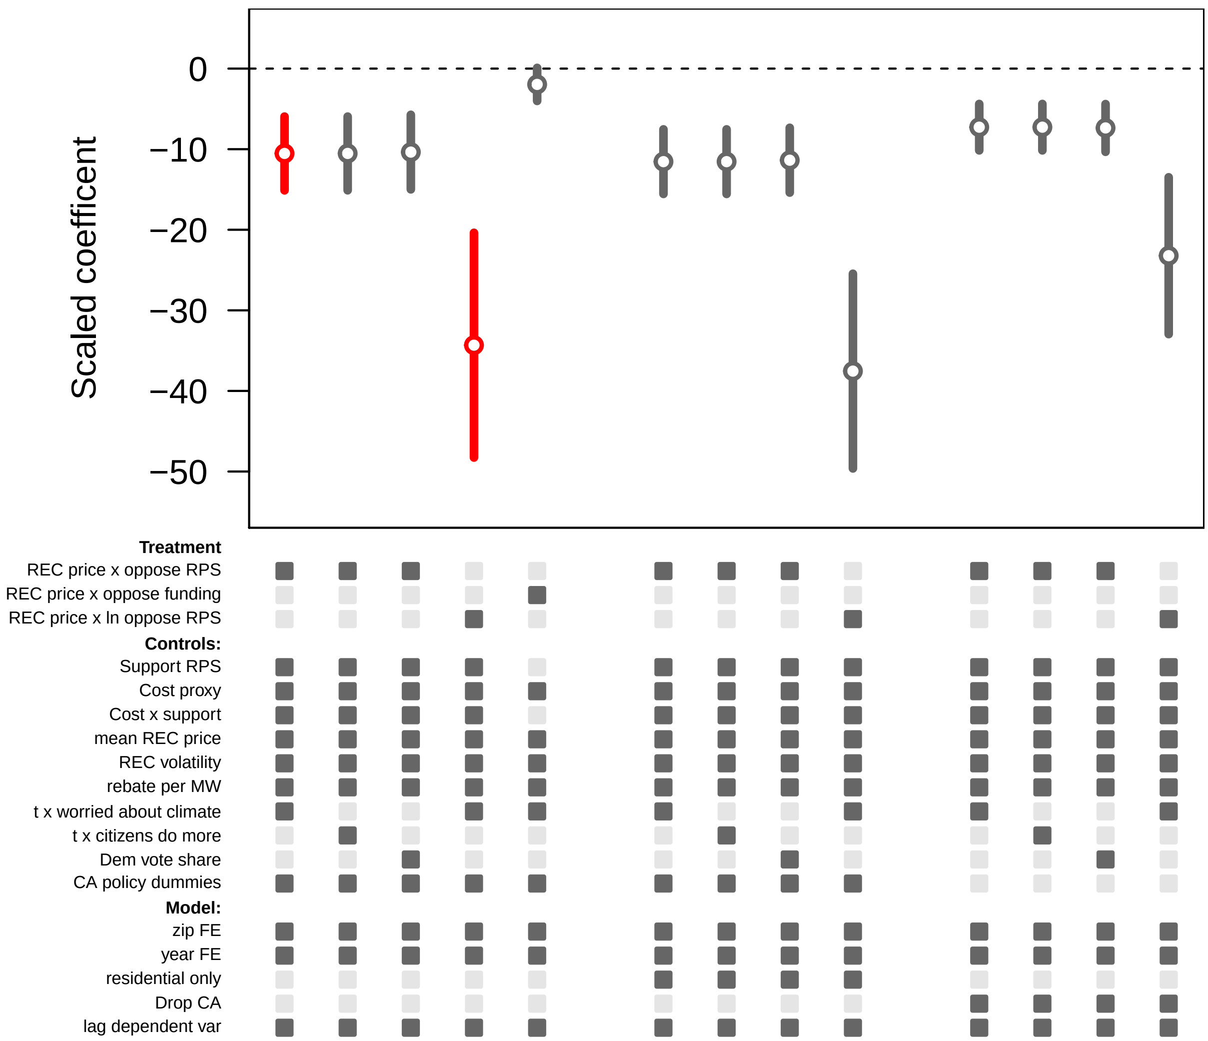

Figure B.3: Fitted values (left) and residuals (right) from linear regression (Model 1)

