# Attribute Production and Technical Change in Automobiles*

Asa Watten† Soren Anderson‡ Gloria Helfand§

January 18, 2023

**Abstract**

We develop a theoretic foundation for analyzing technical change and attribute production for light-duty vehicles, which can be applied to other multi-attribute goods. Applying our theory to survey data of car purchases from 1999 to 2017, we estimate a 10% annual decline in drive-train related production costs. For comparison, EPA recently proposed year-over-year increases in light-duty vehicle greenhouse gas standards of 1.5 percent. We estimate that, holding cost and other attributes fixed, horsepower must be reduced by much more—1.6 times as much—in the later period to increase fuel economy. Fuel economy can be increased either by adding technology or by changing the mix of car attributes. These estimates imply that fuel-economy standards today provide a stronger incentive to add fuel-saving technology and weaker incentive to reduce acceleration and size compared to prior decades. We highlight the role of consumer preferences in driving both observed correlations between car attributes at a point in time—and how consumers readjust their optimal mix of attributes as costs decline. By formalizing a theory of attribute production we provide tools for interpreting prior empirical estimates of attribute trade-offs and technical change in the literature.

JEL classification numbers: XXX

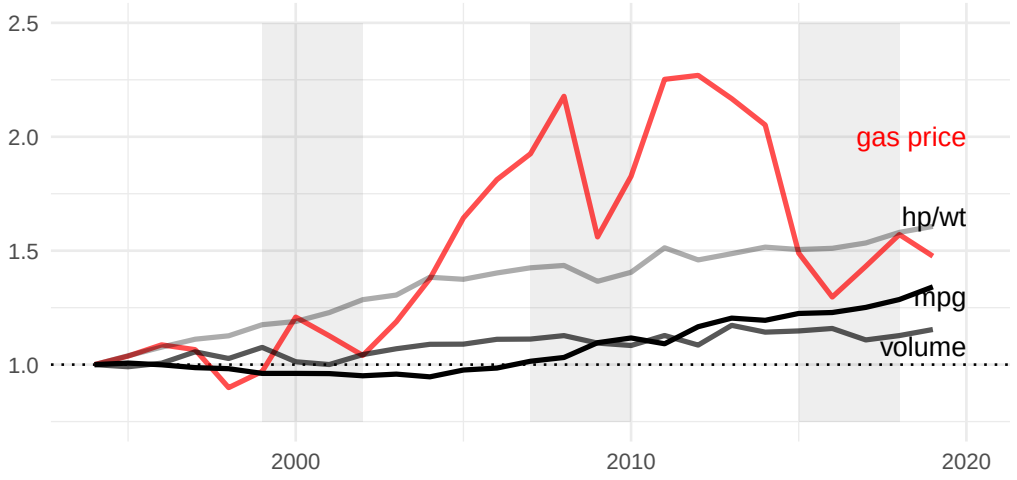
Keywords: XXX

*Acknowledgements: We thank Kevin Bolon and Andrew Moskalik for engineering insights.
†Yale University, 195 Prospect Street, New Haven, CT 06511, asa.watten@yale.edu
‡Michigan State University, 486 West Circle Dr, East Lansing, MI 48824, sta@msu.edu
§University of Michigan, 440 Church Street Ann Arbor, MI 48109, ghelfand@umich.edu

# 1 Introduction

What can be done to make cars more energy-efficient? This question bedevils U.S. policymakers, who seek to reduce greenhouse gas emissions from cars without resorting to politically toxic fuel taxes. Rising oil prices and mandated increases in fuel economy in the late 1970s and early 1980s coincided with large reductions in acceleration, and later size (EPA 2019). This association led to a belief that fuel-economy regulation would necessarily force consumers into smaller, less powerful cars that are less enjoyable and potentially less safe to drive. Falling oil prices in the 1990s then coincided with increases in size and acceleration, and decreases in fuel economy, reinforcing this belief. Yet car performance continued to improve in the 2000s, even under rising oil prices, while size plateaued and fuel economy actually increased—a pattern that continued into the 2010s. See figure 1, which shows trends in new car attributes, and which proxies for acceleration using the horsepower-to-weight ratio.

**Figure 1:** *Attributes for new cars 1994-2019*




Note: Ratio changes in the mean attributes (sales weighted) and gasoline prices compared to 1994. The grey areas show our study years using the National Household Transportation Survey waves 2001, 2009, and 2017. The horsepower-to-weight ratio (HP/weight) is a close proxy for acceleration capacity. Gas prices are core inflation adjusted.

What explains these trends? Cars have simultaneously gotten bigger, faster, more efficient, and arguably more affordable—so clearly there have been large improvements in technology. What has been the nature of this technical change? Performance fell in the 1980s but not the 2000s when oil prices increased. Does this imply an altered relationship between fuel economy and performance? Car technology does not exist in a vacuum, of course. What has been the role of shifting consumer preferences for size and acceleration? How about the role of gas prices and of government policy? Finally, looking to the future: What does a

better understanding of these trends imply for fuel-economy standards and for the looming shift to electric cars going forward? We address these questions using a combination of engineering simulations, economic theory, and empirical tests.

We begin by tracing out the physical relationship between horsepower and fuel economy for six distinct drive-trains from the 1980s to the present, each representing a different vintage of car technology, using an EPA engineering simulation tool. We trace out this relationship by varying engine displacement (total volume of the combustion chambers), which is a key design choice affecting both acceleration and fuel economy. We then develop a simple but powerful theory of car attribute production that accounts for technology, preferences, and policy. On the supply side, we model a competitive car industry facing a Cobb-Douglas-like variable cost function that depends on car size, acceleration, and fuel economy. We show precisely how these attributes must enter the cost function to ensure engineering relationships that align both with our simulations and with economic intuition. On the demand side, we assume heterogeneous marginal willingness-to-pay for car attributes across consumer types, leading to an equilibrium distribution of differentiated car models. We solve for a consumer's optimal choice of fuel economy conditional on other car attributes (as in Knittel 2011), as well as optimal attributes conditional only on model primitives. We then probe the model to derive comparative statics for the equilibrium effects of higher gasoline prices, tighter fuel-economy standards, and various forms of technical change. Finally, we test the model's implications using household-level microdata, exploring trends in the distribution of car attributes over time, while controlling for changes in consumer preferences due to gasoline prices and shifting demographics. In addition, we examine the diffusion of specific engine technologies (e.g. turbo-charging) across the car fleet, testing whether they are adopted uniformly or are targeted at bigger, more powerful cars.

Our data includes three waves of the National Household Transportation Survey (2001, 2009, and 2017) micro consumer data, Environmental Protection Agency attribute data, and detailed vehicle size data from the Canadian government.

Our analysis generates five sets of findings. First, our engineering simulations show that reducing horsepower (via a smaller engine) yields large fuel-economy gains in the 1980s but not so much for modern drive-trains. In particular, we find that a 10% decrease in horsepower leads to a 20% increase in fuel economy in the 1980s but only 12% increase in recent years. These results help explain why car performance fell while fuel economy rose in the 1980s or why performance improved in the 2000s while fuel economy fell slightly despite rising gas prices and why both improved in the 2010s: the marginal cost of saving fuel by choosing a smaller engine has increased. Note, this result does not imply that technology has regressed. To the contrary, our simulations also show that engines of a given size have gotten both more efficient and more powerful over time, which is consistent with across-the-board technical change.

Second, our theory implies that higher gasoline prices lead to proportional reductions in size and acceleration, for a given consumer type (car model). But performance is less affected when fuel economy can be improved easily, i.e. without driving up costs or reducing acceleration. Fuel-economy standards lead to qualitatively similar results. Again, these results highlight the possible role of technical change in explaining the differential outcomes of the 1980s vs. 2000s. Third, our theory implies that unbiased technical change in the abstract leads to equivalent percent gains in size, acceleration, and fuel economy that are identical across consumer types (car models). Yet technical change in practice often manifests as a cost reduction for specific technologies, e.g. turbochargers. We show that this form of technical change is inherently attribute-biased: costly fuel-saving technologies are first installed on bigger and more powerful cars, or by consumers traveling many miles, for which a percent reduction in fuel consumption is most valuable. Such technologies are only added to smaller cars later, as installation costs decline.

Fourth, our empirical analysis shows that trends in car attributes are largely explained by shifts in consumer preferences, including those driven by gasoline prices, while fuel-economy standards explain relatively little. These results highlight the danger of estimating technical change without controlling for costs or preferences. Meanwhile, controlling for preferences, our estimates reveal moderate but persistent trends in car size, acceleration, and fuel economy which are consistent with technical change. Overall, we estimate an greater than 10% rate of technical progress for drive-trains during our sample period. However, these gains are not distributed equally. Controlling for preferences, the gains in size and performance have been greater. Viewed through the lens of our theoretical model, this result implies that biased technical change has flattened the trade-off between fuel economy and performance. Consistent with our engineering simulations, we estimate that a 10% decrees in acceleration lead to a 25% increase in fuel economy in 2001, but only a 16% increase in fuel economy in 2017. Likewise, the effect of a 10% reduction in car size leads to a 25% gain in fuel economy in the early years and a 20% gain in later years.

Finally, our empirical analysis confirms that fuel-saving technologies are installed on bigger and more powerful cars first, and only diffuse to smaller cars gradually over time. This pattern holds for all of the technologies we consider, including turbo-chargers, continuously variable transmissions, direct fuel injection, and hybrid electric drive-trains. These results imply that we cannot regress fuel economy on size and acceleration for car model offerings and hope to recover the slopes of iso-cost curves, as in recent literature (Knittel 2011). Bigger, more powerful cars inevitably feature more advanced technology, implying more costly drive-trains. Controlling for the presence of these technologies in an attempt to limit the bias is also no solution, for such controls absorb the very mechanism of technical change we seek to estimate.

