Attribute Production and Biased Technical Change in Automobiles*

Asa Watten† Soren Anderson‡ Gloria Helfand§

November 20, 2023

PRELIMINARY: PLEASE DO NOT CITE OR DISTRIBUTE

Latest version (link)

Abstract

Cars have simultaneously grown larger, faster, and more efficient over the last 30 years. What explains these trends? We develop and estimate an equilibrium model of car attribute production using U.S. data for 1995–2017. We structurally decompose attribute trends into underlying effects of gas prices, technical change, and consumer preferences. We emphasize the pivotal role of efficiency-biased technical change, which drove efficiency gains and made size and speed less sensitive to gas prices. However, sharply increasing preferences for car size pushed strongly in the opposite direction. Meanwhile, fuel economy and GHG standards had almost no effect on efficiency, indicating that technical change would have allowed much more stringent standards.

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 the University of Alaska Anchorage, University of California Davis, and National Bureau of Economic Research (NBER) Economics of Energy Use in Transportation Workshop in Spring 2023 for their helpful comments.

†Yale University, 195 Prospect Street, New Haven, CT 06511, asa.watten@yale.edu

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

§University of Michigan, 440 Church Street Ann Arbor, MI 48109, ghelfand@umich.edu

1 Introduction

Understanding the production and consumption of cars, houses, and other energy-using durables is essential for climate policy. The durability of these goods implies that today's choices have long-run consequences, with considerable climate impacts: buildings and transportation each account for 30% of U.S. greenhouse gas emissions (EPA 2020). However, the multi-attribute nature of these goods complicates both policy and quantitative analysis. Energy efficiency is physically codetermined with other attributes. This fact raises a persistent concern in public discourse and the economics literature that policies to promote energy efficiency 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, Fowlie, and Skerlos 2017; Leard, Linn, and Zhou 2023). For example, improvements in home insulation may lead to larger homes, while falling costs for hybrid-electric engines may lead to faster cars. Predicting such outcomes is difficult, however, since doing so requires information on both technological trade-offs and consumer preferences.

Figure 1: *Attributes for new cars 1995–2017*

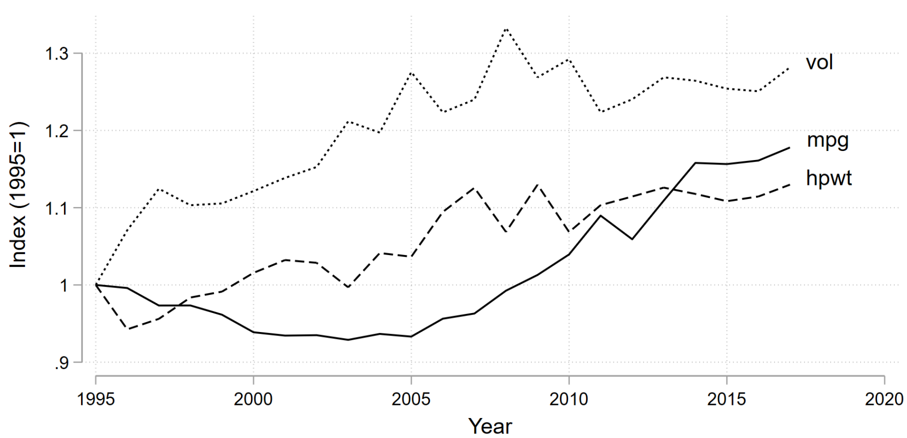

Note: Ratio changes in mean car attributes (sales-weighted) since 1994. The horsepower-to-weight ratio (HP/weight) is a close proxy for acceleration capacity.

The auto industry's history validates these concerns. 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 standards would force consumers into choosing smaller, less powerful cars that are less enjoyable and less safe to drive. Falling oil prices in the 1990s then coincided with increases in size and acceleration and decreases in fuel economy, reinforcing this belief. Nonetheless, car performance continued to improve in the 2000s with the adoption of new engine technologies, even under rising oil prices, while size plateaued and fuel economy actually increased—a pattern that continued into the 2010s. See figure 1, which shows trends in new car attributes and proxies for acceleration with the horsepower-to-weight ratio.

We study these issues 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 into the underlying effects of gas prices, consumer preferences, and changing technology. We show that shifts in car technology and consumer preferences are both pivotal to understanding how the car fleet has evolved in recent decades. We show that technological change was substantially biased, reducing overall costs while making fuel economy cheaper relative to size and acceleration. However, car buyers did not respond simply by buying more fuel economy. Instead, they reallocated some of the technological gains toward larger and faster cars. Simultaneously, preferences for size and acceleration increased sharply over two decades, further eroding the potential gains in fuel economy. To our knowledge, ours is the first paper to emphasize the importance of biased technical change on these outcomes. Biased technical change helps explain why the observed shadow cost of fuel economy and greenhouse gas (GHG) standards was so low and provides lessons for future policies (e.g., electric car subsidies) that aim to drive the direction of innovation.