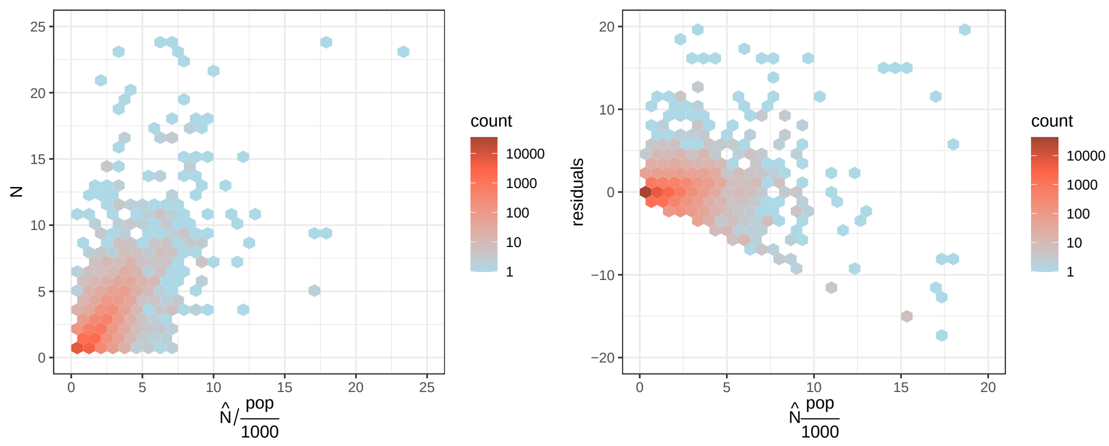

Figure B.4: Fitted values (left) and residuals (right) from Poisson regression (Model 3)

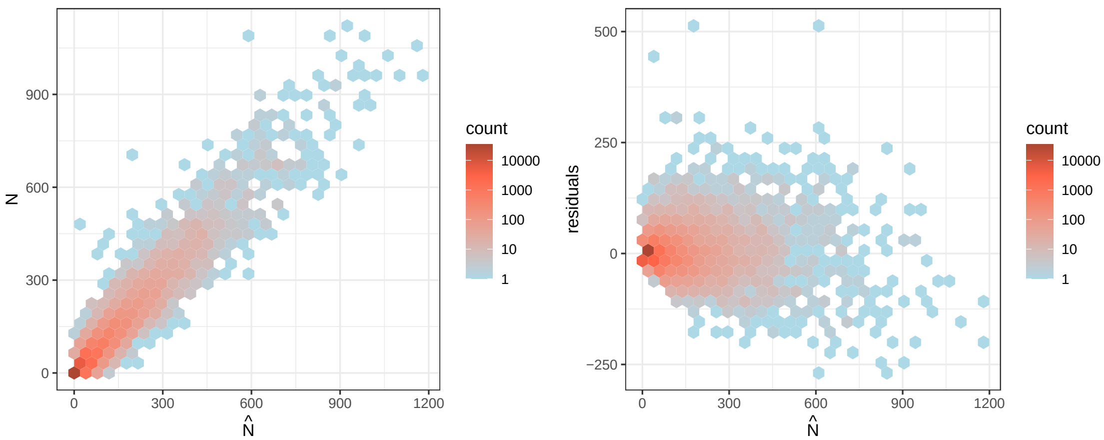

Figure B.5: Bayesian Information Criterion for linear model with an intercept and lags only