A gigantic literature in economics models and estimates technical change in the aggregate economy or for broad industries. A smaller literature, almost entirely empirical, studies technical change for differentiated products with a focus on energy efficiency. The seminal paper in this literature is Newell, Jaffe, and Stavins

(1999), who estimate shifts over time in the relationship between product cost and energy efficiency for air conditioners and water heaters, and relate this technical change to energy prices and policy. We contribute to this literature by providing theoretical micro-foundations, a new set of technical terms, and new empirical methods to better understand technical change for product attributes, including energy efficiency. Such theory is key to understanding the effects of regulation on equilibrium attributes, welfare, and incidence. Using this model, we show that failing to control for costs will lead to biased inferences of technical change. However, we show that controlling for consumer preferences can help compensate for a lack of accurate cost data. Our application to cars reveals attribute-biased technical change, which implies that efficiency gains are now more easily obtained via fuel-saving technology, rather than via reductions in performance as in the 1980s. Additionally, we compile detailed model-level attribute data for the top 50 car brands in the United States (1994-2020) from multiple government sources. These data, which we will make freely available for public use, will be of substantial interest to researchers working on all range of car-related topics.

Our paper is closely related to Knittel (2011), who also uses market data to understand trade-offs among car attributes. Dozens of papers cite Knittel (2011) as the basis for their understanding of technical change and attribute trade-offs in the car industry, whether studying cars specifically (MacKenzie and Heywood 2012; Klier and Linn 2015; Whitefoot, Fowlie, and Skerlos 2017) or technical change more broadly (Newell, Jaffe, and Stavins 1999). We build on and expand the micro-foundations for the empirical specification in Knittel (2011), which leads to a different reinterpretation of his results. In particular, our model clarifies that different bundles of car attributes emerge from heterogeneous consumer preferences for size, power, and fuel economy. Thus, costly engine technologies are not applied at random, since their application is driven by the same consumer preferences that determine car attributes in equilibrium. Indeed, our theory predicts and our empirical results confirm that all of the engine technologies considered in Knittel (2011) are first applied to bigger, more powerful cars. These results imply that we cannot estimate iso-cost curves and technical change using market data as in Knittel (2011) without either controlling for consumer preferences or observing engine costs directly. Knittel (2011) proxies for costs using suggested retail prices, finding that the results are unchanged, but of course prices for cars—unlike water heaters, air conditioners, and other utilitarian appliances—reflect many unrelated add-ons (e.g., luxury trims). Lacking accurate cost data, we instead control for consumer preferences using micro-data over a more recent time period and find substantially different results: the average rate of technical change is much larger, but trends in fuel economy holding other observed attributes fixed are lower implying that consumers are either taking some of the cost savings or applying them to other aspects of cars.[1]

Our paper also contributes to a large literature studying how fuel economy, size, weight, horsepower, and

[1]Knittel (2011) also controls directly for the presence of turbochargers and superchargers. He omits these controls in some models, to allow his estimates of technical progress to reflect the increased penetration of these technologies, but his results again change little, perhaps because these particular technologies never rise above 10% penetration.

other car attributes respond to gasoline prices and fuel-economy standards. See Anderson and Sallee (2016) for a conceptual model and review. Many of these papers use reduced-form methods, implicitly capturing attribute trade-offs but failing to disentangle demand and supply. Some papers estimate discrete-choice structural models of demand, which they then complement with estimates of supply. But almost all such papers take the set of car models and their attributes, often including cost, as given. Thus, they do not consider cost and attribute trade-offs on the supply side. Exceptions include XXX, YYY, and ZZZ, who all make car attributes endogenous, at the expense of much complexity. Yet even these papers fail to consider long-run technical change. Our paper lies somewhere in-between these reduced-form and structural papers. We provide a structure that allows us to disentangle supply and demand. But we abstract away from naming specific car models and their makers, which allows us to consider both short-run attribute trade-offs and long-run technical change. Our theory follows Ito and Sallee (2018), who study weight-based fuel-economy standards in the Japanese car market. Like us, they assume a perfectly competitive car industry with zero entry cost to develop a new car model, such that every consumer type gets a car model optimized to its own preferences. They then quantify short-run trade-offs between fuel economy and weight using a local quadratic approximation for welfare, which reflects both costs and preferences. We provide a global functional form for costs and separately assume linear preferences (as an approximation), which together we use to characterize the full distribution of car models. We show that this simple model provides a rich set of predictions that align closely with empirical evidence.

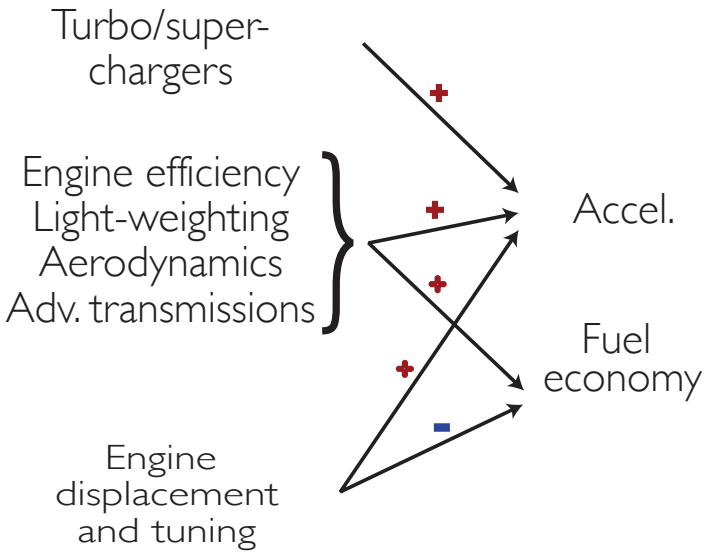
The rest of this paper proceeds as follows. Section 2 discussing engineering and design trade-offs among car attributes, and presents simulation results that illustrate these trade-offs. Section 3 develops a model of car attribute choices, showing how consumer preferences, gasoline prices, production costs, and technical change interact to drive equilibrium size, acceleration, and fuel economy. Section 4 leverages this this theory to estimate production cost parameters and rates of technical change using household-level micro-data on car attribute choices. Section 5, to be added in a draft, discusses implications for electric cars and fuel-economy standards going forward. Section 6 concludes.

## 2 Engineering trade-offs in car attributes

Carmakers modify major car attributes—size, power, and fuel economy—by adding or subtracting discrete technologies, marginally improving drive-train components, choosing different materials, increasing or decreasing engine size, and tuning the engine to achieve different goals. Figure 2 illustrates how design choices, shown in the center column, determine acceleration and fuel economy. The figure illustrates three categories of design choices: (1) those that affect only acceleration (turbo-chargers and super-chargers); (2) those that increase both fuel economy and acceleration (engine efficiency, light-weighting, aerodynamics, advanced transmissions); and (3) those that improve acceleration at the expense of fuel economy (engine displacement

and tuning). We created this figure in discussion with automotive engineers at the U.S. Environmental Protection Agency who have extensive experience testing and modeling design choices and technology options on vehicles.

**Figure 2:** *The determinants of acceleration and fuel economy*




Note: The figure shows a simplified illustration of the determinants of acceleration and fuel economy during the design process. Signs "+" and "−" indicate the sign of the effect of increasing or adding the given vertex.

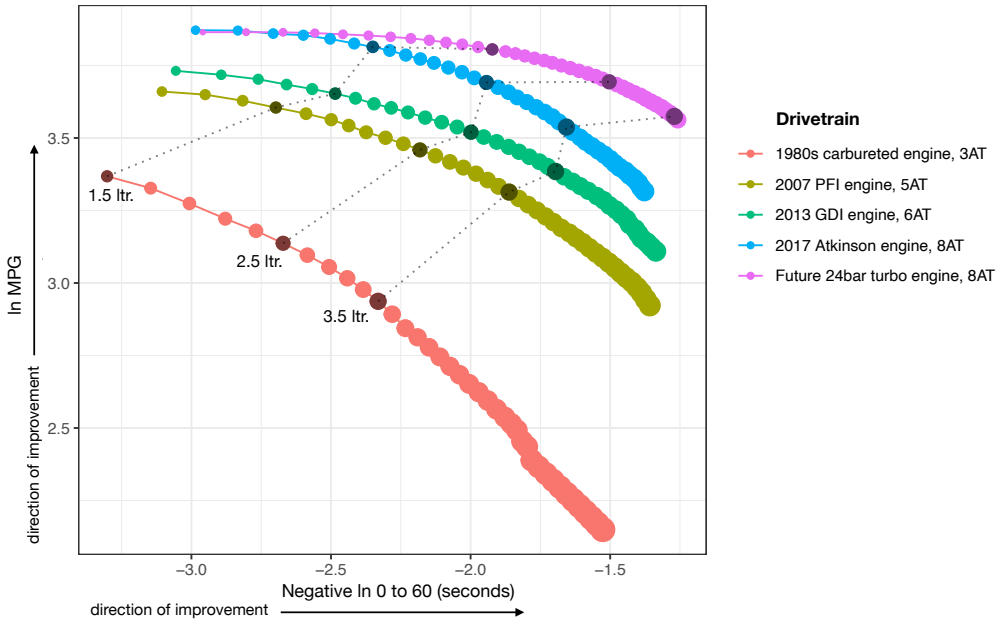
Not shown on the figure are energy-consuming technologies which may reduce fuel economy, but which have little or no effect on power, such as stereo systems, air conditioning, and computation for self-driving technologies. While this figure is not comprehensive, engine displacement and tuning are the only design choices we know of that increase acceleration at the expense of fuel economy.[2] The figure assumes that the size of the vehicle is held fixed. Of course, increasing size and therefore weight lowers both fuel economy and acceleration, assuming no other changes to the drive-train.