Our analysis proceeds in three steps. We begin by developing a theory of car attribute production that accounts for 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.[1] On the demand side, we assume heterogeneous preferences for car attributes across consumer types, leading to an equilibrium distribution of differentiated car models. We solve for a consumer's optimal choice of fuel economy conditional on other car attributes (as in Knittel 2011) and 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 examine the diffusion of specific engine technologies (e.g., turbo-charging) across the car fleet, testing whether they are first adopted among larger, 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 1980–2017, along with data on car attributes from the Environmental Protection Agency (EPA) and detailed data on car size from the Canadian government. Finally, we use our model and a structural decomposition to quantify how rising gas prices, technical change, and shifting preferences have interacted to produce observed trends in car size, acceleration, and fuel economy.

Our analysis yields five main findings. First, attribute-neutral technical change, which reduces the cost of

[1] We show that the assumption of perfect competition is not necessary to identify the parameters of interest as long as markups follow reasonable constraints.

all attributes equally, leads to larger, faster, and more efficient cars. We estimate attribute-neutral technical change equivalent to a 1.5% annual increase in all car attributes holding drivetrain costs fixed from 1995 to 2017. We show that controlling for gas prices is essential both to quantify technical change over time and to recover 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 (e.g., turbochargers for acceleration or hybrid engines for fuel economy), in contrast to incremental changes or process improvements. We show theoretically that the earliest adopters are consumers choosing larger, faster, and more efficient cars. We confirm this prediction empirically for turbochargers, direct fuel injection, continuously variable transmissions, and hybrid-electric drivetrains. Our theory implies that the efficiency gains from an efficiency-boosting technology are partially reallocated to size and speed. Thus, the environmental impact of subsidies for hybrid-electric drivetrains and other engine technologies may be much less than what policymakers intend.

Third, we find strong evidence of biased technical change favoring fuel economy. In particular, we estimate that a 10% reduction in acceleration led to a 2.2% gain in fuel economy in the late 1990s but a 6.7% gain in the 2010s, with costs held fixed. Likewise, a 10% reduction in car size led to a 3.6% gain in fuel economy in the 1980s but a 5.7% gain in the 2010s. Thus, the opportunity cost of size and speed have both risen in terms of forgone fuel economy. We show theoretically that this form of technical change leads to smaller, slower, and more efficient cars. Our simulations imply that cars in 2017 would have been 4% larger, 21% faster, and 9% less efficient had these shifts in technology not occurred. For comparison, we estimate that cars in 2017 would have been 4% larger, 4% faster, and 7% less efficient had gas prices not doubled over this period.

Fourth, biased technical change makes fuel economy more responsive than size and acceleration to gas taxes and fuel economy standards. Thus, achieving efficiency gains through technology alone, rather than reducing size and acceleration, becomes relatively more cost-effective. Indeed, our results imply that a 10% increase in gas prices led to a 1.3% gain in fuel economy in the late 1990s and a 2.4% gain in the 2010s. Meanwhile, for every 10% increase in fuel economy induced by higher gas prices or tighter standards, size and acceleration fell by 11% and 15% in the late 1990s but by only 5% and 7% in the 2010s. These results may help explain why car performance fell sharply in the 1980s under high gas prices and tighter standards but rose throughout the 2010s when gas prices were also high: small reductions in car size and acceleration now yield much larger fuel economy gains, making these attributes less sensitive to shifts in the demand for fuel economy.

Finally, the effects of technical change, energy prices, and fuel economy standards are all mediated by consumer preferences for size and acceleration, and shifts in consumer preferences can undermine policies aimed at reducing carbon emissions. We empirically identify the elasticities of demand for size and acceleration, estimating them to be -0.15 and -0.18. To our knowledge, we are the first to estimate these elasticities; most papers estimate discrete-choice models of car demand assuming constant marginal willingness-to-pay for car attributes. Our simulations imply that cars in 2017 would have been 28% smaller, 30% slower, and 52% more efficient had size and acceleration preferences remained fixed at their 1995 levels. Larger cars increase the risk of death and injury to occupants of other vehicles, which is an important untaxed externality (see Jacobsen 2013; Anderson and Auffhammer 2014; Bento, Gillingham, and Roth 2017). Ironically, surging preferences for larger size may derive from an arms race for safety as drivers seek to protect themselves against other drivers. Further, our simulations indicate a small effect of fuel economy and GHG standards relative to the effects of technical change, gas prices, and preferences. This indicates that standards could have been much more stringent for a modest 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. The seminal paper in this literature is Newell, Jaffe, and Stavins (1999), which estimates shifts over time in the relationship between product cost and energy efficiency for air conditioners and water heaters and relates this technical change to energy prices and policy. We contribute to this literature by providing theoretical microfoundations, a new set of technical terms, and new empirical methods to better understand technical change for product attributes, including energy efficiency. Such theory is key to understanding the effects of regulation on equilibrium attributes, welfare, and incidence. Empirically, we show how controlling for energy prices can compensate for a lack of accurate cost data. Our application to cars reveals attribute-biased technical change favoring fuel economy in recent decades. Additionally, we compile detailed trim-level attribute data (1994–2020) from multiple government sources. These data, which we will make freely available for public use, will be of substantial interest to researchers working on a range of car-related topics.