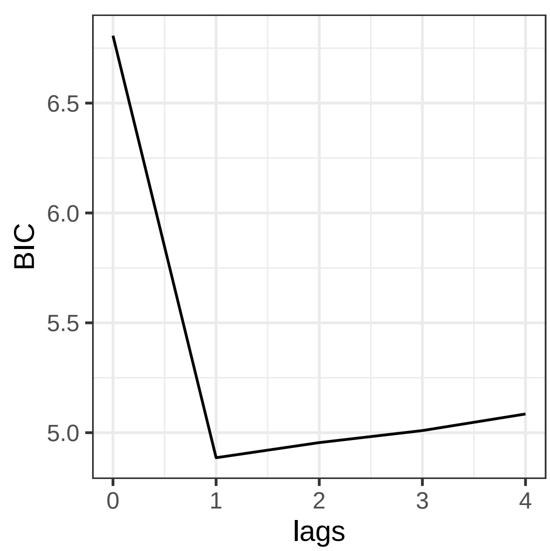

APPENDIX C

CHAPTER 3

Figure C.1: Percent change in technology parameter

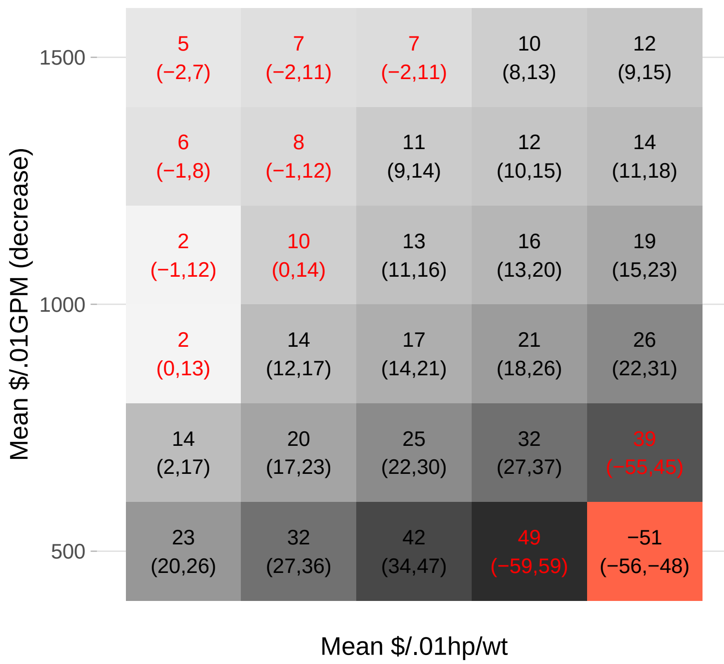

Figure C.2: Sensitivity analysis of λ

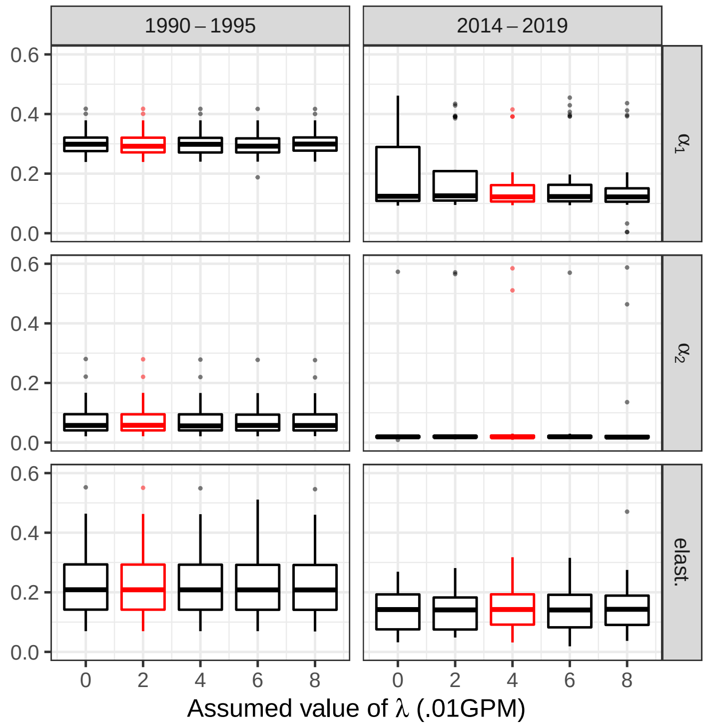

Note: boxes indicate middle quarterlies and median of estimated parameter; red indicates preferred value of λ in each time period.

Figure C.3: Percent of model-year vehicles with turbochargers (sales weighted)

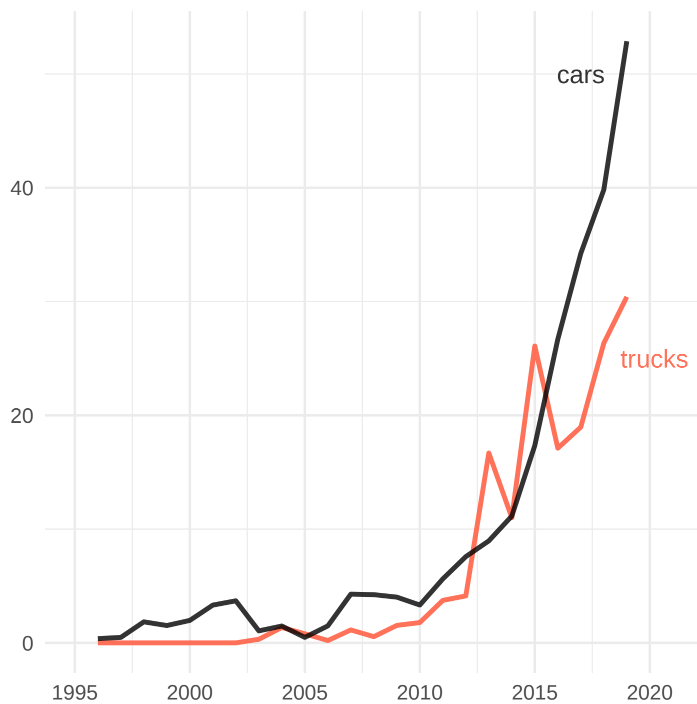

Figure C.4: Mean transmission gears (sales weighted)