How much can fuel economy be improved by decreasing acceleration while holding costs fixed? Below, we define this trade-off as the marginal rate of technical substitution in attributes (MRTSA). Small changes in engine displacement have a trivial impact on costs. Thus, engine displacement is an obvious and important driver of MRTSA.[3] Engine displacement interacts with technology adoption. For example, adding

[2]Transmission upgrades may also allow carmakers to trade off fuel economy and acceleration. However, manufacturers typically choose to improve both attributes.

[3]The existence of competing design levers that differentially affect fuel economy and acceleration are also relevant. For example, a carmaker could choose between turbo-charging the engine, which only improves acceleration—or incurring the same

**Figure 3:** *Simulated fuel economy vs. acceleration (via engine displacement)*




Note: This figure shows simulated log fuel economy vs. log acceleration for five distinct drive-trains, with different combinations of fuel economy and acceleration achieved by changing engine displacement.

a turbocharger and reducing engine displacement ("engine downsizing") is a common method to increase fuel economy while maintaining acceleration (Shahed 2009). Once a turbocharger has been added, however, further reductions in engine size yield little fuel economy gains—but still large reductions in acceleration. Other technologies interact similarly with engine displacement.

To illustrate this interaction, we use evidence from a physics-based full vehicle simulation model—the Environmental Protection Agency's Advanced Light-Duty Powertrain and Hybrid Analysis (ALPHA) tool (Dekraker, Barba, Moskalik, and Butters (2018)), to investigate how the trade-off between fuel economy and power has shifted over time. Figure 3, which is reproduced from Moskalik, Bolon, Newman, and Cherry (2018) data, summarizes this change for five simulated vehicles. Each of the vehicles uses the same midsize sedan; the only difference is the drive-train technology. This allows the simulation to maintain the same road-load across all simulated vehicles.[4] The five power-trains include: (1) a 1980s carbureted engine with a three-speed automatic transmission; (2) a 2007 Toyota port fuel-injected (PFI) engine coupled to a five-speed transmission; (3) a 2013 GM gasoline direct injection (GDI) engine coupled to a six-speed transmission; (4) a 2017 Honda turbo-charged engine coupled with an eight-speed transmission; and (5) a future Ricardo 24

cost to improve aerodynamics, which also improves fuel economy but doesn't boost acceleration very much. The option to do one or the other implies a trade-off between fuel economy and acceleration.

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

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. The figure plots log fuel economy (miles per gallon) against log acceleration time (0-60 miles per hour in seconds) against log fuel economy (miles per gallon).

We highlight two key findings. First, for any given displacement, the later-vintage engines are both faster and more fuel efficient. The gain in acceleration has been greatest for smaller engines, while the gain in fuel economy has been largest for big engines. Engines have clearly improved across-the-board. Second, the correspondence between fuel economy and acceleration has flattened over time. A 1% decrease in acceleration is associated with a 0.75% gain in fuel economy in the 1980s, but only 0.33% gain in 2017. Thus, the opportunity cost improving fuel economy via reductions in engine size has increased—the fuel economy now comes at the expense of a much larger reduction in acceleration.

If displacement were the only pathway for trading-off fuel economy and acceleration, we would then conclude that the trade-off has substantially increased overall. However, while displacement is the most obvious source of a trade-off, there are other potential pathways. Below, we develop a theoretical model that shows how to infer the overall trade-off from car attribute choices, and we estimate changes in this trade-off using household-level microdata.

# 3 Theory of attribute production

We develop a theoretical model of attribute production, considering the interaction of consumer preferences, technology costs, gasoline prices, and policy. We use the model to derive optimal attribute choices in equilibrium comparative statics with respect to gasoline prices and technical change. Using the model, we show how to infer cost parameters and rates of technical change from observed market data on car attribute choices. We also demonstrate that discrete engine technologies are adopted on bigger, faster, more fuel-efficient cars, implying that unobserved costs are correlated with car attributes.

## 3.1 Modeling optimal attribute choice

This subsection sets up the model and derives necessary conditions and optimal attribute choices.

### 3.1.1 Setup

We assume a finite number $J$ of consumer types, indexed by $j$, which have heterogeneous preferences for car attributes; we suppress consumer type $j$ subscripts here and in much of what follows to keep the notation tidy. We explicitly model three car attributes, all bads: smallness ($s$), acceleration time ($a$), and gallons-per-mile ($g$). Smallness is the inverse of interior volume measured in cubic inches: $s \equiv 1/\text{volume}$. This inverse framing is non-standard but ensures proper curvature of costs with respect to preferences. Acceleration is the time

it takes to accelerate from 0 to 60 mph in seconds. Per-mile fuel consumption is the inverse of fuel economy: $g \equiv 1/\text{mpg}$. Each consumer has unit demand for a car and derives utility from their continuous choice of car attributes and spending on other goods. Utility is given by: $u(s,a,g) = v(s,a) + y - pgm - r(s,a,g)$, where $v(s,a)$ is utility derived from car attributes other than fuel economy, $y$ is income, $pgm$ is expenditures on fuel, and $r(s,a,g)$ is the equilibrium price of a car with a given set of attributes. Note that per-mile fuel consumption enters as a purely financial trade-off between fuel expenditures and car price. We assume that miles traveled ($m$) is exogenous, which is consistent with empirical evidence that demand for miles is inelastic; relaxing this assumption would not substantially change our conclusions. However, miles traveled varies across consumer types, which together with fuel prices ($p$) implies heterogeneous preferences for fuel economy.

On the supply side, we assume a common technology across carmakers and a constant variable cost for each car model: $c(s,a,g)$, which depends only on smallness, acceleration, and gallons-per-mile. We focus on these three attributes, since they are bound together in fundamental engineering relationships. Bigger cars are heavier, which slows them down. Larger engines can compensate but use more fuel. Hybrid engines and other technologies can save fuel but cost money. We assume that carmakers are price takers.[5] We further assume that there is free entry and zero fixed cost to develop and produce a new car model. Taken together, these assumptions imply that, in equilibrium, every consumer type chooses a car model custom-tailored to its own unique preferences, with equilibrium prices equal to costs: $r(s,a,g) = c(s,a,g)$. Note that the assumption of zero fixed costs will hold approximately, so long as the number of consumers of each type is sufficiently large.

Finally, we model a fuel-economy standard $\sigma$, which constrains average per-mile fuel consumption across each carmakers's fleet. We assume that this standard is implemented via a perfectly competitive credit-trading system with zero transaction costs, leading to equilibrium credit price $\tau$. Thus, the fuel-economy standard can be modeled as an implicit tax on each car's per-mile fuel consumption relative to the standard: $\tau[g-\sigma]$. Note that this tax becomes an implicit subsidy if the car is more efficient than the standard: $g < \sigma$.[6] If the fuel-economy standard fails to bind in equilibrium, then the credit price is zero ($\tau = 0$), and these implicit taxes and subsidies vanish.[7]

[5] 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.

[6] Under these assumptions, the fuel-economy standard is equivalent to a "feebate" policy: a system of direct fees on inefficient cars and rebates for efficient cars.

[7] Note that future iterations of this paper will consider the possibility, following recent policy changes, that the fuel-economy standard is an increasing function of car size, or increasing function of smallness. This attribute-based policy structure creates an indirect incentive to increase car size, in addition to reducing fuel consumption, as shown in Ito and Sallee (2018), Kellogg (2018), Kellogg (2020), and others.

Thus, in equilibrium, every consumer type faces the following maximization problem:

$$\max_{s,a,g} u(s,a,g) = v(s,a) + y - pgm - c(s,a,g) - \tau[g - \sigma], \tag{1}$$

which is utility from size and acceleration plus spending on other goods, which equals income less fuel expenditures, car technology costs, and policy incentives.

### 3.1.2 Functional-form assumptions

To derive clear predictions, we assume that costs follow a Cobb-Douglas-like functional form:

$$c(s,a,g) = e^{-\theta} s^{-\alpha_s} a^{-\alpha_a} g^{-\alpha_g}, \tag{2}$$

where $\theta$ is an index of attribute-neutral technical progress that scales production costs downward, and parameters $\alpha_s, \alpha_a, \alpha_g > 0$ relate percent changes in costs to percent changes in attributes. We again emphasize that these attributes are all bads: making them bigger reduces utility. Given our functional form, production costs are falling in each of these bads: $c_s, c_a, c_g < 0$. Thus, making a car more desirable by reducing these bads raises production costs. Further note that costs are convex with respect to each of these bads: $c_{ss}, c_{aa}, c_{gg} > 0$. Thus, reducing these bads drives up marginal costs, which makes sense. Intuitively, making a car bigger, faster, or more efficient should become increasingly costly. Finally, note that the cross-partials are all positive: $c_{sg}, c_{ag}, c_{sa} > 0$. Thus, reducing a bad is more costly when the other attributes are less bad. Again, this makes intuitive sense. Making a car bigger or more powerful should be more costly when the car is also more fuel-efficient.[8]

On the demand side, we assume that utility is linear in each of the bads with coefficients (marginal willingness-to-pay values) given by $\beta_s, \beta_a, \beta_g > 0$. Thus, our consumer maximization problem from above becomes:

$$\max_{s,a,g} y + \tau\sigma - \beta_s s - \beta_a a - \beta_g g - e^{-\theta} s^{-\alpha_s} a^{-\alpha_a} g^{-\alpha_g}, \tag{3}$$

where recall that $s$ is smallness as defined above. Note that this formulation implies that utility is nonlinear in car size (interior volume), which implies a strong desire to avoid tiny cars but a diminishing marginal willingness to pay for ever-larger cars. Our derivations above imply that marginal benefits from reductions in gallons-per-mile relate precisely to fuel prices, miles, and policy: $\beta_g \equiv pm + \tau$. However, our formulation based on a generic $\beta_g$ permits various behavioral extensions, such as undervaluation of fuel economy due to myopia, credit constraints, or other factors.

