

Attribute Production and Biased Technical Change in Automobiles*

Asa Watten[†]

Soren Anderson[‡]

June 25, 2025

Abstract

Cars have gotten bigger and faster yet more fuel efficient in recent decades. Why? We estimate an equilibrium model of car attribute production using U.S. household microdata for 1995–2017 and structurally decompose attribute trends into underlying mechanisms. We find that technical change led to gains in all attributes. Rising gas prices boosted efficiency but were offset by surging demand for size and acceleration. Efficiency standards were largely ineffective. We show that using technology alone to meet tighter standards quadruples compliance costs, while half the efficiency gain from a fuel-saving technology subsidy is reallocated to other attributes in equilibrium.

Keywords: technical change, automobiles, fuel economy standards

JEL classification numbers: L50, L62, O30, Q40, R40

*Acknowledgements: We thank Kevin Bolon and Andrew Moskalik for engineering insights. We thank Stephanie Weber and seminar participants at Michigan State University, University of Alaska Anchorage, University of California Davis, University of Maryland, and National Bureau of Economic Research (NBER) Economics of Energy Use in Transportation Workshop (Spring 2023) for their helpful comments. Watten conducted some of this research while serving as Oak Ridge Institute for Science and Education Research (ORISE) Fellow at the U.S. Environmental Protection Agency. The views expressed herein are those of the authors and do not necessarily reflect the views of the Environmental Protection Agency or Electric Power Research Institute.

[†]Electric Power Research Institute, 3420 Hillview Ave, Palo Alto, CA 94304

[‡]Michigan State University, 486 West Circle Dr, East Lansing, MI 48824, sta@msu.edu

1 Introduction

Regulation of energy efficiency in cars, houses, and appliances is a core feature of U.S. energy and environmental policy. These goods last for decades and have big climate impacts, with buildings and transportation each accounting for 30% of U.S. greenhouse emissions (Environmental Protection Agency, 2020). But regulating energy efficiency is tricky, since it is physically co-determined with other attributes. This raises a persistent concern in both public discourse and the economics literature that efficiency standards will force trade-offs with other attributes, negatively impacting consumers, and that the efficiency gains from new technologies may be reallocated to other attributes (see Knittel, 2011; Whitefoot, Fowle and Skerlos, 2017; Leard, Linn and Zhou, 2023). For example, better insulation may lead to bigger homes, while falling costs for hybrid-electric engines may lead to bigger and faster cars. Predicting these policy outcomes is challenging, since doing so requires information on technology costs, attribute trade-offs, and consumer preferences.

Recent history highlights these concerns for the U.S. automobile industry. Rising oil prices and efficiency standards 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 tighter efficiency standards would force consumers into smaller, less powerful cars. Falling oil prices in the 1990s then coincided with increases in size and acceleration and a reduction in fuel economy, reinforcing this belief. Yet size and acceleration continued to increase in the 2000s with the adoption of new engine technologies, even under rising oil prices, and fuel economy actually increased—a pattern that continued into the 2010s. These latter trends suggest that new technologies may have facilitated the purchase of bigger, faster cars favored by consumers.

What explains these historical trends in car attributes, and what can a better understanding of the mechanisms tell us about energy-efficiency policy going forward? We answer these questions by developing and estimating an equilibrium model of car attribute production using U.S. data for 1995–2017. We use the model to structurally decompose trends in fuel economy, size, and acceleration. We show that underlying shifts in gas prices, technology, and preferences are all pivotal to understanding how the car fleet has evolved in recent decades. Steady technical change simultaneously made cars bigger, faster, and more efficient—and also cheaper. Rising gas prices drove large efficiency gains and down-sizing but were offset by surging consumer demand for size and acceleration that pushed strongly in the opposite direction. Efficiency standards were too weak to have had much impact. To our knowledge, we are the first paper to quantify the underlying mechanisms driving long-run trends in car attributes.

Our analysis proceeds in three steps. We begin by developing a theory of car attribute production that accounts for gas prices, technology, preferences, and policy. On the supply side, we model a competitive car industry with production costs that depend on car size, acceleration, and fuel economy; we show that the assumption of perfect competition is not necessary under plausible assumptions for markups. On the demand

side, we assume heterogeneous preferences for car attributes across consumer types, leading to an equilibrium mix of differentiated car models. We solve for a consumer’s optimal choice of fuel economy conditional on other car attributes, micro-founding Knittel (AER 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 gas prices, tighter fuel economy standards, and various forms of technical change.

Next, we estimate the model using household-level microdata, leveraging trends in the distribution of car attributes over time while controlling for changes in consumer preferences due to gas prices and shifting demographics. In addition, we study the diffusion of hybrid engines, turbo-charging, and other specific engine technologies, testing whether they are first adopted in bigger, more powerful cars as our theory predicts. Our data sources include three waves of the National Household Transportation Survey (2001, 2009, and 2017) covering cars originally purchased in 1995–2017, along with data on car attributes from the U.S. Environmental Protection Agency, and detailed data on car size from the Canadian government. Finally, we use our model to quantify how rising gas prices, technical change, and shifting preferences have interacted to produce observed trends in car attributes.

Beyond explaining trends in car attributes, our analysis yields five main findings. First, attribute-neutral technical change leads to bigger, faster, and more efficient cars. We estimate technical change equivalent to a 0.6% annual increase in all car attributes holding costs fixed from 1995 to 2017. We show that controlling for gas prices and other shifters of attribute demand is essential both to quantify technical change over time, and to recover structural cost and preference parameters on size, acceleration, and fuel economy.

Second, technical change often manifests as a reduction in the cost of adopting a discrete, attribute-boosting technology. We show theoretically that the earliest adopters are consumers already choosing bigger, faster, and more efficient cars. We confirm this prediction empirically for turbo-chargers, direct fuel injection, transmissions with 6+ gears, continuously variable transmissions, and hybrid-electric drivetrains. Our results imply that 46% of the direct efficiency gains from a new, fuel-saving technology are reallocated to size, speed, and cost savings. Thus, the environmental benefits of subsidies for hybrid-electric drivetrains and other energy-saving technologies may be much less than what policymakers intend. Yet the converse is also true: innovations targeting size and acceleration are largely used to improve fuel economy.

Third, consistent with engineering simulations, we find suggestive evidence of biased technical change favoring acceleration and size. In particular, a 10% gain in acceleration leads to a 6.1% fuel-economy loss in the late 1990s but only a 5.0% loss in the 2010s, holding costs and other attributes fixed. Meanwhile, a 10% gain in size leads to a 10% fuel-economy loss in the 1990s but only a 8.4% loss in the 2010s. We show theoretically that the falling opportunity cost of size and speed leads to bigger, faster, and less efficient cars and makes down-sizing a relatively less attractive approach for improving fuel economy.

Fourth, the effects of technical change, gas prices, and efficiency standards are all mediated by consumer preferences. We estimate that the price elasticity of demand is -0.36 for size and -0.31 for acceleration. To our knowledge, we are the first to estimate these parameters. When attribute demand is inelastic, new technologies are mainly channeled into improving fuel economy and reducing costs, and shifts in demand can lead to the super-sizing of cars. Our simulations show that cars in 2017 would have been 34% smaller had preferences for size and acceleration remained fixed at 1995 levels. Bigger cars cause death and injury to occupants of other vehicles, implying a large untaxed externality (Jacobsen, 2013; Anderson and Auffhammer, 2014; Bento, Gillingham and Roth, 2017). Ironically, demand for size may derive in part from an inefficient “arms race” for safety as drivers seek to protect themselves against other drivers (Li, 2012).

Finally, while inelastic demand mutes the incentive to down-size cars to comply with efficiency standards, the marginal cost of boosting energy efficiency via technology alone is also steeply increasing. Thus, we find that down-sizing is a relatively low-cost compliance strategy. Indeed, meeting a tighter standard via fuel-saving technology alone, while holding size and acceleration fixed, increases compliance cost by a factor of four. Conversely, meeting a tighter standard via down-sizing, while holding technology fixed, only slightly increases compliance cost.

A large economics literature models and estimates technical change in the aggregate economy or for broad industries. A smaller, mostly empirical literature studies technical change for differentiated products with a focus on energy efficiency. A seminal paper 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 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. Empirically, we show how controlling for energy prices can compensate for a lack of accurate cost data. Our application to cars suggests biased technical change favoring size and acceleration in recent decades.

Our paper is closely related to Knittel (2011), who also uses market data to study attribute trade-offs and technical change in the car industry. Dozens of papers cite Knittel (2011) to inform their understanding of these issues (for example: MacKenzie and Heywood 2012; Klier and Linn 2016; Whitefoot, Fowlie, and Skerlos 2017). We provide explicit micro-foundations for the empirical specification in Knittel (2011), which leads to a different interpretation of his results. In particular, we show that reduced-form regressions of fuel economy on car attributes and a time trend as in Knittel (2011) do not identify iso-cost relationships or the rate of technical change. However, we show how to recover these structural parameters by using household-level microdata, which allows us to control for state-level gas prices and instrument for size and acceleration using household-level demand shifters. We find that iso-cost curves are steeper and that technical change

has been faster than implied by the reduced-form estimates. Consistent with this finding, we show that new engine technologies are not applied to cars at random. Indeed, our theory predicts and our empirical results confirm that new engine technologies like those considered in Knittel (2011) are disproportionately applied to bigger, faster, and more fuel-efficient cars.

Many recent economic and environmental policies seek to direct technical change.¹ Previous literature studies how directed technical change affects the factors of production (Acemoglu, 2002a,b) or explores how taxes and subsidies stimulate technical change (Aghion et al., 2016; Cabel and Dechezleprêtre, 2016). Yet we know relatively little about directed technical change for differentiated products. Aghion et al. (2016) study directed technical change in the auto industry, showing that patents for electric, hybrid, and hydrogen engines increase during episodes of high gas prices. Like Aghion et al. (2016), we show theoretically that the incentive to innovate increases under higher gas prices. However, we differ from Aghion et al. (2016) by exploring how directed technical change interacts with gas prices and other factors to determine consumer-level car attributes in equilibrium.

Finally, our paper contributes to a large empirical literature studying how fuel economy, size, weight, and horsepower respond to gas prices and efficiency standards. See Anderson and Sallee (2016) for a conceptual model and review. Many papers use reduced-form methods, obscuring the underlying mechanisms. Other papers estimate structural models that take the set of car models and their attributes as given, while assuming heterogeneous but constant marginal willingness-to-pay for car attributes, making such models unsuitable for long-run analysis.² 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 long-run attribute trade-offs and technical change. Our theory parallels that of 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 average cost to introduce new car models, such that every consumer type gets a car model optimized to its own unique 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 global functional forms for costs and preferences. In addition, we explicitly model the role of gas prices, interest rates, miles traveled, depreciation, and falling costs for discrete engine technologies. Thus, we are able to characterize the full distribution of car models and how this distribution relates to several key observables. We show that this model provides a rich set of predictions that align closely with empirical evidence, and we argue that patterns of market power are unlikely to affect our qualitative results.

¹Examples include the CHIPS and Science Act of 2022, the Inflation Reduction Act, and the EPA Light- and Medium-Duty Proposed Standards for Model Years 2027 and Later (in the United States); the European Union Recovery Instrument and the European Green Deal (in Europe); and Acemoglu, Aghion, Bursztyn and Hemous (2012) and Aghion, Dechezleprêtre, Hemous, Martin and Van Reenen (2016) (among academic economists).

²Whitefoot and Skerlos (2012) and Lin and Linn (2023) make the choice of car attributes endogenous but do not consider long-run technical change.

The rest of this paper proceeds as follows. Section 2 discusses design trade-offs among car attributes. Section 3 develops a model of car attribute production, showing how consumer preferences, gas prices, production costs, and technical change interact to drive equilibrium choices of size, acceleration, and fuel economy. Sections 4 and 5 leverage this theory to estimate costs, preferences, and rates of technical change using household-level microdata on car attribute choices. Section 6 structurally decomposes trends in car attributes into the effects of gas prices, preferences, and technical change. Section 7 concludes.

2 The physical determinants of car attributes

Carmakers modify major car attributes—size, acceleration, and fuel economy—by adding or subtracting discrete technologies, improving drivetrain components, choosing different materials, altering engine size, and tuning the engine to achieve different goals. Table 1 presents five categories of individual technologies and design choices that: (1) improve acceleration (turbo-chargers and super-chargers); (2) improve fuel economy (engine stop-start); (3) improve both fuel economy and acceleration (engine efficiency, light-weighting, aerodynamics, and advanced transmissions)³; (4) improve acceleration at the expense of fuel economy (engine displacement and tuning); and (5) increase size and therefore weight, lowering both fuel economy and acceleration. We developed this table in discussion with automotive engineers at the Environmental Protection Agency who have extensive experience testing and modeling design choices and technology options on cars. Not shown are energy-consuming technologies that may reduce fuel economy, but which have little or no effect on power, such as stereo systems, air conditioning, and computers for self-driving technologies.

Engine displacement is a key design choice. A bigger engine increases power and acceleration at the expense of fuel economy, while a smaller engine does the opposite. Adding any of the other technologies directly increases manufacturing costs (with the exception of engine tuning). Thus, to change attributes while holding costs fixed entails adding some specific technology and removing others or decreasing engine size and adding some performance-enhancing technology.

Environmental Protection Agency (2019) observes that the relationship between engine displacement, horsepower (the engine’s propulsive force), and fuel economy has shifted substantially over time. A typical modern engine is much more powerful than an engine of the same displacement 40 years ago, while fuel consumption is relatively unchanged. Thus, engines of a given horsepower are both smaller and more efficient today than in the past. Additionally, a given increase in engine size boosts power more and erodes efficiency less today than in the past. To illustrate, appendix A reproduces EPA engineering simulations of representative drivetrain technologies over time and shows that engines of a given displacement achieve both higher fuel economy and better acceleration over time. Meanwhile, the slope in fuel economy vs.

³Transmission upgrades may also allow carmakers to trade off fuel economy and acceleration. However, manufacturers typically choose to improve both attributes. In hybrids, small electric motors are paired with small gasoline motors, but for a given gasoline engine adding an electric motor would improve both attributes.

Table 1: *Technologies used to change attributes*

	acceleration	fuel economy	size
turbo/superchargers	+	0	0
light-weighting	+	+	0
aerodynamics	+	+	0
CVTs/advanced trans.	+	+	0
variable valve timing	+	+	0
gas direct injection	+	+	0
electric/hybrid motors	+	+	0
engine stop-start	0	+	0
performance tuning	+	—	0
Engine displacement	+	—	0
<i>size (not tech.)</i>	—	—	+

Note: +/0/— indicates direction of change. The magnitudes for acceleration and fuel economy may be different or may be adjustable. Table produced from discussions with engineers at EPA.

acceleration space has flattened, such that a given increase in acceleration reduces fuel economy less today than in the past. This implies that the opportunity cost of size and power may have fallen over time, potentially contributing to the super-sizing of cars. We explore this possibility theoretically and empirically in the sections that follow.

3 Theory of attribute production

We develop a theoretical model of attribute production, considering the interaction of consumer preferences, technology, gas prices, and policy. We use the model to derive optimal attribute choices in equilibrium and 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 first adopted on bigger, faster, and more fuel-efficient cars, implying that unobserved costs will be correlated with car attributes.