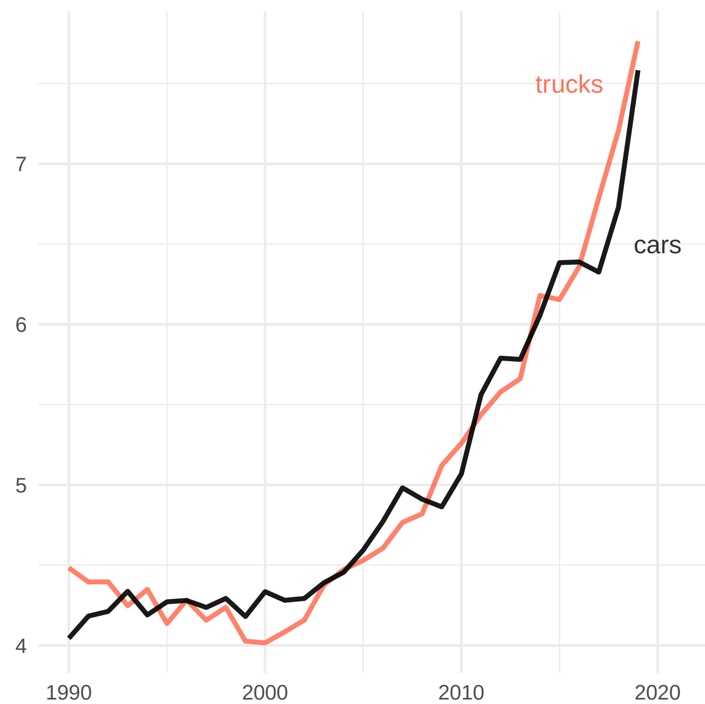

Table C.1: Transmission gears regression results

	Dependent variable:			
	Trasmission gears			
	(1)	(2)	(3)	(4)
ln mpg	1.430***	1.345***	1.327***	1.278***
	(0.038)	(0.054)	(0.038)	(0.051)
ln hp/wt	1.276***	1.375***	1.141***	1.195***
	(0.026)	(0.042)	(0.030)	(0.040)
ln volume	0.282***	−0.106	0.267***	−0.082
	(0.036)	(0.069)	(0.036)	(0.064)
time trend	0.075***	0.076***	0.080***	0.082***
	(0.001)	(0.001)	(0.001)	(0.001)
truck indicator	−0.167***	−3.054***	−0.116***	−1.440***
	(0.024)	(0.337)	(0.023)	(0.318)
ln mpg x truck		0.395***		0.070
		(0.082)		(0.078)
ln hp/wt x truck		−0.176***		−0.120**
		(0.065)		(0.061)
ln vol. x truck		0.753***		0.493***
		(0.084)		(0.079)
2015$/gal			−0.179***	−0.178***
			(0.008)	(0.008)
time trend x truck	0.030***	0.027***	0.025***	0.024***
	(0.001)	(0.002)	(0.001)	(0.002)
constant	2.556***	3.670***	2.844***	3.643***
	(0.136)	(0.218)	(0.150)	(0.206)
Observations	39,255	35,259	33,541	33,541
R^2	0.566	0.569	0.546	0.546
Adjusted R^2	0.566	0.569	0.546	0.546
Note:			*p<0.1; **p<0.05; ***p<0.01	

Mathematical appendix

Derivation of Equation 3.11

To simplify notation, let $z(\cdot) \equiv Tg^{\alpha_g} \prod_j \underline{x}_j^{\alpha_j}$.

1. $\left[\frac{\partial u_i}{\partial \underline{x}_k}\right] : \beta_{ik} = \frac{\alpha_k}{\underline{x}_k} z(\cdot) \implies \underline{x}_k^{1-\alpha_k} = \frac{\alpha_k}{\beta_{ik}} Tg^{\alpha_g} \prod_{j'} \underline{x}_{j'}^{\alpha_{j'}}$, where $j' = \{j \neq k\}$

2. $\left[\frac{\partial u_i}{\partial g_l}\right] : \phi_i p_g + \lambda = \frac{\alpha_g}{g} z(\cdot) \implies g^{1-\alpha_g} = \frac{\alpha_g}{\phi_i p_g + \lambda} T \prod_j \underline{x}_j^{\alpha_j}$.

3. Notice the following three relationships.

   As $\frac{\beta_{ij}}{\beta_{ik}} = \frac{c_{\underline{x}_j}}{c_{\underline{x}_k}} \implies \frac{\beta_{ij}}{\beta_{ik}} = \frac{\alpha_j z(\cdot)}{\alpha_k z(\cdot)} \frac{\underline{x}_k}{\underline{x}_j} \implies \underline{x}_{j'} = \frac{\beta_{ik}}{\beta_{ij}} \frac{\alpha_j}{\alpha_k} \underline{x}_k$

   and, $\frac{\phi_i p_g + \lambda}{\beta_{ij}} = \frac{c_g}{c_{\underline{x}_j}} \implies \underline{x}_j = \frac{\phi_i p_g + \lambda}{\beta_{ij}} \frac{\alpha_j}{\alpha_g} g$

   and also, $g = \frac{\beta_{ik}}{\phi_i p_g + \lambda} \frac{\alpha_g}{\alpha_k} \underline{x}_k$.

4. Define $\psi \equiv 1 - \alpha_g - \sum_j \alpha_j$.

5. Plugging in for $\underline{x}_j$ and g from step 3 gives

   $\underline{x}_k^{1-\alpha_k} = \frac{\alpha_k}{\beta_{ik}} T \left(\frac{\beta_{ik}}{\phi_i p_g + \lambda} \frac{\alpha_g}{\alpha_k} \underline{x}_k \right)^{\alpha_g} \prod_{j'} \left(\frac{\beta_{ik}}{\beta_{ij'}} \frac{\alpha_{j'}}{\alpha_k} \underline{x}_k \right)^{\alpha_{j'}}$

   $\implies \underline{x}_k^{\psi} = \frac{\alpha_k}{\beta_{ik}} T \left(\frac{\beta_{ik}}{\phi_i p_g + \lambda} \frac{\alpha_g}{\alpha_k} \right)^{\alpha_g} \prod_{j'} \left(\frac{\beta_{ik}}{\beta_{ij'}} \frac{\alpha_{j'}}{\alpha_k} \right)^{\alpha_{j'}}$

   $\implies \underline{x}_k^{*} = \left(\frac{\alpha_k}{\beta_{ik}} T \left(\frac{\beta_{ik}}{\phi_i p_g + \lambda} \frac{\alpha_g}{\alpha_k} \right)^{\alpha_g} \prod_{j'} \left(\frac{\beta_{ik}}{\beta_{ij'}} \frac{\alpha_{j'}}{\alpha_k} \right)^{\alpha_{j'}} \right)^{\frac{1}{\psi}}$.