**Marginal Rate of Technical Substitution of Attributes (MRTSA).** Before moving to our necessary conditions, we pause briefly to explore the cost function in (2). Define the marginal rate of technical substitution of attributes (MRTSA) as the slope of the iso-cost curve in two-dimensional attribute space.

[8] Were we to model costs using a Cobb-Douglas-like function, with positive coefficients on goods rather than negative coefficients on bads, we would lose these desirable curvature properties.

The MRTSA measures how much of one attribute must be given up for an increase in another attribute, while holding production costs constant. The MRTSA is simply the (negative) marginal cost ratio for various combinations of attribute pairs:

$$MRTSA_{gs} \equiv \left.\frac{\partial g}{\partial s}\right|_{\Delta c=0} = -\frac{c_s}{c_g} = -\frac{\alpha_s}{\alpha_g}\frac{g}{s} \tag{4}$$

$$MRTSA_{ga} \equiv \left.\frac{\partial g}{\partial a}\right|_{\Delta c=0} = -\frac{c_a}{c_g} = -\frac{\alpha_a}{\alpha_g}\frac{g}{a} \tag{5}$$

$$MRTSA_{as} \equiv \left.\frac{\partial a}{\partial s}\right|_{\Delta c=0} = -\frac{c_s}{c_a} = -\frac{\alpha_s}{\alpha_a}\frac{a}{s}. \tag{6}$$

Note that $MRTSA_{gs}$ in the first row tells us how much gallons-per-mile (a bad) would decrease with a marginal increase in smallness (another bad), holding costs fixed. Note that the last equality follows from our specific functional-form assumption. Likewise, $MRTSA_{ga}$ in the second row tells us how much gallons-per-mile (a bad) would decrease with a marginal increase in acceleration time (another bad), holding costs fixed. Conversely, these values tell us how much fuel economy (a good) would decrease with a marginal increase in interior volume or reduction in acceleration time.

Meanwhile, the elasticity versions are given by:

$$\left.\frac{\partial g}{\partial s}\frac{s}{g}\right|_{\Delta c=0} = -\frac{c_s}{c_g}\frac{s}{g} = -\frac{\alpha_s}{\alpha_g} \tag{7}$$

$$\left.\frac{\partial g}{\partial a}\frac{a}{g}\right|_{\Delta c=0} = -\frac{c_a}{c_g}\frac{a}{g} = -\frac{\alpha_a}{\alpha_g} \tag{8}$$

$$\left.\frac{\partial a}{\partial s}\frac{s}{a}\right|_{\Delta c=0} = -\frac{c_s}{c_a}\frac{s}{a} = -\frac{\alpha_s}{\alpha_a}. \tag{9}$$

Note that the elasticity in the first row tells us the percent decrease in gallons-per-mile for percent increase in smallness, holding costs fixed, and likewise for the other two elasticities. The elasticity formulation is convenient, because the relationships among these bads, when measured in logs, are equivalent to the relationships among their corresponding goods: fuel economy (miles-per-gallon), size (interior volume), and acceleration capacity (proxied by the horsepower-to-weight ratio).

#### 3.1.3 Necessary conditions

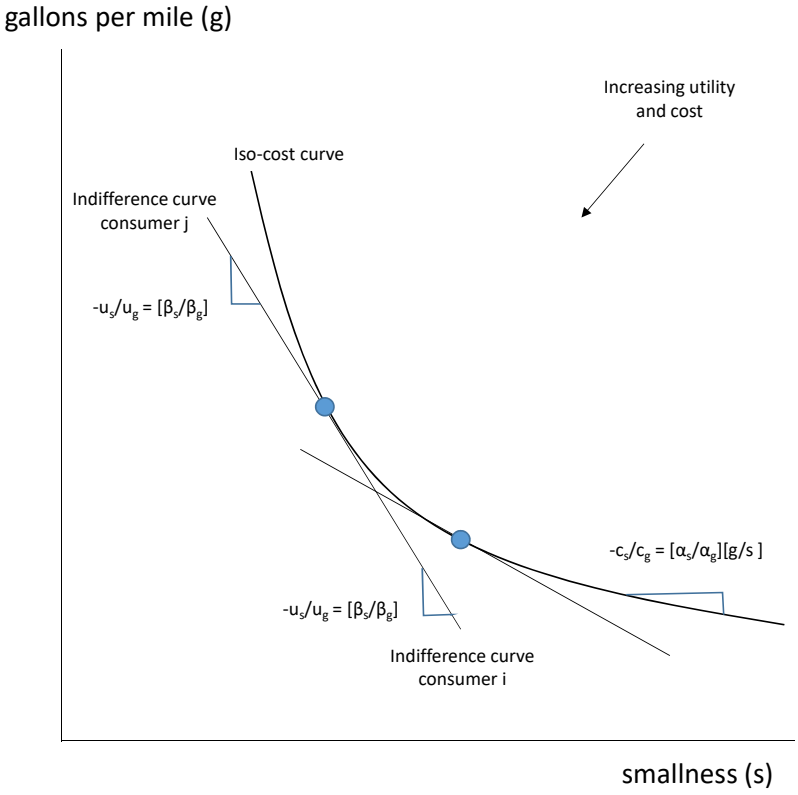
Maximization of equation (3) with respect to each of the continuous car attributes implies the following necessary conditions:

$$\text{s: } \beta_s = \frac{\alpha_s c(s,a,g)}{s} \tag{10}$$

$$\text{a: } \beta_a = \frac{\alpha_a c(s,a,g)}{a} \tag{11}$$

$$\text{g: } \beta_g = \frac{\alpha_g c(s,a,g)}{g}, \tag{12}$$

**Figure 4:** *Optimal attribute choice: MRS = MRTAS*




Note: This figure illustrates optimal attribute choices in two-dimensional (g-s) space as a point of tangency between the MRS for gpm and smallness and the MRTAS for gallons-per-mile and smallness.

where note that we are relying both on our assumtions of linear utility and Cobb-Douglas-like costs. These conditions say that the marginal utility to reduce a bad attribute (left side) should equal the marginal cost of reducing said attribute (right side).

Now divide the left and ride sides for each of the pairwise combinations of (10)–(12) to yield the following three equations involving ratios of marginal utility and marginal costs:

$$-\frac{\beta_s}{\beta_g} = -\frac{c_s(s,a,g)}{c_g(s,a,g)} = -\frac{\alpha_s}{\alpha_g}\frac{g}{s} \tag{13}$$

$$-\frac{\beta_a}{\beta_g} = -\frac{c_a(s,a,g)}{c_g(s,a,g)} = -\frac{\alpha_a}{\alpha_g}\frac{g}{a} \tag{14}$$

$$-\frac{\beta_a}{\beta_s} = -\frac{c_a(s,a,g)}{c_s(s,a,g)} = -\frac{\alpha_s}{\alpha_a}\frac{a}{s}. \tag{15}$$

These equations say that the marginal rate of substitution (MRS) between any combination of attributes should equal the corresponding marginal rate of technical substitution of attributes (MRTSA) in production. Figure 4 illustrates this condition for equation (13) as points of tangency between level sets of utility (indifference curves) and level sets of production costs (iso-cost curves) in $g - s$ space.

### 3.1.4 Optimal attribute choices

Solving equations (10)–(12) for the three unknown car attributes and then taking logs then yields equilibrium attribute choices as a function of underlying preference and technology parameters:

$$\ln s = \frac{1}{\psi}\left[-\theta - (1+\alpha_a+\alpha_g)\ln\beta_s + \alpha_a\ln\beta_a + \alpha_g\ln\beta_g + (1+\alpha_a+\alpha_g)\ln\alpha_s - \alpha_a\ln\alpha_a - \alpha_g\ln\alpha_g\right] \quad (16)$$

$$\ln a = \frac{1}{\psi}\left[-\theta - (1+\alpha_s+\alpha_g)\ln\beta_a + \alpha_s\ln\beta_s + \alpha_g\ln\beta_g + (1+\alpha_s+\alpha_g)\ln\alpha_a - \alpha_s\ln\alpha_s - \alpha_g\ln\alpha_g\right] \quad (17)$$

$$\ln g = \frac{1}{\psi}\left[-\theta - (1+\alpha_a+\alpha_s)\ln\beta_g + \alpha_s\ln\beta_s + \alpha_a\ln\beta_a + (1+\alpha_a+\alpha_s)\ln\alpha_g - \alpha_s\ln\alpha_s - \alpha_a\ln\alpha_a\right], \quad (18)$$

where $\psi \equiv 1+\alpha_s+\alpha_a+\alpha_g$. Note heterogeneity in the preference parameters $(\beta_s, \beta_a, \beta_g)$ across consumer types will give rise to an array of car model offerings in equilibrium. These equations clarify that a consumer's equilibrium choices for a given car attribute depend on their preferences parameters for *all* attributes, along with cost parameters that govern technical trade-offs. Consumers with different preferences in any one dimension will therefore choose different car attributes in all three dimensions.

To illustrate, suppose that the preference parameters on smallness $(\ln\beta_s)$ and acceleration time $(\ln\beta_a)$ are identical across consumer types such that variation in equilibrium car attribute bundles derives solely from heterogeneity in the preference parameter on gallons-per-mile $(\ln\beta_g)$. Inspection of equations (16)–(18) reveals that consumers with stronger preferences to save fuel (larger $\ln\beta_g$) will choose smaller, less powerful, and more efficient cars. Further, because equations (16)–(18) are linear, any pair of logged attributes will be perfectly correlated across car models. In practice, we do not see a perfect correlation of car attributes across car models, implying preference heterogeneity in multiple dimensions.