3.1 Modeling optimal attribute choices

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

3.1.1 Model setup

We assume a finite number of consumer types, which have heterogeneous preferences for car attributes; we keep the notation tidy by suppressing consumer-type subscripts. We explicitly model three car attributes: acceleration (a), size (s), and gallons-per-mile (g). Acceleration (a good) is the rate at which a car can increase its velocity. Gallons-per-mile (a bad) 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

$$u(g, a, s) = v(a, s) + y - pmg - h(g, a, s), \quad (1)$$

where: $v(a, s)$ is utility derived from size and acceleration; y is income; pmg is lifetime fuel expenditures, which equals the price of gas (p) times fuel consumption (g) times lifetime miles (m); and $h(g, a, s)$ is the equilibrium hedonic price of a car with a given set of attributes. Note that gallons-per-mile enters as a purely financial trade-off between lifetime fuel expenditures and car price. We assume that consumers fully value the dollar benefits of fuel economy. More generally, we could allow consumers to undervalue fuel economy due to myopia, misinformation, inattention, credit constraints, or miscalculation. Recent empirical studies suggest a limited role for such factors (Busse, Knittel and Zettelmeyer, 2013; Allcott and Wozny, 2014; Sallee, West and Fan, 2016; Allcott and Knittel, 2019; Gillingham, Houde and van Benthem, 2021). We assume that lifetime miles is exogenous, given that demand for miles is quite inelastic (Anderson and Sallee, 2016); relaxing this assumption would not greatly change our conclusions. However, miles and gas prices may differ across consumer types, which 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: $c(g, a, s)$, which depends only on size, acceleration, and fuel economy. 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 be added to save fuel or boost acceleration but cost money. We assume that carmakers are price takers.⁴ We further assume free entry and zero average 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: $h(g, a, s) = c(g, a, s)$. Note that the assumption of zero average 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 $\sigma(s)$, which constrains average gallons-per-mile across each carmakers's fleet, possibly as a function of each car's size (s). We assume that this standard is implemented via a perfectly competitive credit-trading system with zero transaction costs, leading to equilibrium credit price τ . Thus, the fuel economy standard can be modeled as an implicit tax on each car's gallons-per-mile relative to the size-based standard: $\tau[g - \sigma(s)]$. Note that this tax becomes an implicit subsidy if the car is more efficient than the standard: $g < \sigma$.⁵ If the fuel economy standard fails to bind in equilibrium, then the

⁴Market power changes little so long as markups either are constant or scale with drivetrain costs. In this case, markups are captured by the cost function discussed later (see [C](#)). A problem may arise if there is some other pattern of markups. Grieco, Murry and Yurukoglu (2023) shows that average markups were stable both in aggregate and for specific body styles during our study period 1995–2018, following a period of decline in the 1980s and early 1990s. This alleviates concern that trending markups will substantially affect our estimates for technical change.

⁵Under 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. Note that before 2011 there was a separate standard for cars and light trucks but σ did

credit price is zero ($\tau = 0$), and these implicit taxes and subsidies vanish.

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

$$\max_{g,a,s} u(g, a, s) = v(a, s) + y - pmg - c(g, a, s) - \tau[g - \sigma(s)], \quad (2)$$

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

3.1.2 Functional-form assumptions

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

$$c(g, a, s) = ke^{-\theta} g^{-\alpha_g} a^{\alpha_a} s^{\alpha_s}, \quad (3)$$

where k is a cost shifter, θ is an index of attribute-neutral technical progress that scales production costs downward, and parameters $\alpha_a, \alpha_s, \alpha_g > 1$ relate percent changes in costs to percent changes in attributes.⁶ Given our functional form, costs are rising in acceleration and size ($c_a, c_s > 0$) and falling in gallons-per-mile ($c_g < 0$). Thus, making a car more desirable raises production costs. Costs are convex with respect to each of these attributes ($c_{aa}, c_{ss}, c_{gg} > 0$). Thus, it is increasingly costly to make a car faster, bigger, or more efficient. Finally, the cross-partials imply that making a car more efficient is costlier when the car is fast or big ($c_{ag}, c_{sg} < 0$) and likewise that making a car faster is costlier when the car is big ($c_{as} > 0$). To summarize succinctly: making a car better in any dimension is more costly when the car is already better in any dimension. Note that the elasticities of marginal cost with respect to acceleration, size, and fuel economy are $1 + \alpha_a$, $1 + \alpha_s$, and $1 + \alpha_g$.

On the demand side, we assume that utility is linear and decreasing in the inverse of size, the inverse of acceleration, and gallons-per-mile (all bads). Thus, our consumer problem becomes:

$$\max_{g,a,s} y + \tau\sigma(s) - (pm + \tau)g - \beta_a a^{-\mu_a} - \beta_s s^{-\mu_s} - ke^{-\theta} g^{-\alpha_g} a^{\alpha_a} s^{\alpha_s}, \quad (4)$$

where μ_a and μ_s capture the declining marginal benefits of acceleration and size, and β_a and β_s are preference shifters.⁷ Note that the marginal utility of size is $\beta_s s^{-(1+\mu_s)}$. Thus, the elasticity of demand for size is $-1/(1 + \mu_s)$, and likewise for acceleration. In our empirical application, we assume a uniform elasticity across consumers but allow for heterogeneity in the horizontal demand shifters (β_a, β_s).

not vary as a function of size. From 2011 onward, the level of the standard depended on the footprint of the vehicle. This attribute-based policy structure creates an indirect incentive to increase car size, in addition to improving fuel economy, as shown in Ito and Sallee (2018), Kellogg (2018), Kellogg (2020), and others. While we write the standard as a function of size, the true standard is a function of vehicle footprint, which can be thought of as a function of size on average.

⁶Note that we cannot have cost parameters ($\alpha_a, \alpha_s, \alpha_g > 0$) that sum to one, as in a standard Cobb-Douglas function, since this would imply decreasing marginal costs to improve attributes.

⁷Given that size and acceleration enter the Cobb-Douglas-like cost function with positive coefficients, this formulation ensures that the second-order conditions for utility maximization are satisfied. Note that we implicitly set $\mu_g = 1$ to reflect the fact that marginal reductions in fuel consumption (g) yield constant dollar savings; fuel consumption enters costs with a negative coefficient, thereby ensuring the proper relative curvature.

3.1.3 Necessary conditions and their implications

Maximization of equation (4) with respect to each of the car attributes implies:

$$a: \frac{\mu_a \beta_a a^{-\mu_a}}{a} = \frac{\alpha_a c(g, a, s)}{a} \quad (5)$$

$$s: \frac{\mu_s \beta_s s^{-\mu_s}}{s} = \frac{\alpha_s c(g, a, s)}{s} - \sigma_s \tau \quad (6)$$

$$g^{-1}: \frac{(pm + \tau)(g^{-1})^{-1}}{g^{-1}} = \frac{\alpha_g c(g, a, s)}{g^{-1}}, \quad (7)$$

where $\sigma_s \equiv \partial \sigma / \partial s$ is the change in the standard with respect to size. Note we maximize with respect to fuel economy (g^{-1}) in (7) to emphasize the symmetry among the three necessary conditions. Maximizing with respect to fuel consumption (g) yields an equivalent condition:

$$g: pm + \tau = \frac{\alpha_g c(g, a, s)}{g}. \quad (8)$$

These conditions say that the marginal benefit from improving a car attribute (left side) should equal the marginal cost (right side). At the optimum, the ratios of marginal benefits (slopes of the indifference curves in attribute space) should equal the ratio of marginal costs (slopes of iso-cost curves). We refer to the slopes of the iso-cost curves as the Marginal Rate of Technical Substitution in Attributes (MRTSA) and focus on opportunity cost of acceleration and size in terms of foregone fuel economy. Given our functional-form assumptions, these MRTSA values in elasticity form are given simply by ratios of the corresponding α parameters: $\partial \ln g / \partial \ln a|_{\Delta c=0} = \alpha_a / \alpha_g$ and $\partial \ln g / \partial \ln s|_{\Delta c=0} = \alpha_s / \alpha_g$. See appendix B for more on MRTSA, a graphical interpretation of the optimum, and verification of second-order sufficient conditions.

Solving equations (5)–(7) yields equilibrium attribute choices as a function of underlying preferences and technology parameters; see appendix B. Heterogeneity in preferences (β_a, β_s, pm) then gives rise to an array of car model offerings in equilibrium. Importantly, a consumer's optimal choice for a given car attribute will depend on her preferences for *all* attributes. Thus, consumers with stronger preferences in any one dimension will choose different car attributes in all three dimensions. For example, consumers with a stronger incentive to save fuel (larger pm) will choose slower, smaller, and more efficient cars.

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 strategy for identifying cost parameters 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 (7) using the specific functional form in (3), solve for gallons-per-mile (g), and then take logs to yield:

$$\ln g = \frac{1}{1 + \alpha_g} \ln k - \frac{\theta}{1 + \alpha_g} - \frac{1}{1 + \alpha_g} \ln(pm + \tau) + \frac{\alpha_a}{1 + \alpha_g} \ln a + \frac{\alpha_s}{1 + \alpha_g} \ln s \quad (9)$$

$$\approx \frac{1}{1 + \alpha_g} \ln k - \frac{\theta}{1 + \alpha_g} - \frac{1}{1 + \alpha_g} \left[\ln p + \ln m + \frac{\tau}{pm} \right] + \frac{\alpha_a}{1 + \alpha_g} \ln a + \frac{\alpha_s}{1 + \alpha_g} \ln s, \quad (10)$$

where the approximation in the second row follows from a first-order Taylor expansion of $\ln(pm + \tau)$ around $\tau = 0$. We model lifetime miles (m) as the present-discounted value, assuming a constant rate of time discounting (r), car scrappage (ρ), and decay in annual miles conditional on car survival (δ). Lifetime miles is therefore given by $m = m_0/(r + \delta + \rho)$, where m_0 is miles driven in the car's first year.

This equation suggests that we can recover cost parameters $(\alpha_g, \alpha_a, \alpha_s)$ and the annual rate of attribute-neutral technical change $(\partial\theta/\partial t)$ by regressing log fuel economy on log gas prices, log acceleration, log size, and a time trend. Intuitively, higher gas prices shift the demand for fuel economy, which, conditional on size and acceleration, identifies the cost parameter on fuel economy (α_g). Meanwhile, technical change, acceleration, and size shift the marginal cost of fuel economy. In equilibrium, the change in fuel economy induced by these shifts depends inversely on the cost parameter on fuel economy. Thus, having estimated this parameter, we are able to infer the other cost parameters (α_a and α_s) from the coefficients in this regression. Below, we show how to recover cost parameters from household-level microdata by regressing log fuel economy on log gas prices and log car attributes, while controlling for household-level demographics.

3.2.1 Revisiting Knittel (2011) through the lens of our model

How does our analysis differ from Knittel (2011)? Both analyses seek to estimate the slope of the iso-cost curve at a point in time, as well as the shift in this curve over time due to 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 (3) and rearranging yields:

$$\ln g = \frac{1}{\alpha_g} \ln k - \frac{\theta}{\alpha_g} + \frac{\alpha_a}{\alpha_g} \ln a + \frac{\alpha_s}{\alpha_g} \ln s - \frac{1}{\alpha_g} \ln c, \quad (11)$$

which is the form of Knittel's (2011) baseline regressions: log fuel economy on log attributes and year dummies to capture technical progress (θ).⁸ Thus, our functional form for costs provides explicit micro-foundations for Knittel (2011) and related papers.

The coefficient on log acceleration (α_a/α_g) captures the slope of the iso-cost curve: the percent change in g given a percent change in a holding cost fixed. Likewise for the coefficient on log size (α_s/α_g). Note that technology costs ($\ln c$) are unobserved. Thus, causal identification via OLS requires that technology cost be uncorrelated with size and acceleration. But consumers with a strong desire to save fuel will buy slower and smaller cars (see appendix B.4) and will buy more efficient cars conditional on these attributes, putting them on a higher iso-cost curve (equation 10). Thus, costs are negatively correlated with size and acceleration, biasing estimated iso-cost relationships toward zero.

Knittel (2011) concedes (pg. 3372) that this regression may not yield unbiased estimates for iso-cost curves

⁸Our equation has gallons-per-mile as the outcome, while Knittel's (2011) has miles-per-gallon. Likewise, we model acceleration as the horsepower-to-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.

at a point in time or literal shifts in iso-cost curves over time due to technical change. But he argues that his estimates remain valid for predicting how fuel economy would have evolved in the presence of technical change had car attributes remain fixed. In effect, he argues that he is estimating the reduced-form time trend $(\theta/(1 + \alpha_g))$ in our equation (9). Our model clarifies that, for this argument to hold, it is crucial to control for shifts in the demand for fuel economy over time $(pm + \tau)$. Knittel (2011) partially addresses this concern, estimating a time-series regression that explains the annual rate of technical progress as a function of gasoline prices and fuel economy standards.

Knittel (2011) also concedes (pg. 3386-7) that his estimates cannot be used to forecast fuel economy in equilibrium due to tighter standards but contends that they can be used to quantify what is technically feasible as a function of other car attributes. But this notion of technical feasibility implicitly forces consumers to meet tighter standards *only* by reducing size and power. A broader view of technical feasibility would permit changes in car attributes as well as costs. Indeed, we show below that the equilibrium response to efficiency-forcing mechanisms (gas prices and standards) is not simply to shuffle sideways along a given iso-cost curve, but rather to step up onto a higher iso-cost curve by adding fuel-saving technology, relaxing the trade-off with size and acceleration. Below, we quantify feasibility via the implicit tax (τ) needed to achieve a given fuel-economy target in equilibrium, holding preferences fixed.

3.3 Optimal size and acceleration conditional on other attributes

We now solve for optimal size and acceleration conditional on the other attributes, to motivate our empirical strategy for identifying demand parameters. Substitute for costs in necessary conditions (5) and (6) using the functional form in (3), solve for acceleration (a) and size (s), and then take logs to yield:

$$\ln a = \frac{1}{\alpha_a + \mu_a} \underbrace{(\theta - \alpha_s \ln s + \alpha_g \ln g)}_{\text{cost shifter for } a} + \frac{1}{\alpha_a + \mu_a} (\ln \mu_a \beta_a - \ln \alpha_a k), \quad (12)$$

$$\ln s = \frac{1}{\alpha_s + \mu_s} \underbrace{\left(\theta - \alpha_a \ln a + \alpha_g \ln g - \tau \frac{\sigma_s}{c_s} \right)}_{\text{cost shifter for } s} + \frac{1}{\alpha_s + \mu_s} (\ln \mu_s \beta_s - \ln \alpha_s k) \quad (13)$$

where note that we already know the relevant cost parameters (θ and α 's) based on (10) above.

This equation suggests that we can recover the demand parameter for acceleration (μ_a) by regressing log acceleration on a linear combination of log size and log fuel economy. Intuitively, these variables shift the marginal cost of acceleration. In equilibrium, the change in acceleration induced by this shift depends on the demand parameter (μ_a), along with the corresponding cost parameter (α_a). Thus, having estimated cost parameters in a first step, we can estimate the demand for acceleration. Likewise, we can recover the demand parameter for size (μ_s) by regressing log size on a linear combination of log acceleration and log fuel economy. We take this approach in our empirical analysis below.

3.4 Gasoline prices and efficiency standards

We now explore how attributes respond to higher gas prices. Equilibrium attributes are given by long and tedious expressions (see appendix B.4) but the comparative statics are quite intuitive. Assume that the shadow price on the fuel economy constraint is zero ($\tau = 0$), such that percent changes in gas prices correspond directly to percent changes in the demand for fuel economy. Differentiating log attributes with respect to log gas price then yields:

$$\frac{\partial \ln a}{\partial \ln p} = -\frac{\alpha_g \mu_s}{\psi} < 0 \quad (14)$$

$$\frac{\partial \ln s}{\partial \ln p} = -\frac{\alpha_g \mu_a}{\psi} < 0 \quad (15)$$

$$\frac{\partial \ln g^{-1}}{\partial \ln p} = \frac{\alpha_s \mu_a + \alpha_a \mu_s + \mu_s \mu_a}{\psi} > 0, \quad (16)$$

where $\psi \equiv \alpha_s \mu_a + \alpha_a \mu_s + \mu_s \mu_a + \alpha_g \mu_s \mu_a$ captures the overall insensitivity of equilibrium car attributes due to inelastic demand (large μ 's) or inelastic supply (large α 's). Note that car acceleration and size both decrease, while fuel economy increases.

Consider the percent change in size in (15). Note that larger α 's and μ 's imply more inelastic supply and demand for car attributes. So size decreases more when either the supply of fuel economy is inelastic (larger α_g), or the demand for acceleration is inelastic (larger μ_a), for a given ψ . The percent change in acceleration in (14) follows a similar pattern. Now consider the percent change in fuel economy in (16). The numerator implies that fuel economy increases more when the supply and demand for size and acceleration are more inelastic, again for a given ψ .