Our paper is closely related to that of Knittel (2011), who also uses market data to understand trade-offs among car attributes. Dozens of papers cite Knittel (2011) as the basis for their understanding of technical change and attribute trade-offs in the car industry (for example, MacKenzie and Heywood 2012; Klier and Linn 2015; Whitefoot, Fowlie, and Skerlos 2017). We provide explicit microfoundations for the empirical specification in Knittel (2011) and extend this model in several key dimensions, 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 isocost relationships or the rate of technical change. However, we show how to recover these structural parameters by using household-level microdata and controlling for state-level gas prices. We find that isocost curves are steeper and the average rate of technical change is higher than implied by the reduced-form estimates. In addition, we show that costly engine technologies are not applied at random. Indeed, our theory predicts and our empirical results confirm that all of the engine technologies considered in Knittel (2011) are first applied to larger, faster, and more efficient cars. Finally, we allow technical change to move in an attribute-neutral or attribute-biased direction. We find strong evidence of biased technical change, leading to a more optimistic assessment that it is possible to meet stringent efficiency standards without large reductions in size and acceleration.

Many recent economic and environmental policies seek to direct technical change.[2] While existing literature studies how directed technical change affects the factors of production (e.g., Acemoglu 2002a,b) or explores how taxes and subsidies stimulate technical change (Aghion et al. 2016; Calel and Dechezleprêtre 2016), we know little about directed technical change in markets for multi-attribute goods, such as cars and houses. Our setting connects closely with that of Aghion et al. (2016), who study directed technical change in the auto industry for electric, hybrid, and hydrogen engines. They show that patents for these technologies increase during episodes of high gas prices. Similarly to Aghion et al. (2016), we show theoretically that the incentive to innovate increases with higher gas prices. However, we go beyond Aghion et al. (2016) to explore how directed technical change interacts with gas prices 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 constant marginal willingness-to-pay for size and acceleration, making such models unsuitable for long-run analyses.[3] We provide a structure that allows us to disentangle supply and demand. However, we abstract 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. Similarly to us, they assume a perfectly competitive car industry with zero cost of introducing new car models, such that every consumer type obtains a car model optimized to its own preferences. The authors 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

[2] Examples for the US include the CHIPS and Science Act of 2022, the Inflation Reduction Act, and the EPA's "Light- and Medium-Duty Proposed Standards for Model Years 2027 and Later"; examples for Europe include the European Union Recovery Instrument and the European Green Deal; and among academic economists, related works include, e.g., Acemoglu, Aghion, Bursztyn, and Hemous (2012) and Aghion, Dechezleprêtre, Hemous, Martin, and Van Reenen (2016).

[3] Whitefoot and Skerlos (2012) and Lin and Linn (2023) make car attributes endogenous but do not consider long-run technical change.

key observables. We show that this model provides a rich set of predictions that align closely with empirical evidence.

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 size, acceleration, and fuel economy. Section 4 leverages this theory to estimate production cost parameters 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 Physical determinants of car attributes

Carmakers modify major car attributes—size, power, and fuel economy—by adding or subtracting discrete technologies, marginally improving drivetrain components, choosing different materials, increasing or decreasing engine size, and tuning the engine to achieve different goals. Table 1 illustrates how individual technologies and design choices affect size, acceleration, and fuel economy. The table illustrates five categories of technologies that can be added to a vehicle: (1) technologies that improve only acceleration (turbochargers and superchargers); (2) technologies that improve only fuel economy (engine stop-start); (3) technologies that improve both fuel economy and acceleration (engine efficiency, light-weighting, aerodynamics, and advanced transmissions)[4]; (4) technologies that improve acceleration at the expense of fuel economy (engine displacement and tuning); and (5) increases to size and therefore weight, which lowers both fuel economy and acceleration, absenct any other changes to the drivetrain. We created this table in discussion with automotive engineers at the U.S. EPA 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 that 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 larger 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, changing attributes while holding costs fixed entails adding some technology and removing others or decreasing engine size and adding some performance-enhancing technology.