## 3.2 Optimal fuel economy conditional on other attributes

We now solve for optimal gallons-per-mile *conditional* on other car attributes, to motivate our empirical analysis below, and to connect to Knittel (2011) and other papers that regress log fuel economy on car attributes. Substitute for the cost function in equation (12) using the specific functional form in (2), solve for gallons-per-mile $(g)$, and then take logs to yield:

$$\ln g = -\frac{\theta}{1+\alpha_g} - \frac{1}{1+\alpha_g}\ln\beta_g - \frac{\alpha_s}{1+\alpha_g}\ln s - \frac{\alpha_a}{1+\alpha_g}\ln a \quad (19)$$

$$\approx -\frac{\theta}{1+\alpha_g} - \frac{1}{1+\alpha_g}\ln pm - \frac{\alpha_s}{1+\alpha_g}\ln s - \frac{\alpha_a}{1+\alpha_g}\ln a, \quad (20)$$

where the approximation in the second row follows from $\beta_g = pm + \lambda$ and the assumption that $\lambda \approx 0$. This approximation is consistent with back-of-the-envelope calculations, which suggest that $pm$ is at least an order of magnitude larger than $\lambda$.[9] This equation suggests that we can recover cost parameters $(\alpha_s, \alpha_a, \alpha_g)$ and the annual rate of attribute-neutral technical change $(\partial\theta/\partial t)$ from a regression of log fuel economy on log size,

[9]Consider an improvement in fuel economy from 20 to 25 miles-per-gallon, implying a $1/20 - 1/25 = 0.01$ reduction in gallons-per-mile. Assuming a fuel price of \$3 per gallon and 100,000 miles of lifetime travel (present value), this translates to

log acceleration, and a time trend. However, we must control for fuel-economy preferences ($\ln \beta_g \approx \ln pm$). Why? Even holding size and acceleration fixed, this equation implies that consumers with a strong preference to save fuel will choose more fuel-efficient (and therefore more costly) cars. Meanwhile, equations (16)–(17) show that consumers with strong preference to save fuel will choose smaller, less powerful cars. Thus, the fuel-economy preferences term ($\ln \beta_g/(1+\alpha_g)$) will be correlated with the other car attributes. Failing to control for this term will therefore remand it to the error term, leading to biased coefficient estimates.[10]

Below, we show how to recover cost parameters from household-level micro-data by regressing log fuel economy choices on log gasoline prices, log miles, and log car attributes, while controlling for household-level demographics. Using household-level micro-data (as opposed to model-level attribute data) allows us to control for local gasoline prices, household-level miles traveled, and other observables that might correlate with fuel-economy preferences (e.g., income). Our key insight is that variation in fuel economy choices (and therefore car costs) conditional on other car attributes comes from variation in willingness-to-pay for improved fuel economy, which derives mainly from variation in gasoline prices and miles traveled. Gasoline prices and miles traveled are observable. Thus, we can control for car costs indirectly by controlling for gasoline prices and miles in a household-level regression. Meanwhile, the coefficient on log gasoline prices identifies the cost parameter on fuel economy, allowing us to back out the other cost parameters.

### 3.2.1 Knittel's (2011) regressions are likely biased

How does our analysis differ from Knittel (2011)? We both seek to estimate the slope of the iso-cost curve at a point in time, along with shifts in this curve over time, which reflect technical change. But Knittel jumps directly to the cost function, implicitly ignoring the role of consumer preferences in driving equilibrium attribute bundles. Taking logs of the cost function in equation (2) and rearranging yields:

$$\ln g = -\frac{\theta}{\alpha_g} - \frac{\alpha_s}{\alpha_g} \ln s - \frac{\alpha_a}{\alpha_g} \ln a - \frac{1}{\alpha_g} \ln c, \tag{21}$$

which is the form of Knittel's (2011) baseline regressions: log fuel economy on log attributes.[11] Thus, our functional form for costs provides micro-foundations for Knittel (2011) and related papers.

Note that the coefficient on log smallness ($-\alpha_s/\alpha_g$) captures the slope of the iso-cost curve in two-dimensional product space: percent changes in $g$ that hold costs fixed given a percent change in $s$. Likewise for the coefficient on log acceleration time ($-\alpha_a/\alpha_g$). Note, however, that technology costs ($\ln c$) are unobserved

$0.01 \cdot 3 \cdot 100,000 = \$3000$ of fuel savings. Meanwhile, the non-compliance penalty under U.S. fuel economy standards is \$55 per mile-per-gallon per car, which translates to $5 \cdot 55 = \$225$ for a 5 mile-per-gallon deviation. Note that the non-compliance penalty is an upper bound on marginal compliance costs. Indeed, Anderson and Sallee (2011) show that marginal costs are under \$30 in many years. These calculations imply that $\lambda$ is an order of magnitude (or more) smaller than $pm$. Thus, $\ln(\beta_g) = \ln(pm) + \ln(1 + \lambda/pm) < \ln(pm) + \ln(1 + 0.1) \approx \ln(pm)$.

[10] Ignoring this term will also make it impossible to identify $\beta_g$, which is necessary to infer the other cost parameters.

[11] Our equation has gallons-per-mile as the outcome, while Knittel's 2011 has miles-per-gallon. Likewise, we model acceleration time as the horsepower-weight ratio, putting weight in the denominator, while Knittel includes weight in a numerator. Given logs, these formulations are all equivalent, with signs on the coefficients appropriately flipped.

in Knittel (2011). Thus, causal identification via regression requires that technology costs be uncorrelated with size and acceleration. Our analysis shows that this assumption is unlikely to hold. Conditional on other car attributes, consumers with a strong desire to save fuel will purchase more efficient cars (equation 20), putting them on a higher iso-cost curve. Meanwhile, these same consumers will tend to choose different attributes (equations 16 and 17), which implies different marginal costs for fuel economy (equation 2). Thus, costly fuel-saving technologies will not be added to cars at random. We show this point theoretically below, in our analysis of discrete technology adoption, and we confirm it empirically through an examination of penetration rates for various engine technologies.

### 3.3 Gasoline prices and fuel-economy standards

We now explore equilibrium responses to higher fuel prices. Recall that $\beta_g \equiv pm + \lambda$ and assume that $\lambda \approx 0$, such that percent changes in fuel prices correspond directly to percent changes in $\beta_g$. Differentiating equations (16)–(18) with respect to log fuel-economy preferences ($\ln \beta_g$) yields:

$$\frac{\partial \ln s}{\partial \ln \beta_g} = \frac{\alpha_g}{\psi} > 0 \tag{22}$$

$$\frac{\partial \ln a}{\partial \ln \beta_g} = \frac{\alpha_g}{\psi} > 0 \tag{23}$$

$$\frac{\partial \ln g}{\partial \ln \beta_g} = -\frac{1 + \alpha_s + \alpha_a}{\psi} < 0, \tag{24}$$

where recall that $\psi \equiv 1+\alpha_s+\alpha_a+\alpha_g$. Note that car smallness increases (car size decreases) and acceleration time increases. Meanwhile, gallons-per-mile decreases (fuel economy increases).

Note that the percentage changes in smallness and acceleration are given by $\alpha_g/\psi$, which reflects the relative steepness of the cost function in the direction of gallons-per-mile. The costlier it is to reduce gallons-per-mile relative to other bads, the more smallness increases (size decreases) and acceleration time increases.[12] Meanwhile, the percentage change in gallons-per-mile is given by $(1 + \alpha_s + \alpha_a)/\psi$, which is inversely related to the steepness of the cost function in the direction of miles-per-gallon.[13] Thus, the costlier it is to reduce gallons-per-mile relative to other bads, the smaller the reduction in gallons-per-mile. These results may help explain the divergent responses to gasoline prices in the 1980s vs. 2000s. Our engineering simulations suggest that $\alpha_g$ has fallen over time relative to the other cost parameters, implying that large changes in acceleration now yield relatively modest improvements in fuel economy. Thus, higher fuel prices no longer lead to large reductions in acceleration. [Asa: This comparative statics seems not to be working out. Larger $\alpha_g$ implies that optimal acceleration is more responsive to gasoline prices. What am I missing?]

Note that increasing the stringency of the fuel-economy standard ($\lambda$) has qualitatively similar effects. To see this, observe that $\partial \ln \beta_g / \partial \lambda \approx 1/pm$ when evaluated at $\lambda \approx 0$. Thus, just re-scale each of the effects in

[12] Note that the percentage change is 0 when $\alpha_g = 0$ and approaches 1 as $\alpha_g$ grows to infinity.
[13] Recall that the denominator of this function is $1 + \alpha_s + \alpha_a + \alpha_g$.

equations (24)–(24) by $1/pm$ to yield the marginal percentage effects.

Now consider the effect of higher gasoline prices on fuel economy, conditional on other attributes. Differentiating equation (20)) with respect to log fuel-economy preferences ($\ln \beta_g$) yields:

$$\frac{\partial \ln g}{\partial \ln \beta_g} = -\frac{1}{1+\alpha_g} < 0, \tag{25}$$

which shows that gallons-per-mile decreases conditional size and acceleration. Thus, consumers optimally choose cars that smaller and less powerful overall (equations 22 and 23). In addition, for cars of a given size and performance, consumers optimally choose more efficient and therefore more costly cars. Note that the partial derivative in equation (25) is only a function of the cost parameter $\alpha_g$). This result suggests that exogenous variation in fuel prices, controlling for size and acceleration, can be used to identify the cost coefficient on gallons-per-mile, which is essential to identifying the other cost parameters ($\alpha_s, \alpha_a$), as well as the rate of technical change ($\partial\theta/\partial t$).

## 3.4 Technical change and car attributes

We now consider three different forms of technical change and explore their impacts on equilibrium attributes.