Note that a tighter fuel economy standard ($\tau > 0$) has qualitatively similar effects. To see this, observe that $\partial \ln(pm + \tau)/\partial \tau \approx 1/pm$ when evaluated at $\tau \approx 0$. Thus, just re-scale each of the effects in (14)–(16) by $1/pm$ to yield the percent changes, where recall that m varies across consumers. This re-scaling relates to the finding in Jacobsen, Knittel, Sallee and Van Benthem (2020) that standards imperfectly target car emissions since they fail to account for heterogeneity in lifetime miles across cars and consumers.

Now consider the effect of higher gas prices on fuel economy, conditional on other attributes. Differentiating equation (10) with respect to log gas prices yields:

$$\left. \frac{\partial \ln g^{-1}}{\partial \ln p} \right|_{a,s} = \frac{1}{1 + \alpha_g} > 0, \quad (17)$$

which shows that fuel economy increases conditional on other attributes. This result suggests that exogenous variation in fuel costs, controlling acceleration and size, can be used to identify the cost coefficient on gallons-per-mile, which is essential to identifying the other cost parameters (α_a, α_s) and the rate of technical change ($\partial \theta / \partial t$). We apply this insight in our empirical analysis below.

3.5 Technical change and car attributes

We now consider three different forms of technical change and explore their impacts on equilibrium attributes, along with their implications for empirical estimation.

3.5.1 Attribute-neutral technical change

We model attribute-neutral technical change via an increase in θ . Recall that the expressions for equilibrium attributes are long and tedious (see appendix B.4). Yet it is easy to show that a marginal increase in θ leads to an improvement in all three attributes:

$$\frac{\partial \ln a}{\partial \theta} = \frac{\mu_s}{\psi} > 0 \quad (18)$$

$$\frac{\partial \ln s}{\partial \theta} = \frac{\mu_a}{\psi} > 0 \quad (19)$$

$$\frac{\partial \ln g^{-1}}{\partial \theta} = \frac{\mu_s \mu_a}{\psi} > 0, \quad (20)$$

where recall $\psi \equiv \alpha_s \mu_a + \alpha_a \mu_s + \mu_s \mu_a + \alpha_g \mu_s \mu_a$. If $\mu_a = \mu_s = 1$, we have an equal proportional improvement in all attributes. However, if demand for acceleration and size are sufficiently inelastic ($\mu_a, \mu_s > 1$), then the proportional increase in fuel economy is larger. Meanwhile, if demand for size is more inelastic than demand for acceleration ($\mu_s > \mu_a$), then acceleration improves more than size. These changes are all inversely proportional to the overall insensitivity of equilibrium car attributes (ψ).

Now consider how a marginal increase in parameter θ affects fuel economy conditional on acceleration and size in equation (10). Differentiating with respect to θ yields:

$$\left. \frac{\partial \ln g^{-1}}{\partial \theta} \right|_{a,s} = \frac{1}{1 + \alpha_g} > 0. \quad (21)$$

Thus, the shift in log fuel economy is the rate of technical change itself, scaled by the elasticity of marginal cost with respect to fuel economy ($1 + \alpha_g$). Knittel (2011) interprets this shift directly as technical change, but this is incorrect. In equilibrium, consumers allocate some technical change to cost savings. Thus, the equilibrium shift understates technical change. The correct rate is given by differentiating equation (11) with respect to θ : $1/\alpha_g$. Alternatively, we might wish to know the common change in all attributes that can be obtained per unit increase in θ , holding cost fixed. This rate is given by totally differentiating log cost in equation (3), equating to zero, and solving for the common change in log attributes: $1/(\alpha_a + \alpha_s + \alpha_g)$.

3.5.2 Attribute-biased technical change

Now consider attribute-biased technical change. In our empirical application, we model biased technical change via shifting cost parameters (α 's), which causes iso-cost curves to rotate but which also alters the slope and position of the entire cost function. Here, to isolate the impact of attribute bias, we consider a

linear add-on to our original cost function, which generates a pure rotation in the iso-cost curve:

$$\tilde{c}(g, a, s) = c(g, a, s) - \eta(g - g_0), \quad (22)$$

where $\eta = 0$ implies our original cost function (c) and $\eta > 0$ implies a modified function ($\tilde{c}$) in which costs are rotated around a reference-level gallons-per-mile (g_0).

How does this technical change affect attribute choices? Replace c with modified costs $\tilde{c}$ in equation (4) and note that an increase in the technical change parameter (η) is mathematically equivalent to a decrease in the efficiency standard's credit price (τ), which we considered in the previous sub-section. Thus, biased technical change of this form leads to bigger, faster, and less fuel-efficient cars. Intuitively, biased technical change lowers the opportunity cost of acceleration and size in terms of foregone fuel economy. So fuel economy goes down and these other attributes go up (see appendix B.5 for a graphical interpretation).

3.5.3 Discrete technology adoption and induced innovation

Carmakers periodically invent new technologies that can be added to cars to improve one or more attributes. For example, turbochargers force additional compressed air into the combustion chamber of an engine, boosting acceleration. 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 the carburetor), and electric or hybrid-electric drivetrains, all of which have seen falling costs and increased adoption. How does this form of technical change manifest in the car market?

We consider a discrete technology targeting acceleration—to use the turbocharger example—that suppliers may add to a car at a cost. We make two assumptions. First, installation cost (κ) is the same for all cars regardless of attributes. Second, the technology yields a proportional improvement (ω) in acceleration. Thus, cars with a higher baseline level of acceleration will see a larger absolute improvement. These assumptions will tend to hold for any technology whose performance gains are proportional to engine size or weight, but whose adoption includes a large fixed-cost component; turbochargers are a clear example.⁹

We model acceleration as perceived by the consumer to be $a(1 + \omega) > a$, where a is the acceleration that enters the cost function and $\omega > 0$. Thus, conditional on adoption, and ignoring installation cost, the consumer experiences utility: $u(\omega) = y + \tau\sigma - \beta_s s^{-\mu_s} - \beta_a ((1 + \omega)a)^{-\mu_a} - pmg - c(g, a, s)$, where we have simply scaled perceived acceleration by $1 + \omega$. By the envelope theorem, the utility gain from a marginal

⁹The basic physics of engine power implies a roughly proportional relationship: power = effective pressure $\times$ stroke length $\times$ piston area $\times$ engine speed in rpm. A turbocharger boosts the effective pressure, which interacts with engine size to deliver power. See here: https://www.apexinnovations.co.in/pdf/Engine_theory_and_calculations.pdf. Meanwhile, turbochargers must be matched to engine size for optimal performance. Our casual inspection of prices for various aftermarket turbocharger kits suggests that there is little correlation with engine size.

increase in ω at the optimum is given by: $\partial u / \partial \omega = (1 + \omega)^{-(1+\mu_a)} \mu_a \beta_a a^{-\mu_a}$. Thus, the net benefit from adopting the discrete technology is given by the following approximation:

$$\begin{aligned} u(\omega > 0) - u(\omega = 0) - \kappa &\approx \omega \mu_a \beta_a a^{-\mu_a} - \kappa \\ &= \omega \alpha_a c^* - \kappa, \end{aligned} \tag{23}$$

where c^* is car cost prior to adopting the discrete technology. The first term in the top line is ω times the marginal utility gain evaluated at $\omega = 0$, and equality in the second line follows from replacing the marginal benefit of acceleration with its marginal cost using necessary condition (5).

Thus, for any $\kappa > 0$, consumers choosing more expensive cars are most likely to adopt the discrete technology. This prediction holds regardless of the particular mix of preferences that leads a consumer to choose an expensive car in the first place. Perhaps the consumer has a need for speed. Or craves size. Or is a die-hard conservationist. Whatever the reasons, she will opt to install the discrete technology if her car is sufficiently costly. We confirm this prediction in our empirical analysis below, showing that cars with costlier drivetrains are more likely to feature turboschargers and other advanced engine technologies.

What happens to equilibrium attributes upon adoption? Note that adopting the acceleration-boosting technology is mathematically equivalent to replacing log acceleration preferences ($\ln \beta_a$) with enhanced preferences ($\ln \beta_a - \mu_a \ln(1 + \omega)$) in our expressions for equilibrium attributes (see appendix B.4). Differentiating attributes with respect to the boost ω around $\omega = 0$ then yields:

$$\left. \frac{\partial \ln(1 + \omega)a}{\partial \omega} \right|_{\omega=0} = \frac{\alpha_a \mu_s}{\psi} > 0 \tag{24}$$

$$\left. \frac{\partial \ln s}{\partial \omega} \right|_{\omega=0} = \frac{\alpha_a \mu_a}{\psi} > 0 \tag{25}$$

$$\left. \frac{\partial \ln g^{-1}}{\partial \omega} \right|_{\omega=0} = \frac{\alpha_a \mu_s \mu_a}{\psi} > 0, \tag{26}$$

where the first row is the change in log acceleration as perceived by the consumer.¹⁰ Strikingly, consumers that adopt the discrete technology experience gains in all attributes, not just acceleration. Further, the gains are equivalent to attribute-neutral technical change in equations (18)–(20), scaled by the cost parameter on acceleration (α_a).¹¹ Thus, a falling cost for a discrete technology is one way of motivating attribute-neutral technical change; for this reason we do not follow Knittel (2011) in controlling for turbochargers or the like in our regressions, lest we control-away the very mechanism of technical change we wish to estimate.

¹⁰Note that (i) $\partial \ln(a(1 + \omega)) / \partial \omega = \partial \ln a / \partial \omega + 1$ when $\omega = 0$ and (ii) $\partial \ln a / \partial \omega = -(\alpha_s + \alpha_g \mu_s + \mu_s) \mu_a / \psi$ when $\omega = 0$ via equilibrium log acceleration in the appendix. Replacing $\partial \ln(a(1 + \omega)) / \partial \omega$ in (i) using (ii) and simplifying then yields the expression above.

¹¹The same results obtain when we instead assume that the discrete technology proportionally reduces the cost of supplying the targeted attribute: $c(s, a(1 - \omega), g) = (1 - \omega)^{\alpha_a} c(g, a, s)$, where acceleration perceived by the consumer remains fixed but decreases by fraction ω in the cost function. 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 choosing expensive cars, with the same equilibrium effect as attribute-neutral technical change.

Our analysis implies that consumers with stronger preferences in any dimension are more likely to adopt a technology targeting any attribute. This result helps us understand the diffusion of new engine technologies over time. According to our model, these technologies are first adopted by consumers choosing the costliest cars. Only as installation costs (κ) gradually fall will others adopt, mirroring the stylized fact that new technologies typically appear first in luxury segments. This result also has profound implications for subsidies targeting specific fuel-saving technologies. According to our model, by lowering net installation costs (κ), a subsidy first spurs adoption among cars with costlier drivetrains. The initial efficiency gain is then reallocated, in part, to making these cars bigger and faster. Thus, in equilibrium, fuel savings may be much lower than what policymakers expect, if they naively assume that other attributes remain fixed.

Finally, consider the incentive to innovate new engine technologies. Condition (26) for discrete technology adoption shows that the net payoff increases with car costs. Meanwhile, higher gas prices and efficiency standards both cause car costs to increase.¹² Thus, the payoff to developing a new attribute-boosting technology rises with higher gas prices or tighter efficiency standards, as in Aghion et al. (2016).

4 Data and descriptive statistics

We merge household-level microdata on car choices with model-level data on car attributes, along with state-level data on gasoline prices. We supplement these data with information on the stringency of efficiency standards, interest rates on new car loans, car durability, and discrete technology adoption.

4.1 Data sources

Our car choice data come from the National Household Transportation Survey (NHTS) public use microdata waves for 2001, 2009, and 2017 (Federal Highway Administration, 2001–2017). These data record the make, model, and model year of cars owned by households at the time of the survey, along with the car’s odometer reading. We calculate car age as survey year minus model year. These data also record key household characteristics, including income, # household members by age and sex, owner-occupied, # vehicles, urban vs. rural (census block group), and population density (census tract).¹³ We are unable to differentiate between cars purchased new vs. used. Thus, the parameters we estimate implicitly reflect the combined preferences of original car buyers and current owners, since original buyers channel the preferences of future owners via resale value. We apply NHTS sampling weights throughout our analysis.

We match these car choices to data on car attributes from several sources. Our data on car size come from Transport Canada (a federal agency) as compiled by the Canadian Association of Road Safety Profes-

¹²This result follows from taking logs on both sides of equation (8) and differentiating with respect to $\ln(pm + \tau)$: $\partial \ln c / \partial \ln(pm + \tau) = 1 - \partial \ln g^{-1} / \partial \ln(pm + \tau) > 0$, given inelastic response of fuel economy to gas prices in equation (16).

¹³We define four categories of income that are consistent across the NHTS waves in 2017 dollars: \$0–34k, \$35–49k, \$50–74k, and \$75k+. We redefine urban vs. rural categories for 2001 and 2017 to be consistent with the coarser 2009 categories: town and country, second city, suburban, and urban.

sionals. These are the same data offered by the National Highway Traffic Safety Administration’s Vehicle Identification Number Decoder tool, so we are confident that they apply to U.S. vehicles. However, these data constrain our study period to start in 1995 after NAFTA. With these data, we measure the volume of a box bounding the car’s useful space: length $\times$ width $\times$ height. Length is the distance between the center of the front windshield and rear windshield (or tailgate, in the case of a pickup truck). Width is the maximum distance from left to right. Height is the maximum distance between the ground and the roof of the passenger or cargo compartment. Note that this measure of size is more expansive than narrow measures of passenger space or cargo space or even interior space, which may be influenced by the size and configuration of seats, dashboards, consoles, and other interior features. Our size data are measured at the trim level. Thus, we calculate the national sales-weighted mean size for each make, model, and model year using sales data from the Environmental Protection Agency prior to matching to NHTS data.

Our data on fuel economy and acceleration come from the Environmental Protection Agency (EPA). EPA measures gallons-per-mile in a simulator designed to mimic both city and highway driving. EPA then calculates a weighted average assuming 55% city and 45% highway driving.¹⁴ We proxy for acceleration via the horsepower-to-weight ratio, since the rate of velocity increase is approximately proportional to this ratio (see MacKenzie and Heywood 2012).¹⁵ Weight is measured as curb weight, which includes a full tank of fuel and standard equipment. As with size, our data on fuel economy and acceleration are measured at the trim level. Thus, we again calculate the national sales-weighted mean fuel economy and acceleration for each make, model, and model year prior to matching to our NHTS data.¹⁶

We match car choices to state-average retail gasoline prices from the U.S. Energy Information Administration (EIA). These data are measured annually and include federal, state, and local taxes. We match cars to gas prices by model year and state. We are unable to identify cars that were purchased in one state and then moved across state lines, either with their original owners or through resale. However, our results are robust to excluding cars more than 8–9 years old from the analysis. Of course, what matters for car choice is the expected future price of gasoline over a car’s lifetime. Reassuringly, Anderson, Kellogg and Sallee (2013) find that consumer forecasts move one-for-one with current gasoline prices in monthly time-series data and nearly one-for-one in monthly panel data for states.¹⁷

¹⁴EPA revised the measure for car labeling requirements at dealerships starting in 2012, based on a more realistic simulation. We use the traditional measure, which the EPA and NHTSA continue to use for regulatory purposes, and which is consistent throughout our entire sample period.

¹⁵Acceleration is often measured as time to accelerate from 0 mph to 60 mph. Our concept of acceleration as the “rate of velocity increase” is inversely related to this 0-60 measure. MacKenzie and Heywood (2012) shows that the relationship between horsepower-to-weight and 0–60 time changed little from 1995 to 2010 (the end of their study period).

¹⁶For fuel economy, we first calculate the sales-weighted mean of gallons-per-mile across trims and then take the reciprocal. For acceleration, we directly calculate the sales-weighted mean of the horsepower-to-weight ratio across trims.