6. Plugging in for $\underline{x}_j$ from above gives

   $g^{1-\alpha_g} = \frac{\alpha_g}{\phi_i p_g + \lambda} T \prod_j \left(\frac{\phi_i p_g + \lambda}{\beta_{ij}} \frac{\alpha_j}{\alpha_g} g \right)^{\alpha_j}$

   $\implies g^{\psi} = \frac{\alpha_g}{\phi_i p_g + \lambda} T \prod_j \left(\frac{\phi_i p_g + \lambda}{\beta_{ij}} \frac{\alpha_j}{\alpha_g} \right)^{\alpha_j}$

   $\implies g^{*} = \left(\frac{\alpha_g}{\phi_i p_g + \lambda} T \prod_j \left(\frac{\phi_i p_g + \lambda}{\beta_{ij}} \frac{\alpha_j}{\alpha_g} \right)^{\alpha_j} \right)^{\frac{1}{\psi}}$.

7. Taking logs then gives Equation 3.11.

Derivation of estimator

1. From the relationships in step 3 above, we can re-write the equilibrium attribute levels as a function of preferences and another attribute $g^* = \left(\frac{\alpha_g}{\phi_i p_g + \lambda} T \prod_j \left(\frac{\underline{x}_j}{g}\right)^{\alpha_j}\right)^{\frac{1}{\psi}}$

2. Suppose the researcher is able to observe all equilibrium outcomes of $\underline{x}_j$ $g^* = \left(\frac{\alpha_g}{\phi_i p_g + \lambda} T \prod_j \underline{x}_j^{\alpha_j}\right)^{\frac{1}{1-\alpha_g}} \implies \ln(\phi_i p_g + \lambda) = (\alpha_g - 1)\ln g + \ln \alpha_g + \ln T + \sum_j \alpha_j \ln \underline{x}_j$

3. A first-order Taylor series approximation gives $\ln(\phi_i p_g + \lambda) \approx \ln \lambda + \frac{p_g}{\lambda}\phi_i$

4. Then there are $j+1$ equations which may be estimated using an EM algorithm: $\phi_i = \frac{\lambda}{p_g}\left((\alpha_g - 1)\ln g + \ln \alpha_g + \ln \lambda T + \sum_j \alpha_j \ln \underline{x}_j\right)$ For all $k \in j$, $\ln(\beta_k) = (\alpha_k - 1)\ln \underline{x}_k + \ln \alpha_k + \ln T + \alpha_g \ln g + \sum_{j'} \alpha_{j'} \ln \underline{x}_{j'}$ where $j' \in j \neq k$

Second order conditions

Consider the simplest case, where there are only two vehicle attributes that are not additiavely separable in the cost and value functions, fuel consumption g and another attribute x. To have a local maximum, we must have

1. u_{gg} and u_{xx} are both negative;

2. $u_{gg}u_{xx} - u_{gx}^2 > 0$.

Knittel estimator Proof sketch. The solution f^*, x^* is **not** a local maximum.

1. $u_{ff} = -\alpha_f(\alpha_f - 1)c(\cdot)/f^2$, $u_{xx} = -\alpha_x(\alpha_x - 1)c(\cdot)/x^2$ are negative if $\alpha_f, \alpha_x > 1$;

2. $u_{ff}u_{xx} - u_{fx}^2 = \alpha_f(\alpha_f - 1)\alpha_x(\alpha_x - 1)c(\cdot)^2/f^2x^2 - \left(\alpha_f\alpha_x c(\cdot)/fx\right)^2$; $u_{ff}u_{xx} - u_{fx}^2 > 0$ $\implies \alpha_f(\alpha_f - 1)\alpha_x(\alpha_x - 1)c(\cdot)^2/f^2x^2 > \left(\alpha_f\alpha_x c(\cdot)/fx\right)^2 \implies (\alpha_f - 1)(\alpha_x - 1) > \alpha_f\alpha_x$ $\implies \alpha_f\alpha_x - \alpha_f - \alpha_x + 1 > \alpha_f\alpha_x \implies 1 > \alpha_f + \alpha_x$, which contradicts $\alpha_f, \alpha_x > 1$.

Proposed estimator Proof sketch. The solution $g^*, \underline{x}^*$ is a local maximum under the constraint $\alpha_g + \alpha_x < 1$ and $\alpha_g, \alpha_x \in (0, 1)$.