### 3.4.1 Attribute-neutral or unbiased technical change

We model unbiased or attribute-neutral technical change via a marginal increase in $\theta$, which proportionally reduces costs and marginal costs, and which preserves the ratios of marginal costs in equations (4)–(6). Visually, this is the same as shifting iso-cost curves inward in goods space such that, for any level of cost, all bads may be reduced.

First consider a marginal increase in $\theta$ in equations (16)–(18). We can easily see that this results in an equal proportional decrease in all three bads. Differentiating gallons-per-mile in equation (18) with respect to $\theta$ yields:

$$\frac{\partial \ln g}{\partial \theta} = -\frac{1}{\psi}, \tag{26}$$

and likewise for smallness and acceleration time. Thus, unbiased technical change leads to proportional increase in size, decrease in acceleration time, and increase in fuel economy. These gains are inversely proportional to the exponent parameters in the cost function, which equal the elasticity of marginal costs with respect to reductions in the bad attributes. Steeply increasing marginal costs are a drag on technical change's ability to improve car attributes.

Now consider a marginal increase in $\theta$ in equation (20). Differentiating with respect to $\theta$ yields:

$$\frac{\partial \ln g}{\partial \theta} = -\frac{1}{1+\theta_g}. \tag{27}$$

Thus, technical change manifests as a shift in equilibrium fuel economy conditional on other attributes. The size of this shift in percentage terms is the rate of technical change itself, scaled by the exponent parameter on fuel economy. The larger this parameter, the smaller the shift in fuel economy.

#### 3.4.2 Attribute-biased technical change

COMING SOON

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.

#### 3.4.3 Discrete technology adoption

Periodically, manufacturers invent new technologies that may be added to vehicles to improve one or more attributes. For example, turbochargers temporarily increase the performance of an internal combustion by forcing additional compressed air into the combustion chamber. Adding a turbocharger to a car's drive-train allows it to accelerate faster. Carmakers often add turbochargers and simultaneously decrease engine displacement ("downsizing"), thereby improving fuel economy while holding acceleration constant (Shahed and Bauer 2009). Over time, the cost of adding a turbocharger has fallen—a form of technical change—leading to increased adoption. Other examples include continuously variable transmission, direct injection (replacing carburetor), and electric or hybrid-electric drive-trains, all of which have experienced falling costs and increased adoption. How does this form of technical change manifest in the car market? Are new technologies adopted uniformly across the fleet? Or is their adoption biased toward specific consumer segments or cars with particular attributes? If so, what are the implications?

We consider a discrete technology, targeting a specific car attribute—acceleration, to use the turbocharger example—that consumers may add to their car at a cost. This technology has two important features. First, the installation cost ($k$) in a given year is the same for all cars regardless of baseline attributes. Second, the technology yields a fractional reduction ($\omega$) in the targeted bad attribute. Thus, vehicles with a higher baseline level of the targeted attribute (acceleration time) will see a larger absolute reduction.

We consider a technology that affects the *benefit* of the targeted attribute as perceived by the consumer. Conditional on adopting the technology, and ignoring installation cost, the consumer experiences utility given by:

$$u(s, a, g; \omega) = y + \tau\sigma - \beta_s s - \beta_a(1 - \omega)a - \beta_g g - c(s, a, g), \tag{28}$$

where we have simply scaled the marginal disutility of acceleration time by $1 - \omega$, implying a fractional reduction of $\omega$ in this bad attribute. By the envelope theorem, the gain in maximized utility for a marginal increase in $\omega$ is given by: $\partial u(s, a, g; \omega)/\partial\omega = \beta_a a$. Thus, the net benefit from adopting the discrete technology

is given by the following approximation:

$$u(s,a,g;\omega>0)-u(s,a,g;\omega=0)-k\approx\beta_a\omega a-k, \tag{29}$$

where the first term in the approximation is $\omega$ times the marginal utility gain evaluated at $\omega=0$. For any $k>0$, consumers that like fast cars (large $\beta_a$) but who are stuck in slow cars (large $a$) are the ones most likely to benefit from adoption.

It can be shown that adoption benefits ($\beta_a\omega k$) are increasing in consumer preferences for size ($\beta_g$) and fuel economy ($\beta_g$), as well as acceleration ($\beta_a$). Intuitively, because car attributes are in technical competition, consumers with extreme preferences for size and fuel economy get stuck in slow cars. It's exactly these consumers who would benefit most from a boost in acceleration—especially if they also crave speed. Note, there is nothing special about acceleration in our analysis—these same qualitative results hold for technologies targeting size and fuel economy. Thus, it is generally the consumers who want more of anything and everything who are most likely to adopt an attribute-enhancing technology. Only as costs ($k$) gradually fall over time will the more ambivalent consumers adopt. This result conforms with the stylized fact that new car technologies often first make their appearance in luxury segments, and only gradually filter down to the rest of the market.

A key implication of these results is that 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 technical trade-offs by fitting a line through attribute data as in Knittel (2011), we must assume that unobserved costs are uncorrelated with car attributes. However, we have shown that discrete technologies are more likely to be adopted by consumers with extreme preferences, who choose bigger, faster, more fuel-efficient cars. That is, unobserved costs are inevitably correlated with attributes. We confirm this prediction empirically in the following section.

What happens to equilibrium attributes for consumers that do adopt? Consider equations (16)–(18) and note that adopting the acceleration-boosting technology is equivalent to replacing acceleration preferences ($ln\beta_a$) with modified preferences ($\ln(1-\omega)\beta_a$) in each of these equations. Differentiating each of these equations with respect to $\omega$ and evaluating at $\omega=0$ then yields:

$$\left.\frac{\partial\ln s}{\partial\omega}\right|_{\omega=0}=-\frac{\alpha_a}{\psi}<0 \tag{30}$$

$$\left.\frac{\partial\ln(1-\omega)a}{\partial\omega}\right|_{\omega=0}=-\frac{\alpha_a}{\psi}<0 \tag{31}$$

$$\left.\frac{\partial\ln g}{\partial\omega}\right|_{\omega=0}=-\frac{\alpha_a}{\psi}<0, \tag{32}$$

where note that the relevant attribute in the second row is log acceleration time as perceived by the consumer. Strikingly, consumers that adopt the discrete technology experience improvements in all attributes, i.e. even

those not targeted by the technology. Even more strikingly, the proportional improvement is the same across all three attributes. Thus, the impact of the technology is the same as for attribute-neutral technical change—but only among the adopters.

These same results all obtain when we assume that that the discrete technology affects the *cost* of supplying the targeted attribute: $c(s, a(1+\omega), g) = (1+\omega)^{-\alpha_a} c(s, a, g)$, where the bad attribute as perceived by the consumer remains fixed but this bad attribute increases by fraction $\omega$ in the cost function, driving down costs. Note that scaling the targeted attribute by $(1+\omega) > 1$ is equivalent to scaling the entire cost function by $(1+\omega)^{-\alpha_a} < 1$, given our functional-form assumption. Thus, adoption will be most attractive to consumers facing high technology costs—the ones choosing big, fast, and fuel-efficient cars. Further, since adoption is equivalent to attribute-neutral technical change, it will lead to equal proportional improvements in size, acceleration, and fuel economy among the adopters—same as a technology affecting the perceived benefit of the targeted attribute.

# 4 Data and empirical estimation

## 4.1 Fuel economy conditional on size and acceleration

See Tables 1 2

## 4.2 Trends in fuel economy, size, and acceleration

COMING SOON

## 4.3 Adoption of fuel-saving technologies is biased toward fuel-using attributes

TO BE UPDATED WITH NEW DATA

We test the prediction that new technologies are first adopted by consumer groups with higher overall preferences for attributes in the context of turbocharger adoption. Turbochargers improve a vehicle's horsepower and are often paired with decreasing engine size to improve fuel economy without sacrificing performance. If turbochargers are first adopted by consumers with overall high preferences for all desirable attributes, then their adoption should be correlated with both the attributes that they improve (performance and fuel economy) and those they do not (size). Figure 7 shows that turbochargers were taken up first by heavier vehicles, and later disseminated throughout the fleet.

To test the effect we match vehicle models over time and then test if prior year attributes are predictive of a change in turbocharger status. Since vehicle models are designed for specific consumer groups, the attributes of the prior year are a proxy for that group's preferences.

We use Wards and EPA data for the model years 1996 to 2019. These are years where there was sub-

stantial uptake of turbochargers and for which we have reliable data on turbochargers. However, by 1996 turbochargers were not new.

To test our prediction, we want to see if increasingly lower levels of prior year attributes predict turbocharger adoption in later years. Our empirical model takes the form

$$\text{turbo}_t = \gamma_t + \ln \boldsymbol{X}_{t-1}\zeta_t + \theta \text{turbo}_{t-1} + \epsilon \tag{33}$$

where $\boldsymbol{X_{t-1}}$ are attributes. Logged attribute levels are de-meaned by the sales weighted average log of attributes for that year. De-meaning allows for estimates of the effect of attributes to be relative to the average in a given year.

Table 3 shows the results. Column 1 is the model without a time interaction with attributes. Columns 2 and 3 interact attributes with a time trend.Column 3 interacts an indicator for trucks with attributes in addition to the time trend.Each model includes make-year fixed effects. The fixed effects control for the average rate of adoption for each within each vehicle badge.

Column 1 shows that, averaged over all years, changes in turbocharger status are correlated with lower prior model year fuel economy and volume. In columns 2 and 3 the coefficient on logged attributes interacted with a time trend are all negative and significant, while the un-interacted coefficients are no longer significant. This provides evidence in support of the prediction. As we move forward in time, the vehicles that have yet to adopt turbochargers are those with lower relative attribute levels.