¹⁷In particular, Anderson et al. (2013) show that a \$1 increase in the national gasoline price over one month is associated with a \$0.99 increase in the mean five-year forecast, while a \$1 increase in the state-level gasoline price is associated with a \$0.87 increase in the mean five-year forecast (controlling for both state and month fixed effects).

We collect information on the shadow cost of fuel economy standards (τ) from several sources that have estimated this value. For 1996–2006, we rely on Anderson and Sallee (2011) who estimate shadow costs for carmakers that exploited an option in the standards giving extra compliance credits to flexible-fuel vehicles. Compliance credits were not tradable during this period and so the shadow cost varies across firms. For 2007–2011, the standards appear not to have not been binding (Yeh, Burtraw, Sterner and Greene 2021). For 2012–2017, we rely on Yeh et al. (2021). NHTSA and EPA allowed compliance credits to be traded between firms starting in 2012, such that the shadow cost can be inferred from credit trading.

We measure the real annual interest rate on 48-month new car loans from The Federal Reserve Bank of Saint Louis (U.S. Federal Reserve Bank 2019). We calibrate the annual scrap probability as $\rho = 0.05$ based on data in NHTSA (2006).¹⁸ We estimate first-year miles and miles decay in our NHTS microdata data via the following OLS regression: $\ln \text{miles} = a - \delta \cdot \text{age} + \epsilon$. The coefficient on age yields $\delta = 0.036$, while the fitted value when $\text{age} = 0$ yields $m_0 = \exp(a)$. We use this information to calculate and control for the trend in interest rates plus durability and miles decay ($r + \delta + \rho$) that would otherwise contaminate our estimate of technical change, and we calculate present-value lifetime miles as $m = m_0 / (r + \delta + \rho)$.

Our data on discrete technologies (e.g., turbochargers) come from Wards Automotive for 1995–2019 and record, for a given car production model, the presence of a turbocharger, gasoline direct injection (GDI), automatic transmission speed, continuously variable transmission (CVT), or a gas-electric hybrid engine. We match this information to model-level data on car attributes using the same data sources and matching algorithm as above. Note that our unit of observation is a single production model offered by a carmaker, rather than an individual household’s car choice. We match to EPA data to capture sales weights.


4.2 Correlations in car attributes

Figure 1 shows the distributions of fuel economy, acceleration, and size in our data for model years 1999–2001 (black) and model years 2015–2017 (red). The top panel presents a scatter diagram of fuel economy vs. acceleration across car models in both time periods. Each dot is a different car model. These dots are not equally represented in our household-level choice data, since some car models are more popular than others. However, the histograms along the axes show the full household-weighted distributions of car attributes. The bottom panel presents similar information for fuel economy vs. size.

Echoing Knittel (2011), this figure reveals two striking facts. First, at any given time, there is a strong negative correlation between fuel economy and size and acceleration (inspecting the scatter diagrams). This pattern is consistent with strong technical trade-offs. However, our theory shows that this equilibrium

¹⁸The scrap rate varies with car age: lower than 0.05 when a car is young, higher than 0.05 during middle age, and then low again in the golden years. These patterns presumably reflect a complex mix of heterogeneity across cars and drivers, e.g. beloved old cars that are barely driven. These dynamics are not central to our analysis. Thus, we pick $\delta = 0.05$, which roughly corresponds to the weighted-average scrap rate, with weights given by discount factor $\exp(-0.1age)$ to reflect both time discounting and the annual decay in miles.

Figure 1: *New car attributes in 1999–2001 and 2015–2017*



Note: This figure plots the distribution of new car fuel economy (miles per gallon), acceleration (as proxied by the horsepower per weight-in-pounds ratio), and volume (cubic meters based on $l \times w \times h$ of car). Attributes for model years 1999–2001 are shown in black while attributes for 2015–2017 are shown in red. Scatter diagrams show the correlation between fuel economy and acceleration or size in 1999–2000 (black o's) and 2015–2017 (red x's); each dot plots the sales-weighted average attributes for all trims within a given model and model year. Histograms show the sales-weighted distribution of the attributes for 1999–2001 (shaded red) and 2015–2017 (black outline).

relationship also depends on the cost of adding fuel-saving technology. Second, cars have gotten bigger, faster, and less efficient over time (comparing the red and black histograms). However, cars of a given size and speed have become more efficient over time (comparing the red and black scatter diagrams). These shifts are consistent with rising demand for size and acceleration, along with improved car technology that has allowed even the biggest and fastest cars to become more efficient—interpretations emphasized by Knittel (2011). However, our theory shows that these equilibrium shifts also depend on rising gas prices and falling interest rates, which make fuel-saving technologies more valuable.

4.3 Trends in gas prices, interest rates, and efficiency standards

Figure 2 shows trends in monthly real gasoline prices for three large and differentiated states: California, Michigan, and Texas. The top panel shows that state-level gas prices largely follow a common trend driven by global crude oil prices, with persistent gaps across states. However, the bottom panel shows that there is substantial variation in gas prices, even after removing state and month fixed effects, with California’s price rising relative to Michigan and Texas over time. This residual variation is largely driven by supply-side factors, such as changes in state fuel taxes, mergers, and differential environmental regulation of fuels and refineries (Chouinard and Perloff, 2007), as well as differential pass-through of oil prices into retail gasoline prices (Kilian and Zhou, 2024).

Real interest rates on 48-month new car loans have fallen by a factor of two over our sample period (see appendix D). Thus, to infer technical change from market data, it is essential to control for trends in the discount rate on miles ($r_t + \delta + \rho$). We do not attempt to measure or control for trends in durability via lower scrappage rates (ρ) or weaker decay in miles (δ) since improved durability could be construed as a reduction in drivetrain costs due to technical change. While NHTSA (2006) estimates that scrap rates (ρ) have fallen over time, we find no evidence that first-year miles (m_0) or the decay rate (δ) varies across NHTS waves. Thus, we control only for time-series variation driven by interest rates.

Meanwhile, estimates from the literature suggest that shadow prices on efficiency standards have fluctuated over time (see appendix D). Pre-2007 there was no credit trading. Thus, the shadow cost varies by firm and car vs. light truck fleet. Our empirical analysis considers three values for the shadow cost during this period: (1) the \$55 noncompliance penalty; (2) the mean among constrained firms for which Anderson and Sallee (2011) estimate the shadow cost based on flexible-fuel credits; and (3) zero. Post-2007 our analysis uses a uniform shadow cost based on compliance credit trading (Yeh et al., 2021). We show that the resulting adjustment to log gas prices (τ/pm) is always less than 0.1 (see appendix D). Thus, shadow prices are an order magnitude lower than the direct dollar benefits from improved fuel economy.

Figure 2: *Real gasoline prices*



Note: The top panel shows real monthly retail gasoline prices in California, Michigan, and Texas from 1995–2017. The bottom panel shows the residual variation in gasoline prices for these three states, following an OLS regression of gasoline prices on state and month fixed effects for all 50 states and the District of Columbia during our sample period.

5 Empirical estimation and results

We first present our main results, which estimate key parameters of our model by relating fuel economy to size, acceleration, and gasoline prices. We then present ancillary evidence on the mechanisms of technical change, estimating the correlation between discrete technology adoption and car attributes.

5.1 Fuel economy conditional on size and acceleration

We start by estimating a version of equation (10) above:

$$\ln mpg_{ijt} = f_{\theta}(t) + \gamma_p(\ln p_{ist} - \ln(r_t + \delta + \rho) + \tau_t/\bar{m}p_{ijt}) + \gamma_a \ln hpwt_{ijt} + \gamma_s \ln volume_{ijt} + \phi_j + \varepsilon_{ijt}, \quad (27)$$

where: $\ln mpg_{ijt}$ is log fuel economy, $\ln hpwt_{ijt}$ is log horsepower-to-weight, and $\ln volume_{ijt}$ is log volume for car i in state j bought in year t ; $\ln p_{ijt} - \ln(r_t + \delta + \rho) + \tau/\bar{m}p_{ijt}$ is log gas price in state j and year t with adjustments for the interest rate and shadow cost; $f_{\theta}(t)$ captures attribute-neutral technical change; ε_{ijt} is an error term; and the γ 's are coefficients to be estimated. We include state fixed effects (ϕ_j) in all models. We measure trends in technology (f_{θ}) via year fixed effects; we also try a linear time trend ($\gamma_{\theta}t$). Our gas price variable always subtracts the log discount factor ($r_t + \delta + \rho$). This does not affect the coefficient on gas prices but ensures that the year fixed effects are uncontaminated by the trend in interest rates. In some models we further adjust for the trend in the shadow cost on fuel economy standards ($\tau_t/\bar{m}p_{ijt}$). We do not observe anticipated lifetime miles for individual car buyers. However, in some models we control for log odometer and car age, along with their squares and interaction, plus categories of household income, # adults, home ownership, # vehicles, urban vs. rural, and population density.

Our goal is to identify costs: causal effects of gas prices, car size, and acceleration on fuel economy all else equal. But in practice the error term (ϵ) may include unobserved shifts in efficiency demand (via m) or drivetrain costs (k) that would lead to a re-optimization of all attributes, causing bias in OLS estimates. For example, pickups typically have lower gear ratios for pulling heavy trailers, diminishing both fuel economy and acceleration. Thus, towing capacity will be positively correlated with size and negatively correlated with fuel economy and acceleration, leading to biased estimates. We address this specific concern by including a pickup dummy. But there may be other car attributes, such as advanced safety features, that we do not observe. To identify costs, we therefore require exogenous shifts in the demand for car attributes.

We identify the coefficient on log gas prices using variation across states and over time. Intuitively, gas prices directly shift the demand for fuel economy. Thus, the change in log fuel economy conditional on other attributes ($\gamma_p \equiv 1/(1 + \alpha_g)$) is inversely proportional to the elasticity of marginal costs ($1 + \alpha_g$). Our identification assumption is that unobserved state trends in the demand for fuel economy are uncorrelated with state trends in gas prices. We do not observe exactly when a car is purchased; cars of a given model year could be sold early in the calendar year or even partway through the prior year. Further, car production may respond gradually to gas prices. Thus, we model the log of the three-year moving average of state gas prices; our results are robust to including current and four years of lagged gas prices separately.

We identify costs for size and acceleration by instrumenting for these variables via the total # of household members and the share of members (age 16+) that are male. Bigger households demand bigger cars, shifting the marginal cost of fuel economy. Thus, the equilibrium change in log fuel economy ($\gamma_s \equiv \alpha_s/(1 + \alpha_g)$) is

the underlying shift in log marginal costs (α_s) divided by the elasticity of the marginal cost curve ($1 + \alpha_g$). Likewise, men stereotypically demand faster and more powerful cars, leading to a change in log fuel economy ($\gamma_s \equiv \alpha_s/(1 + \alpha_g)$) equal to the shift in log marginal costs (α_a) divided by the elasticity ($1 + \alpha_g$). Our identification assumption is that household size and maleness only correlate with log fuel economy via their effect on car size and acceleration. Consistent with our two sets of identification assumptions, we show that our results are robust to controlling for a rich set of household-level observables, including odometer, car age, income, # children, home ownership, and urban structure.

Having estimated the coefficients on log gasoline prices ($\gamma_p \equiv 1/(1 + \alpha_g)$), log acceleration ($\gamma_s \equiv -\alpha_s/(1 + \alpha_g)$), and log size ($\gamma_s \equiv -\alpha_s/(1 + \alpha_g)$) in equation (27), we can then identify the underlying cost parameters on fuel economy ($\alpha_g \equiv (1 - \gamma_p)/\gamma_p$), acceleration ($\alpha_a \equiv -\gamma_a/\gamma_p$) and size ($\alpha_s \equiv -\gamma_s/\gamma_p$), along with the ratios capturing attribute trade-offs along iso-cost curves (MRTSA values). The coefficient on log gas prices is also key to recovering trends in the index of attribute-neutral technical change, whether we capture this index via a linear time trend or year dummies ($\Delta\theta \equiv \Delta f_\theta(t)/\gamma_p$).

5.1.1 Main estimation results

Table 2 presents our results. Column (1) parallels Knittel’s (2011) approach. We regress log fuel economy on log attributes, along with a pickup dummy, and capture technical change via year dummies. We also control for state dummies, as in all of our regressions. We estimate this model via OLS. Consistent with figure 1, we find that cars with 1% higher acceleration and size in a given year respectively have 0.495% and 0.467% lower fuel economy. Pickup trucks have 6.1% lower fuel economy conditional on these attributes. Figure 3 plots the coefficients on the year dummies and shows large increases in fuel economy conditional on attributes (solid gray line). Following Knittel, column (2) then replaces the year dummies with a linear time trend; the coefficient implies a 1.6% annual increase in fuel economy conditional on size and acceleration.

Column (3) presents the results of our basic model, derived from theory, which adds log gas price to regression (1). We again estimate the model via OLS. The coefficient on log gas price implies that a 1% increase in gas prices leads to a just-insignificant $\gamma_p = 0.089\%$ increase in fuel economy conditional on other car attributes. The coefficients on acceleration, size, and the pickup dummy change little. Together, these coefficients imply that fuel economy decreases 0.543% and 0.512% for every 1% increase in acceleration and size, holding drivetrain costs fixed (see the iso-cost elasticities at the bottom of the table). However, these OLS estimates are likely biased, since they fail to identify costs via demand shifts.

Column (4) therefore estimates our model using two-stage least squares (2SLS), instrumenting for size and acceleration via # household members and male share. These instruments strongly predict size and acceleration, as indicated by standard F-tests and Sanderson-Windmeijer (SW) conditional F-tests (see

Table 2: *Fuel economy conditional on attributes*


	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	OLS	OLS	OLS	2SLS	2SLS	2SLS	2SLS
Annual trend (γ_θ)		0.016*** (0.000)					
Log gas price (γ_p)			0.089 (0.050)	0.095* (0.040)	0.097* (0.039)	0.093* (0.042)	0.103* (0.047)
Log HP/weight (γ_a)	-0.495*** (0.011)	-0.503*** (0.010)	-0.495*** (0.011)	-0.652*** (0.105)	-0.657*** (0.102)	-0.641*** (0.124)	-0.643*** (0.124)
Log volume (γ_s)	-0.467*** (0.008)	-0.479*** (0.007)	-0.467*** (0.008)	-0.672*** (0.070)	-0.654*** (0.072)	-0.605*** (0.070)	-0.607*** (0.070)
Pickup	-0.061*** (0.006)	-0.055*** (0.007)	-0.061*** (0.006)	-0.019 (0.010)	-0.022* (0.010)	-0.028** (0.009)	-0.028** (0.009)
Year FEs	X		X	X	X	X	X
State FEs	X	X	X	X	X	X	X
Odometer x Age					X	X	X
Demographics						X	X
CAFE							X
$\frac{\partial \ln \text{mpg}}{\partial \ln \text{HP/Wt}} \Big _{\Delta c=0} \left(\frac{\alpha_a}{\alpha_g} = \frac{\gamma_a}{1-\gamma_g} \right)$			-0.543 (0.039)	-0.720 (0.116)	-0.728 (0.113)	-0.707 (0.140)	-0.716 (0.143)
$\frac{\partial \ln \text{mpg}}{\partial \ln \text{volume}} \Big _{\Delta c=0} \left(\frac{\alpha_s}{\alpha_g} = \frac{\gamma_s}{1-\gamma_g} \right)$			-0.512 (0.034)	-0.743 (0.092)	-0.724 (0.092)	-0.668 (0.089)	-0.677 (0.093)
SW F-stat (HP/Wt)				35.768	36.141	21.970	22.006
P-value				(0.000)	(0.000)	(0.000)	(0.000)
SW F-stat (volume)				43.208	46.414	37.854	37.781
P-value				(0.000)	(0.000)	(0.000)	(0.000)

Note: This table presents coefficient estimates from equation (27). $N = 245,277$. Columns (4)–(7) instrument for log HP/Wt and log volume with male share and # household members. Odometer $\times$ age is log odometer (= 0 if missing), car age, and missing odometer dummy, plus their squares and interactions. Demographic controls is income, # adults, home ownership, # vehicles, urban vs. rural, and population density. Column (7) adds the shadow price adjustment (τ/pm) to log gas prices and instruments for the adjusted variable via log gas prices. The table also reports implied iso-cost elasticities. Standard errors robust to state clustering are given in parentheses for both the main coefficients and iso-cost elasticities. Note *, **, and *** indicate statistical significance at the 5%, 1%, and 0.1% levels respectively. The bottom of the table reports Sanderson-Windmeijer (SW) conditional F-tests for weak instruments robust to state clustering.

bottom of table).¹⁹ The coefficient on gas prices increases to $\gamma_g = 0.095\%$ and becomes statistically significant at the 5% level, while the coefficients on acceleration and size get much larger. These patterns make sense. The 2SLS approach leverages shifts in the demand for acceleration and size, which cause these attributes to go up and fuel economy to go down. In contrast, OLS fails to purge unobserved drivetrain costs (k), which cause all attributes to go down, biasing coefficients on acceleration and size toward zero (and the coefficient

¹⁹Sanderson and Windmeijer (2016) applies to cases with multiple endogenous regressors. The test asks whether the excluded instruments jointly predict log(HPWt) after partialing-out variation that is correlated with log(volume) as predicted by the excluded instruments, and vice versa. We report first-stage coefficients and standard unconditional F-tests for the joint significance of the excluded instruments in appendix D. The first-stage coefficients for model (7) imply that males choose 3% more acceleration than females, while each household member boosts size 2%.