1. $u_{gg} = \alpha_g(\alpha_g - 1)c(\cdot)/g^2$, $u_{xx} = \alpha_x(\alpha_x - 1)c(\cdot)/x^2$ are negative if $\alpha_g, \alpha_x < 1$;

2. The second step is the same as above, except in that it does not contradict the implied range of values for the alphas. $u_{gg}u_{xx} - u_{gx}^2 = \alpha_g(\alpha_g - 1)\alpha_x(\alpha_x - 1)c(\cdot)^2/g^2x^2 - \left(\alpha_g\alpha_x c(\cdot)/gx\right)^2$; $u_{gg}u_{xx} - u_{gx}^2 > 0 \implies \alpha_g(\alpha_g - 1)\alpha_x(\alpha_x - 1)c(\cdot)^2/g^2x^2 > \left(\alpha_g\alpha_x c(\cdot)/gx\right)^2$ $\implies (\alpha_g - 1)(\alpha_x - 1) > \alpha_g\alpha_x \implies \alpha_g\alpha_x - \alpha_g - \alpha_x + 1 > \alpha_g\alpha_x \implies 1 > \alpha_g + \alpha_x$.

Derivation of MRAS in goods space

$$\begin{aligned}
\frac{\partial C}{\partial x_j} / \frac{\partial C}{\partial g_j^{-1}} &= -T\frac{\alpha_j C}{\underline{x}_j}(-1) / -T\frac{\alpha_g C}{g}(-(g^{-1})^{-2}) \\
&= \frac{\alpha_j}{\alpha_g}\frac{g^{-1}}{\underline{x}_j}.
\end{aligned}$$

Where the first line makes use of the chain rule and the second line simplifies terms.

APPENDIX D

CHAPTER 4

To receive an indemnity with revenue insurance at the middle terminal value, implies that the distribution of outcome conditional on harvesting is negatively skewed.

To show this, we first prove for v_2 to be less than the initial-period mean, $\mathbb{E}_{(1,1)}[v]$. Then note that for over-abandonment to occur Equation **??** implies that $\mathbb{E}_{(1,1)}[v] \geq \frac{1}{\lambda} v_2 \geq v_2$, where the second inequality comes from the fact that $\lambda \in [0,1]$.

Proposition: $v_2 < \mu \implies \kappa_3 < 0$ where $\mu \equiv \mathbb{E}_{(1,1)}[v]$ and κ_3 is the central third moment of v at node (1,1).

Proof: Define $\phi \equiv v_1 - \mu$. As v_1 is the max of the support of v, $v_1 \geq \mu$ and therefore $\phi \geq 0$. Further $\phi = 0$ only if $pr_1 = 1$, which is an uninteresting case. As $v_1 - d = v_2$ and $v_2 < \mu$, $\phi < d$. 4) So, $v_2 - \mu = v_2 + d - \mu - d = \phi - d$ and $v_3 - \mu = \phi - 2d$.

1) $\kappa_3 \equiv \mathbb{E}\left[(v-\mu)^3\right] = pr_1 (v_1 - \mu)^3 + pr_2 (v_2 - \mu)^3 + pr_3 (v_3 - \mu)^3$. 2) As $d \geq \phi$, $d^3 \geq \phi d^2 \geq \phi^2 d \geq \phi^3$. 3) Re-arranged,[1] This becomes

$$\kappa_3 =$$

$$\phi^3 - d^3 \qquad (\leq 0)$$

$$+ (pr_1 - pr_3 - 1)\left(\phi^2 d - \phi d^2 + d^3\right) \qquad (< 0)$$

$$+ pr_3 \left(2\phi d^2 - 6d^3\right) \qquad (\leq 0)$$

<0 ■

[1] $\kappa_3 = pr_1\phi^3 + pr_2(\phi - d)^3 + pr_3(\phi - 2d)^3 = pr_1\phi^3 + pr_2(\phi^3 - \phi^2 d + \phi d^2 - d^3) + pr_3(\phi^3 - \phi^2 d + \phi d^2 - d^3) = \phi^3 + pr_2(-\phi^2 d + \phi d^2 - d^3) + pr_3(-2\phi^2 d + 4\phi d^2 - 8d^3) = \phi^3 + (1 - pr_1 - pr_3)(-\phi^2 d + \phi d^2 - d^3) + pr_3(-2\phi^2 d + 4\phi d^2 - 8d^3)$

BIBLIOGRAPHY

BIBLIOGRAPHY

[1] Acemoglu, D. (2002). Directed technical change. *The Review of Economic Studies*, 69(4):781–809.

[2] Achdou, Y., Han, J., Lasry, J.-M., Lions, P.-L., and Moll, B. (2017). Income and wealth distribution in macroeconomics: A continuous-time approach. Technical report, National Bureau of Economic Research.

[3] Aigner, D., Lovell, C. K., and Schmidt, P. (1977). Formulation and estimation of stochastic frontier production function models. *Journal of econometrics*, 6(1):21–37.