These results are important for interpreting empirical models in the literature. New technologies are not randomly assigned. However, previous research has assumed a constant average cost across the range of attributes to estimate attribute trade-offs for a fixed cost. If technologies are unevenly applied across attribute over time, this assumption cannot be true.

# 5 Implications for efficiency standards and electric vehicles

[COMING SOON]

# 6 Conclusion

[COMING SOON]

# References

Anderson, Soren T and James M Sallee. 2011. "Using Loopholes to Reveal the Marginal Cost of Regulation: The Case of Fuel-Economy Standards." 101 (4):1375–1409. URL http://pubs.aeaweb.org/doi/10.1257/aer.101.4.1375.

———. 2016. "Designing Policies to Make Cars Greener." *Annual Review of Resource Economics* 8:157–180.

Dekraker, Paul, Daniel Barba, Andrew Moskalik, and Karla Butters. 2018. "Constructing Engine Maps for Full Vehicle Simulation Modeling." URL http://www.sae.org/content/2018-01-1412/.

Ito, Koichiro and James M. Sallee. 2018. "The Economics of Attribute-Based Regulation: Theory and Evidence from Fuel Economy Standards." *The Review of Economics and Statistics* 100 (2):319–336.

Kellogg, Ryan. 2018. "Gasoline Price Uncertainty and the Design of Fuel Economy Standards." *Journal of Public Economics* 160:14–32.

———. 2020. "Output and Attribute-based Carbon Regulation Under Uncertainty." *Journal of Public Economics* 190:104246.

Klier, Thomas and Josh Linn. 2015. "The Effect of Vehicle Fuel Economy Standards on Technology Adoption." :56.

Knittel, Christopher R. 2011. "Automobiles on Steroids: Product Attribute Trade-Offs and Technological Progress in the Automobile Sector." *American Economic Review* 101 (7):3368–3399.

MacKenzie, Don and John Heywood. 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. URL http://journals.sagepub.com/doi/10.3141/2287-15.

Moskalik, Andrew, Kevin Bolon, Kevin Newman, and Jeff Cherry. 2018. "Representing GHG Reduction Technologies in the Future Fleet with Full Vehicle Simulation." URL http://www.sae.org/content/2018-01-1273/.

Newell, Richard G, Adam B Jaffe, and Robert N Stavins. 1999. "The Induced Innovation Hypothesis and Energy-saving Technological Change." *The Quarterly Journal of Economics* 114 (3):941–975.

Shahed, SM and Karl-Heinz Bauer. 2009. "Parametric studies of the impact of turbocharging on gasoline engine downsizing." *SAE International Journal of Engines* 2 (1):1347–1358.

USEPA. 2019. "The 2019 Automotive Trends Report EPA-420-R-20-006." URL https://www.epa.gov/automotive-trends/download-automotive-trends-report.

Whitefoot, Kate S., Meredith L. Fowlie, and Steven J. Skerlos. 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. URL http://pubs.acs.org/doi/10.1021/acs.est.7b03743.

# 7 Tables and figures

**Table 1:** *Conditional regressions*

| height | Model 1 | Model 2 | Model 3 | Model 4 | Model 5 |
|---|---|---|---|---|---|
| time trend ($\frac{\theta}{1+\alpha_g}$) | 0.015 | 0.015 | 0.015 | 0.015 | 0.011 |
| | (0.0001) | (0.0001) | (0.0001) | (0.0001) | (0.0003) |
| log gas price × miles ($\frac{1}{1+\alpha_g}$) | | −0.002 | | | |
| | | (0.0008) | | | |
| log gas price ($\frac{1}{1+\alpha_g}$) | | | 0.016 | 0.016 | −0.074 |
| | | | (0.002) | (0.002) | (0.007) |
| log gas price t-1 | | | | | 0.047 |
| | | | | | (0.008) |
| log gas price t-2 | | | | | 0.079 |
| | | | | | (0.006) |
| log miles | | | −0.004 | | |
| | | | (0.0008) | | |
| log HP/Wt ($-\frac{\alpha_a}{1+\alpha_g}$) | −0.527 | −0.528 | −0.526 | −0.526 | −0.525 |
| | (0.002) | (0.002) | (0.002) | (0.002) | (0.002) |
| log volume ($-\frac{\alpha_s}{1+\alpha_g}$) | −0.439 | −0.439 | −0.438 | −0.438 | −0.439 |
| | (0.003) | (0.003) | (0.003) | (0.003) | (0.003) |
| (Intercept) | 3.203 | 3.220 | 3.220 | 3.187 | 3.137 |
| | (0.007) | (0.010) | (0.010) | (0.007) | (0.009) |
| Num.Obs. | 48 698 | 48 698 | 48 698 | 48 698 | 48 698 |
| R2 | 0.565 | 0.565 | 0.566 | 0.566 | 0.569 |
| RMSE | 0.14 | 0.14 | 0.14 | 0.14 | 0.14 |
| Sum of log gas price | | | | | 0.051 |
| $\theta$ (annual rate of $\Delta T$) | 0.015 | -8.659 | 0.904 | 0.907 | 0.222 |
| $\frac{\alpha_a}{\alpha_g}$ (MRATS FE for accel) | 0.527 | 0.527 | 0.534 | 0.534 | 0.553 |
| $\frac{\alpha_s}{\alpha_g}$ (MRATS FE for size) | 0.439 | 0.438 | 0.445 | 0.445 | 0.462 |

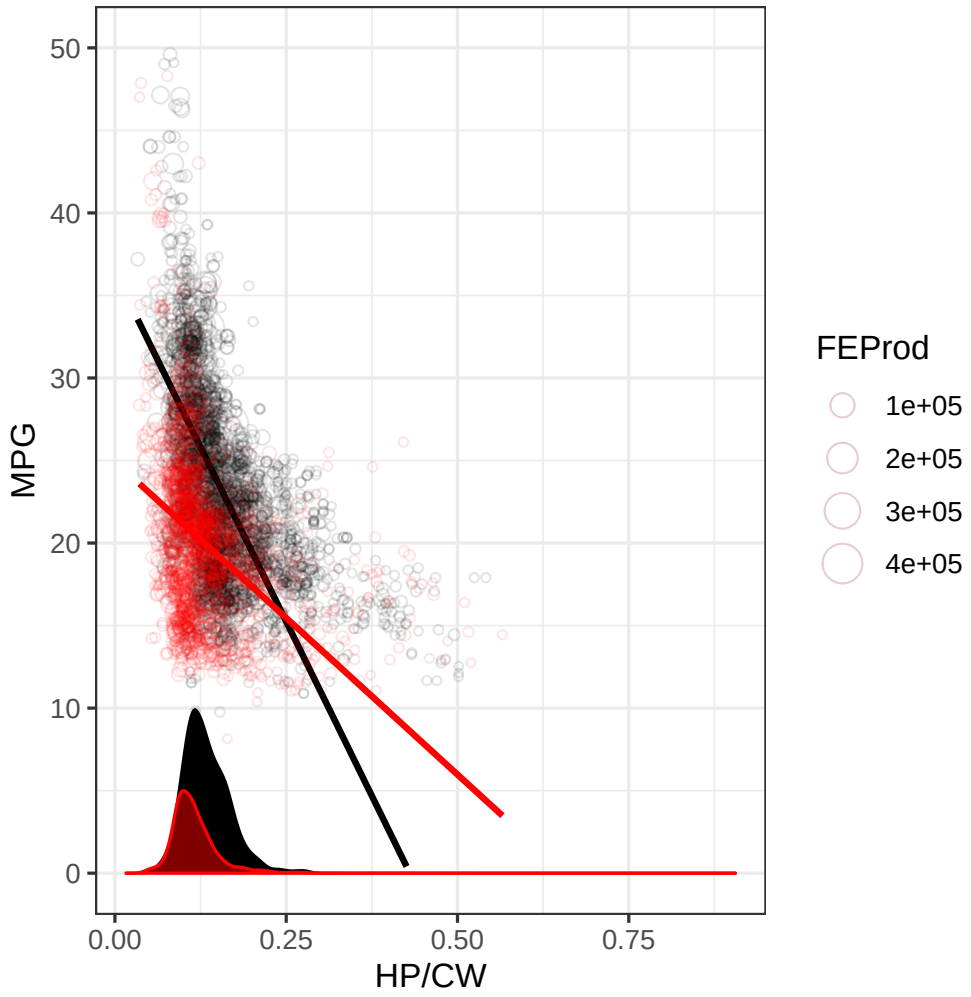
Note:

**Table 2:** *Conditional regressions with biased tehcnical change*

| height | Model 1 | Model 2 | Model 3 | Model 4 |
|---|---|---|---|---|
| time trend ($\frac{\theta}{1+\alpha_g}$) | 0.012 | 0.013 | 0.014 | 0.011 |
| | (0.001) | (0.001) | (0.001) | (0.001) |
| log gas price ($\frac{1}{1+\alpha_g}$) | 0.093 | 0.034 | 0.040 | 0.060 |
| | (0.007) | (0.019) | (0.008) | (0.020) |
| log gas price t-1 | | 0.008 | | −0.156 |
| | | (0.023) | | (0.027) |
| log gas price t-2 | | 0.047 | | 0.132 |
| | | (0.011) | | (0.013) |
| log HP/Wt ($-\frac{\alpha_a}{1+\alpha_g}$) | −0.524 | −0.525 | −0.521 | −0.522 |
| | (0.003) | (0.003) | (0.003) | (0.003) |
| log volume ($-\frac{\alpha_s}{1+\alpha_g}$) | −0.451 | −0.454 | −0.446 | −0.449 |
| | (0.004) | (0.004) | (0.004) | (0.004) |
| trend × log gas price | 0.016 | 0.012 | 0.013 | 0.019 |
| | (0.0008) | (0.002) | (0.0008) | (0.002) |
| trend × log gas price t-1 | | −0.002 | | −0.012 |
| | | (0.002) | | (0.002) |
| trend × log gas price t-2 | | −0.0005 | | 0.013 |
| | | (0.001) | | (0.002) |
| trend × log HP/Wt | −0.0001 | −0.0001 | −0.0003 | −0.0003 |
| | (0.0004) | (0.0004) | (0.0004) | (0.0004) |
| trend × log volume | −0.002 | −0.003 | −0.002 | −0.002 |
| | (0.0005) | (0.0005) | (0.0005) | (0.0005) |
| (Intercept) | 3.152 | 3.147 | 3.180 | 3.166 |
| | (0.012) | (0.013) | (0.026) | (0.026) |
| Num.Obs. | 48 698 | 48 698 | 48 698 | 48 698 |
| R2 | 0.569 | 0.570 | 0.574 | 0.576 |
| RMSE | 0.14 | 0.14 | 0.14 | 0.14 |
| Demographics and state fixed effects | NO | NO | YES | YES |
| Sum of log gas price | | 0.089 | | 0.036 |
| Sum of log gas price trend | | 0.009 | | 0.019 |
| $\theta$ (annual rate of $\Delta T$) | 0.127 | 0.148 | 0.345 | 0.311 |
| $2017\frac{\alpha_a}{\alpha_g}$ (MRATS FE for accel) | 0.578 | 0.576 | 0.543 | 0.542 |
| $2001\frac{\alpha_a}{\alpha_g}$ (MRATS FE for accel) | 0.452 | 0.495 | 0.442 | 0.407 |
| $\Delta\frac{\alpha_a}{\alpha_g}$ | 0.017 | 0.01 | 0.015 | 0.02 |
| $2017\frac{\alpha_s}{\alpha_g}$ (MRATS FE for size) | 0.497 | 0.498 | 0.465 | 0.465 |
| $2001\frac{\alpha_s}{\alpha_g}$ (MRATS FE for size) | 0.361 | 0.392 | 0.352 | 0.324 |
| $\Delta\frac{\alpha_s}{\alpha_g}$ | -0.034 | 0.014 | 0.019 | 0.024 |

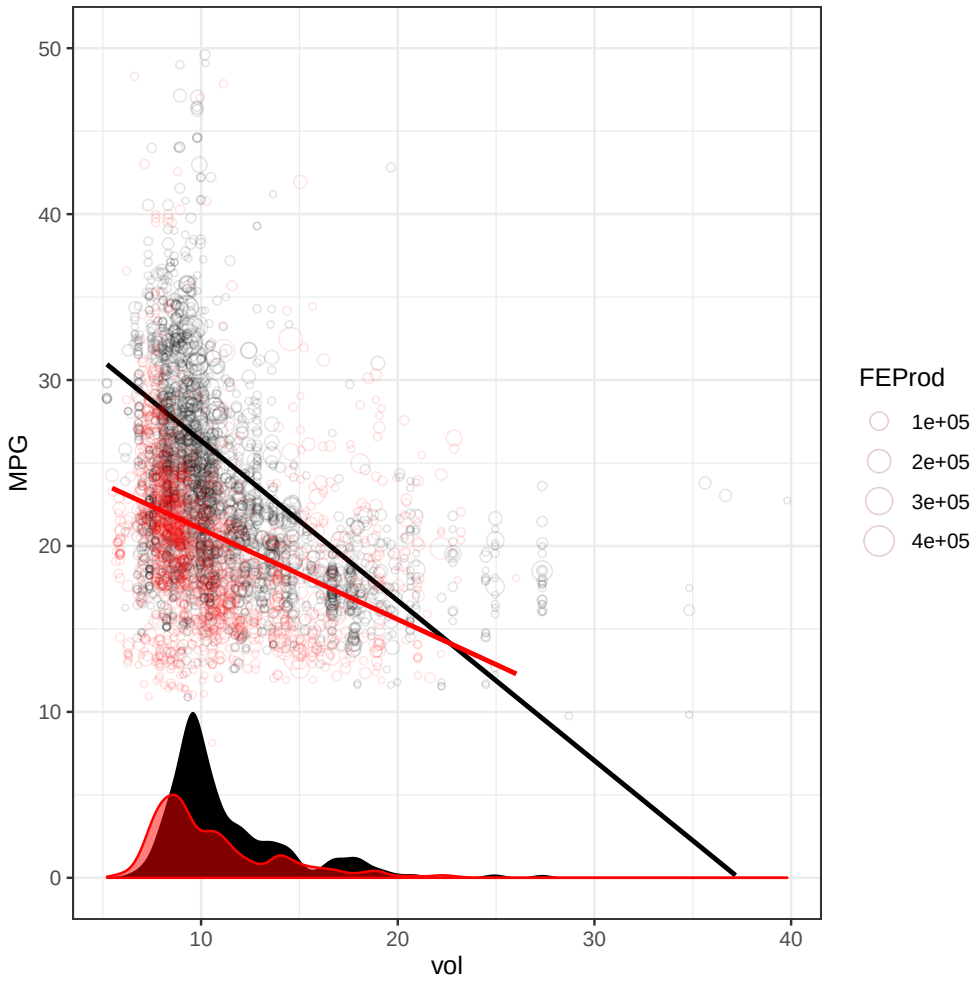
Note:

**Figure 5:** *HP/WT vs MPG*



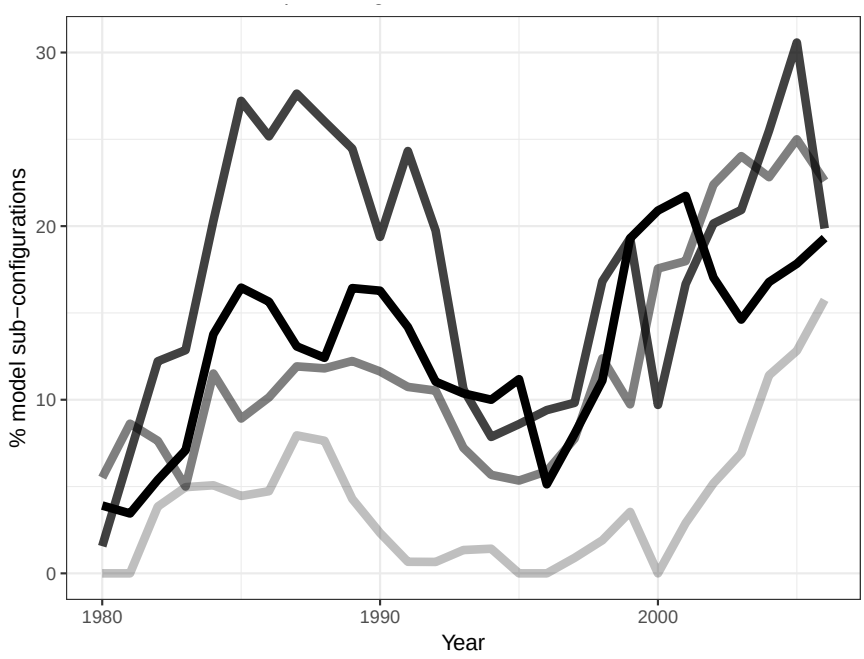

Note: Red shows 1999-2001; black shows 2015-2017.

**Figure 6:** *Size vs MPG*




Note: Red shows 1999-2001; black shows 2015-2017.

**Figure 7:** *Turbocharger adoption by weight quartile*




Note: Darker lines indicate higher weight quartiles. Aggregation is at the fleet level. The figure shows that heavier vehicles adopted turbochargers first.

**Table 3:** *Turbocharger regressions*

| | *Dependent variable:* | | |
|---|---|---|---|
| | $\text{Turbo}_t$ | | |
| | (1) | (2) | (3) |
| Truck | | | 0.074 (0.192) |
| ln $\text{mpg}_{t-1}$ | $-0.022^{*}$ (0.013) | 0.030 (0.034) | 0.044 (0.036) |
| ln $\text{volume}_{t-1}$ | $-0.050^{***}$ (0.015) | 0.039 (0.032) | 0.050 (0.039) |
| ln $\text{hp/wt}_{t-1}$ | 0.017 (0.013) | $-0.039$ (0.029) | $-0.078^{**}$ (0.039) |
| $\text{Turbo}_{t-1}$ | $0.754^{***}$ (0.011) | $0.748^{***}$ (0.011) | $0.744^{***}$ (0.011) |
| t×ln $\text{mpg}_{t-1}$ | | $-0.004^{*}$ (0.002) | $-0.005^{**}$ (0.002) |
| t×ln $\text{volume}_{t-1}$ | | $-0.007^{***}$ (0.002) | $-0.007^{***}$ (0.002) |
| t×ln $\text{hp/wt}_{t-1}$ | | $0.004^{**}$ (0.002) | $0.005^{**}$ (0.002) |
| Truck×ln $\text{mpg}_{t-1}$ | | | 0.015 (0.028) |
| Truck×ln $\text{volume}_{t-1}$ | | | $-0.029$ (0.039) |
| Truck×ln $\text{hp/wt}_{t-1}$ | | | 0.017 (0.040) |
| Constant | $0.321^{***}$ (0.088) | $0.514^{***}$ (0.134) | $0.410^{**}$ (0.195) |
| Fixed effects | Make-year | Make-year | Make-year height |
| Observations | 4,960 | 4,960 | 4,960 |
| $R^2$ | 0.866 | 0.866 | 0.867 |
| Adjusted $R^2$ | 0.836 | 0.836 | 0.836 |
| F Statistic | $28.642^{***}$ | $28.671^{***}$ | $28.572^{***}$ |
| *Note:* | | | $^{*}$p<0.1; $^{**}$p<0.05; $^{***}$p<0.01 |