Figure 3: *Trend in log fuel economy conditional on car attributes*



Note: This figure plots the trend in log fuel economy holding attributes (but not cost) fixed based on the year dummies from regressions (1), (3), (4), (6), and (7) from table 2. In addition, the figure plots the trend in log fuel economy holding both attributes and cost fixed by dividing the year dummies from regression (7) by $1 - \gamma_g = \alpha_g / (1 + \alpha_g)$.

on pickup away from zero). The results imply that fuel economy decreases 0.720% and 0.743% for every 1% increase in acceleration and size holding drive-train cost fixed. The iso-cost elasticities are much larger than the reduced-form correlations in column (1).

Column (5) adds controls for log odometer (= 0 if missing), car age, and a missing odometer dummy, along with their squares and interactions. These variables are intended to capture heterogeneity in the present value of lifetime miles; cars driven more intensively early in their lives have higher effective miles. The coefficients on gas prices and car attributes barely budge. Column (6) further adds our full suite of demographic controls, including categories of income, # adults, home ownership, # vehicles, urban vs. rural, and population density. These variables are intended to capture unobserved car attributes or demand for fuel economy that would otherwise bias our cost estimates. The coefficients on size and acceleration fall slightly but the coefficient on gas prices barely changes. Finally, column (7) adds our shadow-cost adjustment to logged gas prices. We instrument for the adjusted variable using log gas prices, which ensures that we use the same identifying variation as before while purging the time trend in shadow costs from our estimated

technical change.²⁰ The coefficients in the table barely budge. Overall, these results highlight the danger of inferring iso-cost curves from reduced-form equilibrium relationships as in Knittel (2011).

Figure 3 again plots the coefficients on the year dummies. Model (3) controls for gas prices. The year coefficients (long-dash gray line) are substantially smaller than those from the Knittel approach in model (1). This difference is largely driven by the strong upward trend in gas prices, with falling interest rates playing a minor role.²¹ Model (4) instruments for acceleration and size. The year coefficients (short-dash gray line) shift back upward. Steeper iso-cost curves imply that the super-sizing of cars has hurt fuel economy more. Thus, more rapid technical change is needed to explain rising fuel economy over time. Model (6) controls for car age, odometer, and full demographics. The year coefficients (short-dash black line) are somewhat lower. Model (7) adjusts for the shadow cost. The year coefficients (long-dash black line) are lower still.

To recover technical change ($\Delta\theta$) we multiply the year dummies by $1/\hat{\gamma}_p$. We then divide by the cost parameter on fuel economy ($\hat{\alpha}_g = 8.7$) to yield the annual gain in fuel economy possible holding both drivetrain costs and other car attributes fixed. Figure 3 reports the results (solid black line). Note that the raw year dummies from model (7) understate technical change because the equilibrium value does not hold cost fixed. Based on this approach, we calculate technical change averaging 1.4% annually over our 22-year sample period; the rate appears to increase post-2005 under rising gas prices and efficiency standards. Alternatively, we divide by the sum of our cost parameters ($\hat{\alpha}_s + \hat{\alpha}_a + \hat{\alpha}_g = 20.9$) to yield the annual gain in all car attributes possible under technical change holding drivetrain costs fixed: 0.6%.

5.1.2 Robustness to alternative assumptions

We show in appendix D that our results are generally robust to a variety of alternative assumptions. We first replace the three-year moving average of gas prices with the current year’s price plus four lags. The coefficient on the current price is 0.013 and statistically insignificant; this is not surprising, given that many cars may be sold early in a calendar year before that year’s gas price is even known. However, the coefficient on the lagged price is large (0.063) and significant, and the cumulative effect across all five years is virtually identical to our main results based on a three-year moving average.

Next we include make dummies to capture variation in markups across major firms and their brands; recall that our log-linear model already captures any common markup that scales with drivetrain costs (see appendix C). That is, for example, we allow Ford’s markups to differ from Toyota’s markups, and for markups on Toyota’s Toyota-branded cars to differ from the markups on its Lexus-branded cars. Thus, we only use within-make variation to identify cost parameters. The coefficient on gas prices falls from 0.103 to

²⁰Note that the shadow-cost adjustment includes state gas prices in the denominator, which we find leads to nonsensical results when we do not instrument for adjusted gas prices.

²¹To confirm, we drop our adjustment for interest rates and depreciation from the gas price variable, such that trends in interest rates load onto the year dummies. The year coefficients in this case closely match those shown in figure 3.

0.082; the coefficient on acceleration shrinks moderately while the coefficient on size grows slightly.

Next we consider different estimation samples. We start by focusing on newer cars that are more likely to remain with their original owners. We omit cars in the 2009 NHTS from 2000 or earlier and cars in the 2017 NHTS from 2008 or earlier.²² The coefficient on gas prices falls from 0.102 to 0.069 (insignificant) while the coefficients on size and acceleration change relatively little. Next we drop NHTS survey weights. The coefficient on gas price is virtually unchanged, while the coefficients on size and acceleration change little. This finding is consistent with our theoretical assumption of a universal cost function and demographic controls that capture most of the relevant preference heterogeneity across households.

Finally, we adjust gas prices using our high-end assumption for the shadow cost on efficiency standards. Our coefficients barely budge. The higher shadow prices lead to modestly lower attribute-neutral technical change over time. But they do not affect our estimates for attribute trade-offs at a point in time.

5.1.3 Biased technical change

Have the slopes of the iso-cost curves shifted over time? Table 3 shows 2SLS estimates for a version of equation (27) in which we allow the coefficients on size and acceleration to differ over time but otherwise follow model (7) from table 3. Columns (1)–(3) report the results from a single regression which allows the coefficients on acceleration and size to differ across three eras (1995–2001, 2002–2009, and 2010–2017) but which imposes constant coefficients on all other variables; we use three columns to report the coefficients from each era separately.²³ The results show that iso-cost curves have become less steep (see the bottom of the table) albeit with a curious episode of non-monotonicity in the middle era.

Column (4) captures the overall shift by interacting acceleration and size with a linear time trend (re-scaled to range from 0 to 1). The results suggest that the iso-cost curves are only half or three-quarters as steep in 2017 as they were in 1995. The t-statistic on the trend in acceleration is only slightly bigger than one. Yet the point estimates are consistent with engineering simulations which show that the fuel economy vs. acceleration trade-off has flattened over time with the accumulation of technology (see appendix A).

5.2 Consumer preferences

We now estimate preferences for size and acceleration (μ_a and μ_s). We follow equations (12) and (13) above, regressing log acceleration and log size on their marginal-cost shifters, which we construct from the other log attributes and our estimated cost parameters. Note that the error term in the regression for log acceleration includes unobserved demand and cost shifters ($\ln \beta_a$ and k) that affect equilibrium choices of size and efficiency. Likewise, the error term in the regression for log size includes unobservables ($\ln \beta_s$ and k) that

²²We preserve the one year of overlap with the prior NHTS wave (2001 and 2009) so that we can still identify the coefficients on car age, which would otherwise be collinear with the year dummies.

²³We interact acceleration and size with each of the three era dummies; we likewise interact our instruments measuring male share and # household members with each of the era dummies.

Table 3: *Biased technical change*

	(1)	(2)	(3)	(4)
	1995-01	2002-09	2010-17	Trend
Log gas price (γ_p)	0.131 (0.067)	0.131 (0.067)	0.131 (0.067)	0.107 (0.057)
Log HP/weight (γ_a)	-0.528* (0.255)	-0.945*** (0.276)	-0.435*** (0.095)	-0.826** (0.290)
Log HP/weight $\times$ trend				0.401 (0.358)
Log volume (γ_s)	-0.866*** (0.179)	-0.498*** (0.119)	-0.729*** (0.104)	-0.753*** (0.158)
Log volume $\times$ trend				0.147 (0.272)
Pickup	-0.026** (0.009)	-0.026** (0.009)	-0.026** (0.009)	-0.024** (0.009)
$\frac{\partial \ln \text{mpg}}{\partial \ln \text{HP/Wt}} _{\Delta c=0} \left(\frac{\alpha_a}{\alpha_g} = \frac{\gamma_a}{1-\gamma_p} \right)$	-0.607 (0.288)	-1.087 (0.353)	-0.500 (0.106)	-0.926 (0.332)
In 2017				-0.476 (0.132)
$\frac{\partial \ln \text{mpg}}{\partial \ln \text{volume}} _{\Delta c=0} \left(\frac{\alpha_s}{\alpha_g} = \frac{\gamma_s}{1-\gamma_p} \right)$	-0.997 (0.214)	-0.572 (0.133)	-0.838 (0.155)	-0.843 (0.198)
In 2017				-0.679 (0.161)

Note: This table presents coefficient estimates from equation (27) in which we interact acceleration and size with era dummies (columns 1–3) or a time trend re-scaled to range from 0 to 1 (column 4). $N = 245,277$. Columns (1)–(3) report coefficients from the same model; we use three columns to report the era-specific coefficients on acceleration and size. Standard errors in parentheses are clustered by state. Note *, **, and *** indicate statistical significance at the 5%, 1%, and 0.1% levels respectively.

affect equilibrium choices of acceleration and efficiency. Thus, in both cases we instrument for marginal costs using log gas prices and our odometer-based variables (log, log squared, log $\times$ car age, missing odometer, and missing $\times$ age). Intuitively, higher gas prices and miles cause households to choose more efficient cars, raising the marginal cost of size and acceleration. Our identification assumption is that gas prices and miles only affect size and acceleration indirectly via the demand for fuel economy. This assumption is natural for gas prices. It cannot literally be true for miles, since consumers jointly choose miles and fuel economy. Thus, the assumption is that major factors that determine miles like where to live and work have large variance relative to unobserved factors shifting demand for size and acceleration. Intuitively: the “need to drive a lot of miles, so choose an efficient car” effect dominates the “want a big car, so drive few miles” effect. Previous research generally finds that miles driven is highly insensitive to gas prices and fuel economy, which supports

Table 4: *Auxiliary regressions to identify demand elasticities*

	(1)	(2)	(3)
	OLS	2SLS	2SLS
log HP/weight			
Cost shifter ($\hat{\theta} - \hat{\alpha}_s \ln s + \hat{\alpha}_g \ln g$)	0.107 [0.018,0.172]	0.124 [0.024,0.189]	0.123 [0.023,0.185]
Elasticity	-0.224 [-0.336,-0.041]	-0.319 [-0.726,-0.077]	-0.308 [-0.652,-0.077]
F-stat		53.653	64.500
P-value		[0.000]	[0.000]
log Volume			
Cost shifter ($\hat{\theta} - \hat{\alpha}_a \ln a + \hat{\alpha}_g \ln g$)	0.095 [0.016,0.150]	0.123 [0.026,0.163]	0.130 [0.027,0.168]
Elasticity	-0.177 [-0.305,-0.033]	-0.309 [-0.512,-0.054]	-0.357 [-0.612,-0.001]
F-stat		60.372	44.905
P-value		[0.000]	[0.000]

Note: This table presents coefficient estimates for log HP/Wt (top panel) and log volume (bottom) regressed on their constructed cost shifters as in equations (12) and (13). $N = 245,277$. Columns (2)–(3) instrument for the cost shifter with log gas prices and odometer-based variables (log, log squared, log $\times$ car age, missing odometer, and missing $\times$ age). All columns control for state and year fixed effects, the pickup dummy, car age and its square, male share, and # household members. Column (3) additionally controls for categories of income, # adults, home ownership, # vehicles, urban vs. rural, and population density. Middle rows report corresponding elasticity estimates. Bootstrapped 95% confidence intervals robust to state clustering are given in brackets for both the cost shifter and elasticity. Bottom rows report first-stage F-tests for weak instruments robust to state clustering (given cost parameters).

this assumption²⁴ We include state and year fixed effects and control for male share, # household members, pickup, and vehicle age and its square across all models; we also control for our full suite of demographics in one model. The instruments strongly predict size and acceleration conditional on other controls, as indicated by standard robust F-tests.²⁵

Table 4 reports our results. Column (1) presents coefficients for OLS regressions of log acceleration and log size on their constructed cost shifters ($\hat{\theta} - \hat{\alpha}_a \ln a + \hat{\alpha}_g \ln g$ and $\hat{\theta} - \hat{\alpha}_s \ln s + \hat{\alpha}_g \ln g$) and controls; we draw our cost parameters from model (7) in table 2. Column (2) presents 2SLS results, instrumenting for the cost shifters with log gas prices and our odometer-based variables. The coefficients grow moderately in magnitude. Column (3) adds controls for income, # adults, home ownership, # vehicles, urban vs. rural, and population density. The coefficients barely budge. Below each coefficient we report the implied elasticity of demand, which is a function of the corresponding demand and cost parameters (μ and α). These elasticities

²⁴See Anderson and Sallee (2016) and Gillingham, Rapson and Wagner (2016) for reviews covering response of miles to gas prices. West, Hoekstra, Meer and Puller (2017) show that the Obama-era cash-for-clunkers program induced households at the margin of eligibility to purchase more fuel-efficient cars, but there was zero effect on miles driven.

²⁵We report first-stage results and corresponding F-tests for the joint significance of the instruments in appendix D.

Table 5: *Equilibrium responses to gas prices and technology*

Underlying causal mechanism	Attribute changes			
	$\Delta \ln mpg$	$\Delta \ln a$	$\Delta \ln s$	$\Delta \ln c$
$\Delta \ln(pm + \tau)$	0.46	-0.28	-0.30	0.54
$\Delta \theta$ (per $\Delta \ln mpg$)	1.00	0.53	0.56	-1.00
$\Delta \omega_g$ fuel economy	0.54	0.28	0.30	-0.54
$\Delta \omega_a$ acceleration	0.38	0.20	0.21	-0.38
$\Delta \omega_s$ size	0.36	0.19	0.20	-0.36

Note: This table shows equilibrium changes in car attributes and drivetrain costs (column headings) associated with underlying causal mechanisms (row labels). Both attribute changes and causal mechanisms are given in percentage terms; for attribute-neutral technical change ($\Delta \theta$), percent changes are scaled by the percent change in fuel economy. Table assumes cost parameters based on regression (7) from table 2 and demand parameters are based on column (3) from table 4.

are -0.308 for acceleration and -0.357 for size. Both are statistically significant based on bootstrapped 95% confidence intervals that are robust to state clustering.²⁶

5.3 Equilibrium responses to gas prices and technical change

We now use our estimated cost and preference parameters in combination with our theoretical comparative statics to calculate long-run equilibrium responses to gas prices and technical change. Table 5 reports the results: percent changes in fuel economy, size, acceleration, and vehicle cost (indicated by the column headings) in response to changes in gas prices, attribute-neutral technical change, and adoption of discrete attribute-boosting technologies (given by the row headings).