[4] Anderson, S. T. and Sallee, J. M. (2011). Using loopholes to reveal the marginal cost of regulation: The case of fuel-economy standards. 101(4):1375–1409.

[5] Babcock, B. A. and Hennessy, D. A. (1996). Input Demand under Yield and Revenue Insurance. 78(2):416.

[6] Ballew, M. T., Leiserowitz, A., Roser-Renouf, C., Rosenthal, S., Kotcher, J. E., Marlon, J. R., Lyon, E., Goldberg, M. H., and Maibach, E. W. (2019). Climate change in the american mind: Data, tools, and trends. *Environment: Science and Policy for Sustainable Developmen*, 61(3):4–18.

[7] Barbose, G. (2019). U.s. renewables portfolio standards 2019 annual status update. Technical report, Lawrence Berkeley National Lab.(LBNL), Berkeley, CA (United States).

[8] Barbose, G. (2021). U.s. renewables portfolio standards 2021 annual status update. Technical report, Lawrence Berkeley National Lab.(LBNL), Berkeley, CA (United States).

[9] Barbose, G., Darghouth, N., Elmallah, S., Forrester, S., Kristina SH, K., Millstein, D., Rand, J., Cotton, W., Sherwood, S., and O'Shaughnessy, E. (2019). Tracking the sun: Pricing and design trends for distributed photovoltaic systems in the united states-2019 edition.

[10] Binswanger, H. P. et al. (1981). Attitudes toward risk: Theoretical implications of an experiment in rural india. *Economic Journal*, 91(364):867–890.

[11] Bolinger, M., Seel, J., Robson, D., and Warner, C. (2020). Utility-scale solar data update: 2020 edition. Technical report, Lawrence Berkeley National Lab.(LBNL), Berkeley, CA (United States).

[12] Borges, R. B. and Thurman, W. N. (1994). Marketing quotas and random yields: marginal effects of inframarginal subsidies on peanut supply. *American Journal of Agri-*

cultural Economics, 76(4):809–817.

[13] Braun, M., Fader, P. S., Bradlow, E. T., and Kunreuther, H. (2006). Modeling the âPseudodeductibleâ in Insurance Claims Decisions. 52(8):1258–1272.

[14] Center, S. (2020). Climate deregulation tracker.

[15] Chamberlin, E. H. (1949). *Theory of monopolistic competition: A re-orientation of the theory of value*. Oxford University Press, London.

[16] Coble, K. H. and Barnett, B. J. (2013). Why do we subsidize crop insurance? *American Journal of Agricultural Economics*, 95(2):498–504.

[17] Cullenward, D. and Victor, D. G. (2020). *Making Climate Policy Work*. John Wiley & Sons.

[18] De Groote, O. and Verboven, F. (2016). Subsidies and myopia in technology adoption: Evidence from solar photovoltaic systems.

[Dekraker et al.] Dekraker, P., Barba, D., Moskalik, A., and Butters, K. Constructing engine maps for full vehicle simulation modeling.

[20] Dempster, A. P., Laird, N. M., and Rubin, D. B. (1977). Maximum likelihood from incomplete data via the em algorithm. *Journal of the Royal Statistical Society: Series B (Methodological)*, 39(1):1–22.

[21] Dempster, A. P., Rubin, D. B., and Tsutakawa, R. K. (1981). Estimation in covariance components models. *Journal of the American Statistical Association*, 76(374):341–353.

[22] Di Crescenzo, A. and Meoli, A. (2016). On a fractional alternating poisson process. *AIMS Mathematics*, 1(3):212–224.

[23] Dixit, A. K., Dixit, R. K., Pindyck, R. S., and Pindyck, R. (1994). *Investment under uncertainty*. Princeton university press.

[24] Dohmen, T. J., Falk, A., Huffman, D., Sunde, U., Schupp, J., and Wagner, G. G. (2005). Individual risk attitudes: New evidence from a large, representative, experimentally-validated survey.

[25] Doidge, M. K. (2019). *Farmers' Behaviour in the Face of Uncertainty: A Behavioural Economics Approach to Farmers' Land and Insurance Decisions*. Michigan State University.

[26] Du, X., Feng, H., and Hennessy, D. A. (2016). Rationality of Choices in Subsidized Crop Insurance Markets. page aaw035.

[27] Fowlie, M., Greenstone, M., and Wolfram, C. (2018). Do energy efficiency investments deliver? evidence from the weatherization assistance program. *The Quarterly Journal of Economics*, 133(3):1597–1644.

[28] Friend, I. and Blume, M. E. (1975). The demand for risky assets. *The American Economic Review*, pages 900–922.

[29] Funds, C. I. (2011). *Guidance Paper on Subsidies for Renewable Energy Technologies Under SREP*.

[30] Glauber, J. W. (2004). Crop insurance reconsidered. *American Journal of Agricultural Economics*, 86(5):1179–1195.

[31] Gomez-Limon, J. A., Riesgo, L., and Arriaza Balmón, M. (2002). Agricultural risk aversion revisited: A multicriteria decision-making approach. Technical report.