First consider gas prices. The first row indicates that a 1% increase in gas prices (technically $pm + \tau$) leads to a 0.46% increase in fuel economy, a 0.28% decrease in acceleration, a 0.30% decrease in size, and a 0.54% increase in cost. The cost increase conflicts with the notion of technical feasibility in Knittel (2011) that fuel economy can only be improved by reducing size and acceleration holding cost fixed.

Next consider attribute-neutral technical change (θ). The second row indicates that a 1% increase in fuel economy induced by such change is associated with a 0.53% increase in acceleration, a 0.56% increase in size, and a 1% reduction in cost. Technical change is mostly channeled into fuel economy and cost savings, due to inelastic demand for size and acceleration. The percent gain in fuel economy equals the percent reduction in cost; this result is implied by the optimizing condition for fuel economy given our functional form.²⁷

²⁶We bootstrap confidence intervals to account for the two-step procedure: (1) estimating cost parameters (α 's); (2) estimating demand parameters ($\hat{\mu}$'s) conditional on the estimated cost parameters. We apply Rubin's (1981) Bayesian bootstrap, drawing random Dirichlet weights for each state (iid exponential), which we combine with the NHTS survey weights in both steps of the estimation for each bootstrap iteration.

²⁷Consider a small gain in technology $\Delta \theta > 0$. Suppose log fuel economy and log cost change by different amounts in equilibrium. Then $\Delta \ln g \neq -\Delta \theta + \alpha_s \Delta \ln s + \alpha_a \Delta \ln a - \alpha_g \Delta \ln g$, which violates equation (9) above.

Figure 4: *Discrete technology adoption*



Note: This figure shows the sales-weighted shares of car models with the given technology by cost index decile and model year. Deciles are computed within each 5-year bin of model year and the estimated cost index uses coefficients from specification (7) in table 2.

Finally, consider a discrete attribute-boosting technology (ω 's). The third row indicates that a 1% boost to fuel economy leads to a 0.54% increase in fuel economy, a 0.28% increase in acceleration, a 0.30% increase in size, and a 0.54% decrease in cost (ignoring adoption cost κ). Thus, a full 46% of the direct gain in fuel economy is reallocated elsewhere in equilibrium. Conversely, the fourth and fifth rows indicate that more than a third of the initial gain from an acceleration- or size-boosting technology is reallocated to fuel economy. What about a discrete technology that simultaneously boosts more than one attribute? Simply add rows vertically, weighted by the corresponding boosts, to get the overall effect.

5.4 New engine technologies are not evenly adopted in the fleet

Figure 4 shows the fraction of cars adopting specific technologies by drivetrain cost and model year. To construct this figure, we compute a cost index based on the log of estimated drivetrain costs:

$$\text{cost index} = \hat{\alpha}_g \ln \text{mpg} + \hat{\alpha}_a \ln \text{HP}/\text{Wt} + \hat{\alpha}_s \ln \text{volume}, \quad (28)$$

with parameters ($\hat{\alpha}$'s) drawn from model (7) in table 2 above. We compute this index at the trim-level using Wards Automotive data, which contains detailed information on technology adoption. We split trims into cost deciles and 5-year bins and calculate the fraction of trims within each cell that adopt each of

Table 6: *Discrete technology adoption correlations*

	Turbo	GDI	CVT	Hybrid	6+Speed
Cost index	0.016*** (0.001)	0.004** (0.001)	0.024*** (0.001)	0.015*** (0.001)	0.025*** (0.002)

Note: This table presents OLS regression results for technology dummies regressed on the constructed cost index and year fixed effects. Note *, **, and *** indicate statistical significance at the 5%, 1%, and 0.1% levels.

five technologies: turbo and superchargers, gasoline direct injection (GDI), continuous variable transmission (CVT), 6+ transmission speeds (for trims lacking CVT), and hybrid gasoline-electric motors.

Figure 4 shows that costlier trims are more likely to feature advanced drivetrain technologies in any given year, while adoption rates increase over time, which is consistent with falling adoption costs. Note that adoption rates for gasoline-direct injection (GDI), continuous variable transmissions (CVT), 6+ transmission speeds, and hybrid gasoline-electric motors were all close to zero at the start of the sample, and that adoption since then has skewed toward high-cost trims. These findings are consistent with our theoretical prediction that new technologies are first adopted by consumer types choosing high-cost drivetrains.

To confirm this finding, we estimate linear probability models of the following form:

$$technology_{jt} = \gamma \cdot cost\ index_{jt} + \phi_t + \varepsilon_{jt}, \quad (29)$$

where: $technology_{jt}$ is a technology adoption dummy for car trim j in year t ; γ captures the correlation between adoption and cost; ϕ_t is a year fixed effect; and ε_{jt} is an additive error term. For this exercise, we opt not to match vehicle trims with the NHTS micro-data as we would lose crucial trim-level detail; we instead weight by trim-level sales from EPA.

Table 6 presents the OLS estimation results. Consider a doubling of drivetrain costs (a one-unit increase in the index). This increase is associated with higher adoption rates across the board: 0.016 for turbo-chargers, 0.004 for GDI, 0.024 for CVT, 0.015 for hybrids, and 0.025 for 6+ speeds (for trims lacking CVT). These findings further emphasize the danger of inferring iso-cost relationships from reduced-form correlations. Costly technologies are disproportionately applied to attribute-rich cars and can be used to enhance all attributes, not just fuel economy. Thus, OLS regressions of fuel economy on size and acceleration will tend to understate true iso-cost trade-offs.

6 Counterfactual simulations

We use our model to quantify the role of gas prices, technology, preferences, and policy in driving trends in car attributes and costs during 1995–2017, and to estimate marginal abatement costs for 2012–2017.

6.1 Model estimation and calibration

Cost parameters. We estimate cost parameters via model (7) in table 2. See the table for the estimated parameter on fuel economy (α_g) and corresponding iso-cost relationships (α_a/α_g and α_s/α_g). Figure 3 plots attribute-neutral technical change (θ) based on a transformation of the year dummies. We do not consider attribute-biased technical change in our simulations, given the suggestive nature of these estimates.

Demand for fuel economy. We calibrate demand as log gas price times lifetime miles plus the shadow cost: $\ln(pm + \tau)$, where p is the three-year moving average of state-level gas prices; lifetime miles $m = m_0/(r + \delta + \rho)$ trends over time with the interest rate (r) given our estimates of first-year miles (m_0), miles decay (δ), and scrap rate (ρ); and τ is the mid-range estimate for the shadow cost on efficiency standards.

Cost residual. Given the above values, we calibrate the cost shifter (k) via equation (10) for every car in our dataset. Note that whatever we choose for lifetime miles, we calibrate a cost residual (k) that rationalizes a car’s baseline fuel economy. Thus, allowing for heterogeneous miles would change little.


Demand for size and acceleration. Given costs, we estimate the utility parameters (μ ’s) via column (3) from table 4. Finally, given the full set of cost and demand parameters (α ’s, μ ’s, θ , and k), we calibrate the heterogeneous preference shifters (β ’s) for each car in our sample as the unique solutions to equations (12)–(13). Note that the equation for size includes the shadow-cost incentive ($\tau\sigma'(s)/c_s$) in 2011 and later. We estimate $\sigma'(s)$ via an OLS regression of EPA’s footprint-based standard on our measure of car size (volume) and a set of year dummies, using our matched NHTS-EPA data. Preferences for size and acceleration vary widely across household and are strongly positively correlated. Large, high-income households in low-density census tracts have the highest demand for both acceleration and size; meanwhile, male drivers like fast cars, and female drivers like big cars (see appendix D).

6.2 Structural decomposition via counterfactual simulation

Figure 5 presents the results of our structural decompositions for fuel economy (top left), acceleration (top right), size (bottom left), and cost (bottom right). Baseline trends (labeled “base”) show changes in log mean attributes relative to 1995. The decomposition ascribes these changes to underlying trends in gas prices (“gas”), technical change (“tech”), preferences (“prefs”), and efficiency standards (“stand”). To do this, we perform three counterfactual simulations that reset gas prices (p), technical change (θ), and preferences (β_a & β_s) to their 1995 levels all-else-equal; we perform a fourth counterfactual that sets the shadow cost (τ) to zero in all years. We make these counterfactual changes and then re-calculate equilibrium size, acceleration, and fuel economy for each car in our dataset (see appendix B.4). We finally compute the mean values in each year and calculate the deviation from baseline in logs.

Figure 5 shows that gas prices, technology, and preferences are all important for explaining trends in car

Figure 5: *Structural decompositions for attribute trends*



Note: This figure plots structural decompositions for trends in fuel economy (top left), acceleration (top right), size (bottom left), and cost (bottom right). Bold lines indicate baseline trends (“base”), which are decomposed into the underlying effects of gas prices (“gas”), attribute-neutral technical change (“tech”), consumer preferences (“prefs”), and efficiency standards (“stand”). See text for details.


attributes, while efficiency standards have done little. Fluctuating gas prices caused fuel economy and cost to rise in the mid-2000s, plateau in the 2010s, and then fall in recent years, with size and acceleration moving in the opposite direction. By the end of the sample period, gas prices had caused fuel economy and cost to rise by 15% and 17% and size and acceleration to each fall by 9%.

Meanwhile, attribute-neutral technical change caused steady gains in all dimensions. By the end of the sample period, technical change had caused fuel economy to rise and cost to fall by 17% and size and acceleration to each rise by 9%. These results illustrate how changes in technology are mediated by consumer preferences. Demand for size and acceleration are both inelastic. Thus, most of the attribute-neutral technical change has been channeled into fuel economy gains and cost reductions.

Yet consumer preferences have driven striking increases in both size and acceleration. By the end of the sample period, surging preferences had caused size to rise by 34%, acceleration to rise by 21%, fuel economy to fall by 34%, and cost to rise by 34%.²⁸ These results illustrate how shifting demand for performance can

²⁸These end-of-sample numbers likely understate the true shift in preferences. Note that our model identifies preferences for size and acceleration relative to fuel economy, such that absolute preferences fall off with the short-run drop in gas prices at

Figure 6: *Marginal abatement costs for 2012–2017*



Note: This figure plots marginal abatement cost vs. percent abatement via fuel economy standards while (a) allowing both drivetrain costs and attributes to adjust (labeled “Equilibrium”), (b) holding drivetrain cost fixed (“dCost=0”) but allowing attributes to adjust, or (c) holding attributes fixed (“dAttributes=0”) but allowing for more costly drivetrains. We assume no size-basing.

lead to the super-sizing of cars, especially when demand is inelastic. Our calibrated preference shifters ($\ln \beta_a$ and $\ln \beta_s$) are positively correlated with income and negatively correlated with density, suggesting that these shifts may be driven in part by rising incomes and suburbanization (see appendix D).

Finally, efficiency standards have had a small effect on car attributes, limited mainly to 2012–2016 when positive shadow prices indicate a binding standard. Standards caused fuel economy and cost to be roughly 5% higher and size and acceleration to be 3% lower during these years. Ancillary simulations indicate that the size-based component of current standards has had a negligible effect on car attributes for two reasons: the policy incentive to increase size ($\tau\sigma'(s)/c_s$) is small relative to underlying marginal willingness to pay among consumers ($\ln \beta_s$), and demand for size is inelastic.

6.3 Marginal abatement costs in equilibrium

Our structural decomposition shows that fuel-economy standards historically were too weak to have had much effect. How costly would it have been to reduce emissions via tighter standards? Figure 6 shows marginal abatement costs (\$/tCO₂) in equilibrium for efficiency standards applied to new cars during 2012–2017. Mechanically, we generate this figure by ratcheting up the shadow price in our model (\$ per lifetime emissions) starting from zero and tracing out the equilibrium relationship with total abatement the end of the sample.

(percent reduction in lifetime emissions). For comparison, we also calculate marginal abatement costs while holding drivetrain cost fixed and allowing attributes to adjust (labeled “dCost = 0”) as in Knittel (2011), and while holding size and acceleration fixed but allowing for more costly drivetrains (labeled “dAttributes = 0”) as in EPA regulatory impact analyses.

The figure shows that marginal abatement costs are only slightly higher when we hold drivetrain costs fixed but roughly four times higher when we hold attributes fixed. These results suggest that EPA’s regulatory analyses may greatly overstate abatement costs by holding attributes fixed, since down-sizing offers relatively low-cost abatement. Meanwhile, the results imply that equilibrium behavior closely approximates the attribute-only response as in Knittel (2011), despite inelastic demand for size and acceleration, since adding fuel-saving technology is an even costlier compliance strategy.

7 Conclusion and discussion

Why have car size, acceleration, and fuel economy all increased in recent years, reversing the trend of the previous three decades? We find that technical change has facilitated a moderate increase in all car attributes, with inelastic demand for size and acceleration channeling much of this change into efficiency gains and cost savings. Meanwhile, gas prices have pushed cars to be more efficient at the expense of size and power, while surging demand for size and power have pushed strongly in the opposite direction. On net, trends in fuel economy and acceleration have mostly tracked the underlying technical change, while preferences for size have overwhelmed gas prices, leading to the super-sizing of cars.

Implications for electric vehicles. We do not directly study electric vehicles (EVs), which post-date our sample. Yet our model yields valuable insights. Adding EVs would require modeling another car attribute: range. Electric drivetrains reduce cost-per-mile and improve acceleration but reduce range. Larger batteries can extend range but cost more and increase weight. Our analysis of discrete technology suggests that the adoption of EVs should first occur among drivers who prioritize fuel costs, crave speed, and, to a lesser extent, prefer big cars—but do not especially value range. Indeed, EVs are already poised to replace highly efficient gas-powered cars, sports cars, and some SUVs. As EV costs fall and efficiency standards tighten, EVs will gradually replace larger, slower, and less efficient cars among range-loving drivers. Our analysis suggests that much of the direct improvement in cost of driving and acceleration will be reallocated to other car attributes. We suspect that demand for range is highly inelastic, which explains why EV manufacturers have focused on extending range for mid-sized EVs before moving into the larger segments.

Implications for efficiency standards. Our equilibrium approach contrasts with EPA’s regulatory impact analyses at one extreme and Knittel’s (2011) empirical analysis at the other. EPA holds car attributes fixed and calculates the cost of adding new technologies to meet standards. Knittel (2011) attempts to hold costs

fixed and asks whether standards can be met via reductions in size and acceleration. Meanwhile, our analysis quantifies the optimal combination of cost increases and attribute changes to meet standards. Our finding of inelastic demand for size and acceleration is consistent with the EPA approach. Yet a steeply increasing marginal cost for improving fuel economy is consistent with the Knittel (2011) approach. On net, the latter effect dominates, such that the EPA approach vastly overstates compliance costs.

Implications for directed technical change. We conclude with four lessons for policies that aim to speed and steer the direction of technical change. First, attribute-neutral technical change causes consumers to buy more of all attributes. If externalities are not priced, overall welfare may increase in some dimensions (less pollution) but decrease in others (more fatal accidents). Second, technology subsidies boost adoption among consumers that choose high levels of all attributes, making subsidies regressive. Further, the environmental benefit from an efficiency-boosting technology will be reallocated to other attributes in equilibrium, eroding the initial gain. Third, directed technical change that reduces the marginal cost of efficiency relative to other attributes will pull consumers in a more efficient direction, and dampen the effect of a carbon tax or efficiency standard on other attributes, perhaps making carbon regulation more politically palatable. Finally, the equilibrium effects of carbon taxes, efficiency standards, and technical change are all mediated by consumer preferences, and shifts in demand for other attributes can undermine policies to improve efficiency. This concern is especially salient when the demand shifts are associated with other untaxed externalities, such as vehicle size (see Jacobsen, 2013; Anderson and Auffhammer, 2014; Bento et al., 2017).

References

- Acemoglu, Daron**, “Directed Technical Change,” *The Review of Economic Studies*, 2002, *69* (4), 781–809.
- , “Technical Change, Inequality, and the Labor Market,” *Journal of Economic Literature*, 2002, *40* (1), 7–72.
- , **Philippe Aghion**, **Leonardo Bursztyn**, and **David Hemous**, “The Environment and Directed Technical Change,” *American Economic Review*, 2012, *102* (1), 131–166.
- Aghion, Philippe**, **Antoine Dechezleprêtre**, **David Hemous**, **Ralf Martin**, and **John Van Reenen**, “Carbon Taxes, Path Dependency, and Directed Technical Change: Evidence from the Auto Industry,” *Journal of Political Economy*, 2016, *124* (1), 1–51.
- Allcott, Hunt** and **Christopher Knittel**, “Are Consumers Poorly Informed about Fuel Economy? Evidence from Two Experiments,” *American Economic Journal: Economic Policy*, February 2019, *11* (1), 1–37.


- **and Nathan Wozny**, “Gasoline Prices, Fuel Economy, and the Energy Paradox,” *The Review of Economics and Statistics*, 2014, *96* (5), 779–795.
- Anderson, Michael L and Maximilian Auffhammer**, “Pounds That Kill: The External Costs of Vehicle Weight,” *Review of Economic Studies*, 2014, *81* (2), 535–571.
- Anderson, Soren T and James M Sallee**, “Using Loopholes to Reveal the Marginal Cost of Regulation: The Case of Fuel Economy Standards,” *American Economic Review*, 2011, *101* (4), 1375–1409.
- **and —**, “Designing Policies to Make Cars Greener,” *Annual Review of Resource Economics*, 2016, *8*, 157–180.
- Anderson, Soren T., Ryan Kellogg, and James M. Sallee**, “What Do Consumers Believe About Future Gasoline Prices?,” *Journal of Environmental Economics and Management*, 2013, *66* (3), 383–403.
- Bento, Antonio, Kenneth Gillingham, and Kevin Roth**, “The Effect of Fuel Economy Standards on Vehicle Weight Dispersion and Accident Fatalities,” 2017. NBER Working Paper w23340.
- Busse, Meghan R., Christopher R. Knittel, and Florian Zettelmeyer**, “Are Consumers Myopic? Evidence from New and Used Car Purchases,” *American Economic Review*, February 2013, *103* (1), 220–56.
- Calel, Raphael and Antoine Dechezleprêtre**, “Environmental Policy and Directed Technological Change: Evidence from the European Carbon Market,” *Review of Economics and Statistics*, 2016, *98* (1), 173–191.
- Chouinard, Hayley and Jeffrey Perloff**, “Gasoline Price Differences: Taxes, Pollution Regulations, Mergers, Market Power, and Market Conditions,” *The B.E. Journal of Economic Analysis Policy*, 2007, *7* (1).
- Dekraker, Paul, Daniel Barba, Andrew Moskalik, and Karla Butters**, “Constructing Engine Maps for Full Vehicle Simulation Modeling,” *WCX World Congress Experience*, 2018.
- Environmental Protection Agency**, “The 2019 Automotive Trends Report EPA-420-R-20-006,” 2019.
- , “Inventory of U.S. Greenhouse Gas Emissions and Sinks: 1990-2018,” 2020.
- Federal Highway Administration, U.S. Department of Transportation**, “National Household Travel Survey,” 2001–2017.
- Gillingham, Kenneth, David Rapson, and Gernot Wagner**, “The Rebound Effect and Energy Efficiency Policy,” *Review of Environmental Economics and Policy*, 2016, *10* (1), 68–88.

- Gillingham, Kenneth T., Sébastien Houde, and Arthur A. van Benthem**, “Consumer Myopia in Vehicle Purchases: Evidence from a Natural Experiment,” *American Economic Journal: Economic Policy*, August 2021, *13* (3), 207–38.
- Grieco, Paul L.E., Charles Murry, and Ali Yurukoglu**, “The Evolution of Market Power in the U.S. Automobile Industry,” *The Quarterly Journal of Economics*, 09 2023, *139* (2), 1201–1253.
- Ito, Koichiro and James M. Sallee**, “The Economics of Attribute-Based Regulation: Theory and Evidence from Fuel Economy Standards,” *The Review of Economics and Statistics*, 05 2018, *100* (2), 319–336.
- Jacobsen, Mark R.**, “Fuel Economy and Safety: The Influences of Vehicle Class and Driver Behavior,” *American Economic Journal: Applied Economics*, 2013, *5* (3), 1–26.
- , **Christopher R Knittel, James M Sallee, and Arthur A Van Benthem**, “The Use of Regression Statistics to Analyze Imperfect Pricing Policies,” *Journal of Political Economy*, 2020, *128* (5), 1826–1876.
- Kellogg, Ryan**, “Gasoline Price Uncertainty and the Design of Fuel Economy Standards,” *Journal of Public Economics*, 2018, *160*, 14–32.
- , “Output and Attribute-based Carbon Regulation Under Uncertainty,” *Journal of Public Economics*, 2020, *190*, 104246.
- Kilian, Lutz and Xiaoqing Zhou**, “Heterogeneity in the Pass-through from Oil to Gasoline Prices: A New Instrument for Estimating the Price Elasticity of Gasoline Demand,” *Journal of Public Economics*, 2024, *232*, 105099.
- Klier, Thomas and Joshua Linn**, “The Effect of Vehicle Fuel Economy Standards on Technology Adoption,” *Journal of Public Economics*, 2016, *133*, 41–63.
- Knittel, Christopher R.**, “Automobiles on Steroids: Product Attribute Trade-Offs and Technological Progress in the Automobile Sector,” *American Economic Review*, 2011, *101* (7), 3368–3399.
- Leard, Benjamin, Joshua Linn, and Yichen Christy Zhou**, “How Much Do Consumers Value Fuel Economy and Performance? Evidence from Technology Adoption,” *Review of Economics and Statistics*, 2023, *105* (1), 158–174.
- Li, Shanjun**, “Traffic Safety and Vehicle Choice: Quantifying the Effects of the ‘Arms Race’ on American Roads,” *Journal of Applied Econometrics*, 2012, *27* (1).
- Lin, Yujie and Joshua Linn**, “Environmental Regulation and Product Attributes: The Case of European Passenger Vehicle Greenhouse Gas Emissions Standards,” *Journal of the Association of Environmental and Resource Economists*, 2023, *10* (1), 1–32.

- MacKenzie, Don and John Heywood**, “Acceleration Performance Trends and Evolving Relationship between Power, Weight, and Acceleration in U.S. Light-Duty Vehicles: Linear Regression Analysis,” *Transportation Research Record: Journal of the Transportation Research Board*, 2012, 2287 (1), 122–131.
- Moskalik, Andrew, Kevin Bolon, Kevin Newman, and Jeff Cherry**, “Representing GHG Reduction Technologies in the Future Fleet with Full Vehicle Simulation,” *WCX World Congress Experience*, 2018.
- Newell, Richard G, Adam B Jaffe, and Robert N Stavins**, “The Induced Innovation Hypothesis and Energy-saving Technological Change,” *The Quarterly Journal of Economics*, 1999, 114 (3), 941–975.
- NHTSA**, “Vehicle Survivability and Travel Mileage Schedules,” 2006.
- Rubin, Donald B.**, “The Bayesian Bootstrap,” *The Annals of Statistics*, 1981, 9 (1), 130 – 134.
- Sallee, James M., Sarah E. West, and Wei Fan**, “Do Consumers Recognize the Value of Fuel Economy? Evidence from Used Car Prices and Gasoline Price Fluctuations,” *Journal of Public Economics*, 2016, 135, 61–73.
- Sanderson, Eleanor and Frank Windmeijer**, “A weak Instrument F-test in Linear IV Models with Multiple Endogenous Variables,” *Journal of Econometrics*, 2016, 190 (2), 212–221.
- Shahed, SM and Karl-Heinz Bauer**, “Parametric Studies of the Impact of Turbocharging on Gasoline Engine Downsizing,” *SAE International Journal of Engines*, 2009, 2 (1), 1347–1358.
- U.S. Federal Reserve Bank**, “Consumer Credit-G. 19, January 2019,” 2019. Retrieved August 20, 2021.
- West, Jeremy, Mark Hoekstra, Jonathan Meer, and Steven Puller**, “Vehicle Miles (not) Traveled: Fuel Economy Requirements, Vehicle Characteristics, and Household Driving,” *Journal of Public Economics*, 2017, 145 (C), 65–81.
- Whitefoot, Kate S. and Steven J. Skerlos**, “Design Incentives to Increase Vehicle Size Created from the U.S. Footprint-Based Fuel Economy Standards,” *Energy Policy*, 2012, 41, 402–411.
- , **Meredith L. Fowlie, and Steven J. Skerlos**, “Compliance by Design: Influence of Acceleration Trade-offs on CO₂ Emissions and Costs of Fuel Economy and Greenhouse Gas Regulations,” *Environmental Science & Technology*, 2017, 51 (18), 10307–10315.
- Yeh, Sonia, Dallas Burtraw, Thomas Sterner, and David Greene**, “Tradable Performance Standards in the Transportation Sector,” *Energy Economics*, 2021, 102, 105490.

A Engineering simulation

Figure 7: *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. From bottom to top the five engines and transmissions are: 1980s carbureted (3-speed), 2007 Toyota PFI (5-speed), 2013 GM GDI (6-speed), 2017 Honda turbo (8-speed), future Ricardo 24 bar EGR (8-speed).

To illustrate the changing relationship between displacement, power, and fuel economy, we present evidence from the EPA’s Advanced Light-Duty Powertrain and Hybrid Analysis (ALPHA) tool (Dekraker, Barba, Moskalik and Butters (2018)), which is a physics-based all-car simulation model. Figure 7, which is reproduced from data in Moskalik, Bolon, Newman and Cherry (2018), summarizes this changing relationship for five simulated cars. Each car uses the same midsize sedan body; the only difference is the drivetrain technology. This allows the simulation to maintain the same road-load across all simulated cars.²⁹ 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

²⁹Road-load is the combined measure of a car’s weight and aerodynamic resistance.

coupled to an eight-speed transmission; and (5) 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 drivetrain technology package using data gathered from real-world cars in a laboratory setting. The figure plots log fuel economy (miles per gallon) against negative log acceleration time (0-60 miles per hour in seconds). Thus, car attributes improve moving up and to the right.

For any given displacement, the later-vintage engines are both faster and more fuel efficient. The gain in acceleration over time is largest for small (low-displacement) engines, while the gain in fuel economy is largest for big (high-displacement) engines (see the connecting dotted lines). Thus, while engines have improved across-the-board, the correspondence between fuel economy and acceleration has flattened. On average, a 1% increase in acceleration is associated with a 0.75% fuel economy loss in the 1980s but only a 0.33% loss in 2017, implying that the opportunity cost of acceleration has decreased.

B Theory appendix

B.1 Marginal Rate of Technical Substitution in Attributes

Here we explore the cost function in equation (3). Define the **marginal rate of technical substitution in attributes (MRTSA)** as the slope of the iso-cost curve in two-dimensional attribute space. The MRTSA measures how much one attribute must change to increase another attribute, while holding production costs constant. The MRTSA is simply the (negative) marginal cost ratio for a given pair of attributes:

$$MRTSA_{ga} \equiv \left. \frac{\partial g}{\partial a} \right|_{\Delta c=0} = -\frac{c_a}{c_g} = \frac{\alpha_a}{\alpha_g} \frac{g}{a} > 0 \quad (30)$$

$$MRTSA_{gs} \equiv \left. \frac{\partial g}{\partial s} \right|_{\Delta c=0} = -\frac{c_s}{c_g} = \frac{\alpha_s}{\alpha_g} \frac{g}{s} > 0 \quad (31)$$

$$MRTSA_{as} \equiv \left. \frac{\partial a}{\partial s} \right|_{\Delta c=0} = -\frac{c_s}{c_a} = -\frac{\alpha_s}{\alpha_a} \frac{a}{s} < 0, \quad (32)$$

where the equalities to the right invoke our specific functional form. Note that $MRTSA_{ga}$ in the first row tells us how much gallons-per-mile (a bad) increases with a marginal increase in acceleration holding costs fixed. Likewise, $MRTSA_{gs}$ in the second row tells us how much gallons-per-mile (a bad) increases with a marginal increase in size holding costs fixed. Conversely, these values tell us how much fuel economy (a good) decreases with a marginal increase in size or acceleration.

Meanwhile, the elasticity versions are given by:

$$\left. \frac{\partial g}{\partial a} \frac{a}{g} \right|_{\Delta c=0} = \frac{c_a}{c_g} \frac{a}{g} = \frac{\alpha_a}{\alpha_g} \quad (33)$$

$$\left. \frac{\partial g}{\partial s} \frac{s}{g} \right|_{\Delta c=0} = \frac{c_s}{c_g} \frac{s}{g} = \frac{\alpha_s}{\alpha_g} \quad (34)$$


$$-\left. \frac{\partial a}{\partial s} \frac{s}{a} \right|_{\Delta c=0} = \frac{c_s}{c_a} \frac{s}{a} = \frac{\alpha_s}{\alpha_a}, \quad (35)$$

where we have scaled the MRTSAs by the corresponding attribute ratios. Note that the elasticity in the first row tells us the percent change in gallons-per-mile for a percent increase in acceleration, holding costs fixed, and likewise for the other two elasticities. The elasticity formulation is convenient because the relationships among these attributes are equivalent when gallons-per-mile is replaced with its inverse (miles-per-gallon, a good). Signs flip but nothing else changes.

B.2 Optimizing conditions and graphical interpretation

Assume for simplicity that the fuel-economy standard is slack ($\tau = 0$). Then divide the left and right sides for each of the pairwise combinations in (5)–(7) to yield the following three equations involving ratios

Figure 8: *Optimal attribute choices for two consumers*



Note: This figure illustrates optimal attribute choices in two-dimensional (fuel economy and size) space as points of tangency between consumers' indifference curves (u 's) with slopes equal to the MRS, and the iso-cost curve in production (c) with slope equal to the MRTSA.

of marginal utility and marginal costs:

$$-\frac{\mu_a}{\mu_g} \frac{\beta_a}{pm} \frac{(g^{-1})^{1+\mu_a}}{a^{1+\mu_a}} = -\frac{\alpha_a}{\alpha_g} \frac{g^{-1}}{a} \quad (36)$$

$$-\frac{\mu_s}{\mu_g} \frac{\beta_s}{pm} \frac{(g^{-1})^{1+\mu_s}}{s^{1+\mu_s}} = -\frac{\alpha_s}{\alpha_g} \frac{g^{-1}}{s} \quad (37)$$

$$-\frac{\mu_a}{\mu_s} \frac{\beta_a}{\beta_s} \frac{s^{1+\mu_s}}{a^{1+\mu_a}} = -\frac{\alpha_a}{\alpha_s} \frac{s}{a}, \quad (38)$$

where we continue to focus on fuel economy (a good) rather than fuel consumption (a bad). These equations say that the marginal rate of substitution (MRS) between any combination of attributes (left side) should equal the corresponding marginal rate of technical substitution in attributes (MRTSA) in production (right side), where the MRTSA is defined as the the slope of the iso-cost curve in attribute space (see appendix section B.1). Figure 8 illustrates this condition for equation (37) as points of tangency between level sets of utility (indifference curves) and level sets of production costs (iso-cost curves) in fuel economy and size

space. Note that $MRTSA_{gs}$ and $MRTSA_{ga}$ on the right side of equations (36) and (37) tells us how much fuel economy decreases with marginal increase in acceleration and size, holding cost fixed. We estimate these attribute trade-offs in elasticity form (see section B.1) in our empirical analysis.