[32] Greene, D., Hossain, A., Hofmann, J., Helfand, G., and Beach, R. (2018). Consumer willingness to pay for vehicle attributes: What do we know? 118:258–279.

[33] Hochman, G. and Zilberman, D. (2021). Optimal environmental taxation in response to an environmentally-unfriendly political challenger. *Journal of Environmental Economics and Management*, 106:102407.

[34] Howe, P. D., Mildenberger, M., Marlon, J. R., and Leiserowitz, A. (2015). Geographic variation in opinions on climate change at state and local scales in the usa. *Nature climate change*, 5(6):596–603.

[35] Kachelmeier, S. J. and Shehata, M. (1992). Examining risk preferences under high monetary incentives: Experimental evidence from the people's republic of china. *The American Economic Review*, pages 1120–1141.

[36] Klier, T. and Linn, J. (2012). New-vehicle characteristics and the cost of the corporate average fuel economy standard. 43(1):186–213.

[37] Knittel, C. R. (2011). Automobiles on steroids: Product attribute trade-offs and technological progress in the automobile sector. 101(7):3368–3399.

[38] Lade, G. E., Lin Lawell, C.-Y. C., and Smith, A. (2018). Policy shocks and market-based regulations: Evidence from the renewable fuel standard. *American Journal of Agricultural Economics*, 100(3):707–731.

[39] Lancaster, T. (1990). *The Econometric Analysis of Transition Data*. Number 17. Cambridge university press.

[Leard and Linn] Leard, B. and Linn, J. Have US fuel economy and greenhouse gas emissions standards improved social welfare? page 51.

[41] MacKenzie, D. and Heywood, J. (2012). Acceleration performance trends and evolving relationship between power, weight, and acceleration in u.s. light-duty vehicles: Linear regression analysis. 2287(1):122–131.

[42] Matisoff, D. C. and Johnson, E. P. (2017). The comparative effectiveness of residential solar incentives. 108:44–54.

[43] Moskalik, A., Bolon, K., Newman, K., and Cherry, J. (2018). Representing GHG reduction technologies in the future fleet with full vehicle simulation.

[44] Mulungu, K. and Tembo, G. (2015). Effects of Weather Variability on Crop Abandonment. 7(3):2858–2870.

[45] Newell, R. G., Jaffe, A. B., and Stavins, R. N. (1999). The induced innovation hypothesis and energy-saving technological change. *The Quarterly Journal of Economics*, 114(3):941–975.

[46] Nickell, S. (1981). Biases in dynamic models with fixed effects. *Econometrica: Journal of the econometric society*, pages 1417–1426.

[47] on Climate Change Communication, Y. P. and for Climate Change Communication, G. M. U. C. (2020). Climate change in the american mind: National survey data on public opinion (2008-2018) [data file and codebook].

[48] Pigou, A. C. (1920). *The economics of welfare. 2013 reprint.* Palgrave Macmillan.

[49] Roberts, M. J., Key, N., and O'Donoghue, E. (2006). Estimating the extent of moral hazard in crop insurance using administrative data. *Review of Agricultural Economics*, 28(3):381–390.

[50] Scott Hacker, R. and Hatemi-J, A. (2008). Optimal lag-length choice in stable and unstable var models under situations of homoscedasticity and arch. *Journal of Applied Statistics*, 35(6):601–615.

[51] Seo, S., Mitchell, P. D., and Leatham, D. J. (2005). Effects of federal risk management programs on optimal acreage allocation and nitrogen use in a texas cotton–sorghum system. *Journal of Agricultural and Applied Economics*, 37(3):685–699.

[52] Shahed, S. and Bauer, K.-H. (2009). Parametric studies of the impact of turbocharging on gasoline engine downsizing. *SAE International Journal of Engines*, 2(1):1347–1358.

[53] Stokes, L. C. (2020). *Short circuiting policy: Interest groups and the battle over clean energy and climate policy in the American States.* Oxford University Press, USA.

[54] Szpiro, G. G. (1986). Measuring risk aversion: an alternative approach. *The review of economics and statistics*, pages 156–159.

[55] USDA (2019). Risk management agency, summary of business.

[56] USEPA (2009). Epa optimization model for reducing emissions of greenhouse gases from automobiles (omega): Model documentation 1.0. Environmental Protection Agency, 420-B-09-035.

[57] USEPA (2019). The 2019 automotive trends report epa-420-r-20-006.

[58] Weitzman, M. L. (1974). Prices vs. quantities. *The review of economic studies*, 41(4):477–491.

[59] Whitefoot, K. S., Fowlie, M. L., and Skerlos, S. J. (2017). Compliance by design: Influence of acceleration trade-offs on CO_2 emissions and costs of fuel economy and greenhouse gas regulations. 51(18):10307–10315.

[60] Wooldridge, J. M. (1997). Multiplicative panel data models without the strict exogeneity assumption. *Econometric Theory*, pages 667–678.

[61] Zacks, S. (2012). Distribution of the total time in a mode of an alternating renewal process with applications. *Sequential Analysis*, 31(3):397–408.