B.3 Second-order sufficient conditions

The second derivatives of the utility function are

$$u_{ss} = -2\frac{\beta_s}{s^3} - \alpha_s(\alpha_s - 1)\frac{c}{s^2} \quad (39)$$

$$u_{aa} = -\alpha_a(\alpha_a + 1)\frac{c}{a^2} \quad (40)$$

$$u_{gg} = -\alpha_g(\alpha_g + 1)\frac{c}{g^2} \quad (41)$$

$$u_{sa} = \alpha_s\alpha_a\frac{c}{sa} \quad (42)$$

$$u_{sg} = \alpha_s\alpha_g\frac{c}{sg} \quad (43)$$

$$u_{ag} = -\alpha_a\alpha_g\frac{c}{ag}. \quad (44)$$

Note at the optimum, $\beta_s = \alpha_s c(s^*)s^*$. Substituting into u_{ss} above,

$$u_{ss} = -\alpha_s(\alpha_s + 1)\frac{c}{s^2} \quad (45)$$

We then find the sign of the determinants of the leading principle minors of the Hessian matrix.

The determinant of the first leading principle minor is negative.

The determinant of the second leading principle minor is

$$D_2 = u_{ss}u_{gg} - u_{sg}^2 \quad (46)$$

$$= \alpha_s\alpha_g\frac{c^2}{s^2g^2}\left((\alpha_s + 1)(\alpha_g + 1) - \alpha_s\alpha_g\right) > 0. \quad (47)$$

The determinant of the third leading principle minor is

$$D_3 = u_{ss}u_{aa}u_{gg} - u_{ss}u_{ag}^2 + 2u_{sa}u_{sg}u_{ag} - u_{sa}^2u_{gg} \quad (48)$$

$$= -\alpha_s\alpha_a\alpha_g\frac{c^3}{s^2a^2g^2}\left(\alpha_s + \alpha_a + \alpha_g + 2\alpha_s\alpha_a\alpha_g + 2\alpha_a\alpha_g\right) < 0. \quad (49)$$

Therefore the Hessian is negative definite and our solutions are maxima.

B.4 Optimal attribute choices

Log-linearizing the necessary conditions in equations (5)–(7) and then solving for the three unknown car attributes yields equilibrium attribute choices as a function of underlying preference and technology

parameters:

$$\ln a \approx \frac{1}{\psi} \left[(\alpha_s + \alpha_g \mu_s + \mu_s) \ln \beta_a - \alpha_s \left(\ln \beta_s + \tau \frac{\hat{\sigma}_s}{\hat{c}_s} \right) - \alpha_g \mu_s \ln(pm + \tau) + \mu_s \theta \right. \quad (50)$$

$$\left. + (\alpha_s + \alpha_g \mu_s + \mu_s)(\ln \mu_a - \ln \alpha_a) + \alpha_s(\ln \alpha_s - \ln \mu_s) + \alpha_g \mu_s \ln \alpha_g - \mu_s \ln k \right]$$

$$\ln s \approx \frac{1}{\psi} \left[(\alpha_a + \alpha_g \mu_a + \mu_a) \left(\ln \beta_s + \tau \frac{\hat{\sigma}_s}{\hat{c}_s} \right) - \alpha_a \ln \beta_a - \alpha_g \mu_a \ln(pm + \tau) + \mu_a \theta \right. \quad (51)$$

$$\left. + (\alpha_a + \alpha_g \mu_a + \mu_a)(\ln \mu_s - \ln \alpha_s) + \alpha_a(\ln \alpha_a - \ln \mu_a) + \alpha_g \mu_a \ln \alpha_g - \mu_a \ln k \right]$$

$$\ln g^{-1} \approx \frac{1}{\psi} \left[(\alpha_s \mu_a + \alpha_a \mu_s + \mu_s \mu_a) \ln(pm + \tau) - \alpha_s \mu_a \left(\ln \beta_s + \tau \frac{\hat{\sigma}_s}{\hat{c}_s} \right) - \alpha_a \mu_s \ln \beta_a + \mu_a \mu_s \theta \right. \quad (52)$$


$$\left. - (\alpha_s \mu_a + \alpha_a \mu_s + \mu_a \mu_s) \ln \alpha_g + \alpha_s \mu_a (\ln \alpha_s - \ln \mu_s) + \alpha_a \mu_s (\ln \alpha_a - \ln \mu_a) - \mu_a \mu_s \ln k \right],$$

where $\psi \equiv \alpha_s \mu_a + \alpha_a \mu_s + \mu_s \mu_a + \alpha_g \mu_s \mu_a$ and we here model fuel economy (a good) rather than fuel consumption (a bad). The terms $\hat{\sigma}_s$ and $\hat{c}_s$ are the partial derivatives of the fuel economy standard and costs with respect to car size, evaluated at equilibrium attributes. Thus, these equations hold exactly at equilibrium attribute values (given by the solution to the nonlinear equations 5–7), and can be used to calculate approximate changes in equilibrium attributes due to changes in gas prices, policy, and other factors. These changes are exact when fuel economy constraint is slack ($\tau = 0$). In practice, even when the constraint is binding ($\tau > 0$), the implied incentive to boost car size given by $\tau \sigma_s / c_s$ is so small relative to $\ln \beta_s$ that it has a trivial impact on equilibrium attribute levels.

Note that heterogeneity in gasoline prices and the preference parameters (β_a, β_s, pm) 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 for *all* attributes, along with the cost parameters that govern technical trade-offs. Consumers with stronger preferences in any one dimension will therefore choose different car attributes in all three dimensions.

To illustrate, suppose that the preference parameters on acceleration ($\ln \beta_a$) and size ($\ln \beta_s$) are identical across consumer types such that variation in equilibrium car attribute bundles derives solely from heterogeneity in the preference parameters on gallons-per-mile ($\ln pm$), with a slack fuel economy constraint ($\tau = 0$). Inspection of equations (50)–(52) reveals that consumers with stronger preferences to save fuel (larger $\ln pm$) will choose smaller, slower, and more efficient cars. Further, because equations (50)–(52) 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, which is consistent with preference heterogeneity in multiple dimensions.

Figure 9: *Biased technical change*



Note: This figure illustrates the effects of biased technical change in two-dimensional (fuel economy and acceleration) space. The figure shows an optimal attribute bundle as the point of tangency between a consumer's indifference curve (u) and an iso-cost curve prior to technical change ($c(\eta = 0)$). In the presence of technical change ($\eta > 0$), the new iso-cost curve passing through this bundle ($c(\eta > 0)$) is flatter.

B.5 Attribute-biased technical change

Figure 9 illustrates the effect of attribute-biased technical change as described in main text section 3.5.2. As in section B.2 above, the figure illustrates a consumer's optimal attribute bundle as a point of tangency between her indifference curve (u) and an iso-cost curve (c) prior to technical change ($\eta = 0$). In the presence of biased technical change ($\eta > 0$), the new iso-cost curve passing through this attribute bundle is flatter. The $MRTSA_{ga}$ has decreased, implying that the opportunity cost of acceleration is now lower. Thus, the consumer now strictly prefers attribute bundles to the immediate southeast of the initial bundle: lower fuel economy and higher acceleration. As noted in the main text, this form of technical change will interact with gasoline prices and efficiency standards, making fuel economy less responsive relative to size and acceleration.

C Accounting for markups

Let M denote the profit margin above cost for each vehicle produced. The consumer's problem becomes

$$\max_{g,a,s} u = y + v(g, a, s) - (1 + M) \times c(g, a, s). \quad (53)$$

The first order condition for g is then

$$[g] : \frac{\partial v}{\partial g} = (1 + M) \frac{\partial c}{\partial g}. \quad (54)$$

Solving for g^{-1*} and taking logs, we see that $(1 + M)$ is linearly separable and is relegated to the error of our regression model.

It is reasonable to assume that the gas price shocks in our data are not correlated with M . Thus, M should not cause our estimates of α_g to be biased. For regressions that include demographic controls (D) and make and pickup dummy controls (X), the identifying assumption for unbiased $\hat{\alpha}_a$, $\hat{\alpha}_s$, and $\hat{\theta}$ is

$$\mathbf{E}[M|D, X, a, s, t] = \mathbf{E}[M|D, X]. \quad (55)$$


This assumption holds trivially if M is uncorrelated with a , s , and t , which implies the mean of the markup multiplier M is fixed over time and across acceleration and size levels. As long as D and X controls absorb correlations between M and a , s , and t , as would be the case if luxury brands have higher markups or if demographic variables are correlated with models within brands with higher markups, then omitted M does not bias our estimates. We can further relax the assumption if we only wish to identify technical change. The assumption for identifying technical change is simply

$$\mathbf{E}[M|D, X, s, a, t] = \mathbf{E}[M|D, X, a, s]. \quad (56)$$

That is, if the bias of $\hat{\gamma}_{at}$ and $\hat{\gamma}_{st}$ is constant, then changes in the coefficient over time identify changes in the cost parameters. For the purposes of simulation, we do not need to be too concerned about this bias since we are modeling consumer choices and the bias reflects changes in car prices, over manufacturing costs, that caused by changes in attributes.


D Tables and figures

Figure 10: *Trends in interest rates on new-car loans*



Note: This figure plots the trend in nominal and real interest rates on 48-month new auto loans using data from Federal Reserve Bank of St. Louis.

Figure 11: *Estimated shadow cost for efficiency standards*



Note: The top figure shows the estimated shadow cost and the statutory penalty for fuel economy and GHG standards from the literature, expressed in dollars per MPG (Yeh et al., 2021). The bottom figure expresses this shadow cost as a fraction of lifetime fuel savings from a 1MPG improvement ($\tau/1mpg$) assuming 100,000 discounted lifetime miles and the mean gasoline price and fuel economy for each year. We make three alternative assumptions for the pre-2007 period. High assumes a shadow price equal to the non-compliance penalty of \$55. Mid assumes the mean value reported in ?. Low assumes zero.

Table 7: *First-stage regression results for table 2*

	(4)	(4)	(5)	(5)	(6)	(6)	(7)	(8)
	lnHPWt	lnVol	lnHPWt	lnVol	lnHPWt	lnVol	lnHPWt	lnVol
Male share	0.033*** (0.004)	0.004 (0.004)	0.034*** (0.004)	0.003 (0.004)	0.029*** (0.005)	0.004 (0.004)	0.029*** (0.005)	0.004 (0.004)
# household members	0.006*** (0.001)	0.013*** (0.001)	0.006*** (0.001)	0.012*** (0.001)	0.009*** (0.001)	0.019*** (0.002)	0.009*** (0.001)	0.019*** (0.002)
F-stat	39.022	98.433	41.757	101.469	36.457	83.576	24.938	55.662
P-value	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]

Note: This table presents first-stage results for the 2SLS regressions reported in table 2: $\log(\text{HPWt})$ and $\log(\text{volume})$ regressed on the excluded instruments and exogenous controls. $N = 245,277$. Column numbers here match those in the main results table. We only report coefficients on the excluded instruments: male share and # household members. Standard errors in parentheses are clustered by state. Note *, **, and *** indicate statistical significance at the 5%, 1%, and 0.1% levels respectively. The bottom rows report standard F-tests for the joint significance of the excluded instruments robust to clustering by state.

Table 8: *Fuel economy conditional on attributes: robustness*

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Main (6)	Gas Lags	Main (7)	Brand	Overlap	Weights	Shadow
Log gas price (γ_p)	0.093* (0.042)		0.103* (0.047)	0.082* (0.039)	0.069 (0.046)	0.100* (0.045)	0.109* (0.048)
Log gas price (lag 0)		0.013 (0.023)					
Log gas price (lag 1)		0.063* (0.025)					
Log gas price (lag 2)		0.018 (0.028)					
Log gas price (lag 3)		-0.018 (0.032)					
Log gas price (lag 4)		0.016 (0.033)					
Log HP/weight (γ_a)	-0.641*** (0.124)	-0.641*** (0.124)	-0.643*** (0.124)	-0.517*** (0.150)	-0.594*** (0.105)	-0.612*** (0.051)	-0.642*** (0.123)
Log volume (γ_s)	-0.605*** (0.070)	-0.605*** (0.070)	-0.607*** (0.070)	-0.684*** (0.133)	-0.642*** (0.067)	-0.697*** (0.033)	-0.604*** (0.069)
Pickup	-0.028** (0.009)	-0.028** (0.009)	-0.028** (0.009)	-0.035* (0.014)	-0.028** (0.010)	-0.019** (0.007)	-0.028** (0.009)
Sum gas price (γ_g)		0.092 (0.060)					
$\frac{\partial \ln \text{mpg}}{\partial \ln \text{HP/Wt}} _{\Delta c=0} \left(\frac{\alpha_a}{\alpha_g} = \frac{\gamma_a}{1-\gamma_g} \right)$	-0.707 (0.140)	-0.706 (0.146)	-0.716 (0.143)	-0.563 (0.166)	-0.637 (0.121)	-0.680 (0.067)	-0.720 (0.143)
$\frac{\partial \ln \text{mpg}}{\partial \ln \text{volume}} _{\Delta c=0} \left(\frac{\alpha_s}{\alpha_g} = \frac{\gamma_s}{1-\gamma_g} \right)$	-0.668 (0.089)	-0.667 (0.095)	-0.677 (0.093)	-0.745 (0.152)	-0.689 (0.083)	-0.774 (0.055)	-0.678 (0.093)


Note: This table presents coefficient estimates from equation (27). N = 245,277. Column (1) reproduces model (6) from table 2 that instruments for size and acceleration and includes full demographics but does not adjust gas prices for the shadow cost of efficiency standards. Column (2) replaces the log three-year moving average of gas price with current and lagged log gas prices. Column (3) reproduces model (7) from table 2. Relative to this model, column (4) adds brand fixed effects, column (5) omits all but one overlapping model year across NHTS waves (2001, 2009, and 2017), column (6) omits NHTS survey weights, and column (7) adjusts log gas prices using the “high” shadow cost value in place of the “mid” value. Standard errors in parentheses are clustered by state. Note *, **, and *** indicate statistical significance at the 5%, 1%, and 0.1% levels respectively.

Table 9: *First-stage regression results for table 4*

	(2)	(2)	(3)	(3)
	CostHPWt	CostVolume	CostHPWt	CostVolume
Log gas price	-0.391 (0.642)	-1.019* (0.467)	-0.374 (0.653)	-0.969* (0.479)
Log odometer	0.229*** (0.065)	0.038 (0.086)	0.260*** (0.068)	0.069 (0.085)
Log odometer ²	-0.009* (0.004)	-0.002 (0.005)	-0.011** (0.004)	-0.004 (0.005)
Log odometer $\times$ age	-0.012*** (0.004)	0.013*** (0.003)	-0.012*** (0.004)	0.013*** (0.003)
Odometer missing	1.463*** (0.308)	0.089 (0.390)	1.555*** (0.318)	0.187 (0.387)
Odometer missing $\times$ age	-0.139*** (0.040)	0.158*** (0.033)	-0.140*** (0.039)	0.155*** (0.033)
F-stat	53.653	60.372	64.500	44.905
P-value	(0.000)	(0.000)	(0.000)	(0.000)

Note: This table presents first-stage results for the 2SLS regressions reported in table 4: marginal cost shifters regressed on the excluded instruments and exogenous controls. $N = 245,277$. Column numbers here match those in the main results table. We only report coefficients on the excluded instruments: log gas price and the odometer-related variables (log, log squared, log $\times$ car age, missing odometer, and missing $\times$ age). Standard errors in parentheses are clustered by state (conditional on the estimated cost parameters). Note *, **, and *** indicate statistical significance at the 5%, 1%, and 0.1% levels respectively. The bottom row reports standard F-tests for the joint significance of the excluded instruments robust to clustering by state (conditional on the estimated cost parameters).

Figure 12: *Consumer preferences*



Note: The top figure plots estimated preference shifters for acceleration ($\ln \beta_a$) and size ($\ln \beta_s$) for every car purchase in our final estimation sample. The bottom figure plots coefficients from OLS regressions of these preference shifters on dummy variables for household income, population density, and # household members, along with continuous male share using car purchases from 2017. The hollow circles show point estimates for acceleration (dependent variable $\ln \beta_a$) while the hollow squares show point estimates for volume (dependent variable $\ln \beta_s$). Coefficient estimates are relative to the excluded categories of income < \$35,000, density < 100 people per square mile, and one household member. The 95% confidence intervals (horizontal bars) are based on standard errors clustered by state.