USEPA (2019) observes that the relationship between engine displacement, horsepower (the propulsive force coming from the engine), and fuel economy has substantially shifted over time. A typical modern

[4]Transmission 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
ICE displacement	+	–	0
size (not tech.)	–	–	+

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.

engine is much more powerful than was an engine of the same size 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, carmakers today can add more power with relatively smaller changes in engine size than in the past. The combined effect is a trend toward faster, more efficient engines that can be made faster still by sacrificing relatively little fuel economy. Figure 9 in the appendix illustrates these shifts. It reproduces EPA engineering simulations of representative drivetrain technologies over time and shows that engines have improved in both the dimensions of acceleration and fuel economy space while the slope has flattened out. Some have taken this to mean that there is no longer a trade-off between power and fuel economy.

While these trends are compelling, they do not tell us much about the opportunity cost of improving fuel economy or acceleration. We need two pieces of information to trace out an isocost curve: (1) the cost of adding technology to marginally increase power or fuel economy and (2) the cost savings from reducing engine size. If the cost of adding technology is large or the cost savings from reducing engine size is small, then we could be stuck at a frontier where we are unable to improve efficiency without enormous reductions to power. On the other hand, if the cost of adding technology is small or the cost savings from reducing engine size is large, we can shed relatively little horsepower while gaining efficiency and keeping costs fixed. Understanding why drivers buy the mix of attributes that we observe and how they respond to policy, fuel costs, and technological shifts requires additional information on driver preferences for attributes.

Unfortunately, it is not easy to observe detailed production costs, especially over the long run, or consumer preferences. In the sections that follow, we develop a theory that allows us to estimate how technological changes both have improved vehicles and have altered the opportunity cost of attributes.

3 Theory of attribute production

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

3.1 Modeling optimal attribute choices

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

3.1.1 Model setup

We assume a finite number J of consumer types, indexed by j, which have heterogeneous preferences for car attributes; we suppress the type j subscripts at times to keep the notation tidy. We explicitly model three car attributes: size (s), acceleration (a), 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 her continuous choice of car attributes and spending on other goods. Utility is given by

$$u(g, a, s) = v(s, a) + y - \beta_g pmg - h(g, a, s), \tag{1}$$

where $v(s, a)$ is utility derived from size and acceleration; y is income; pgm is lifetime fuel expenditures, which equal the price of gas (p) times fuel consumption (g) times lifetime miles (m); β_g is the marginal utility for expenditures such that $\beta_g = 1$ implies full valuation of discounted lifetime fuel savings and $\beta_g < 1$ implies undervaluation;[5] 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 fuel expenditures and car price. We assume that the number of lifetime miles is exogenous, consistent with empirical evidence that demand for miles is inelastic; relaxing this assumption would not substantially change our conclusions. However, miles and fuel prices may vary 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. Larger cars are heavier, which slows them down. Larger engines can compensate but use more fuel. Hybrid engines

[5] This allows for undervaluation of fuel economy due to myopia, credit constraints, or other factors and valuation above pecuniary benefits due to prosocial behavior.

and other technologies can be added to save fuel or boost acceleration but cost money. We assume that carmakers are price takers.[6] We further assume that there is free entry and zero fixed cost to develop and produce a new car model. Taken together, these assumptions imply that, in equilibrium, every consumer type chooses a car model custom-tailored to its own unique preferences, with equilibrium prices equal to costs: $h(g, a, s) = c(g, a, s)$. Note that the assumption of zero fixed costs will hold approximately as 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 as a function of each car's footprint (f). 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 standard: $\tau[g - \sigma(s)]$. Note that this tax becomes an implicit subsidy if the car is more efficient than the standard: $g < \sigma$.[7] If the fuel economy standard fails to bind in equilibrium, then the credit price is zero ($\tau = 0$), and these implicit taxes and subsidies vanish.[8]

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

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

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:

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

where k is a common cost shifter, θ is an index of attribute-neutral technical progress that scales production costs downward, and parameters $\alpha_s, \alpha_a, \alpha_g > 1$ relate percent changes in costs to percent changes in attributes.[9] Given our functional form, costs are rising in size and acceleration ($c_s, c_a > 0$) and falling in gallons-per-mile ($c_g < 0$). Thus, making a car more desirable raises production costs. Costs are convex with

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

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

[8] Note that, before 2011, there was a separate standard for cars and light trucks but σ did 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.

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

respect to each of these attributes ($c_{ss}, c_{aa}, c_{gg} > 0$). Thus, it is increasingly costly to make a car larger, faster, or more efficient. Finally, the cross-partials imply that making a car more efficient is costlier when the car is large or fast ($c_{ag}, c_{sg} < 0$) and, likewise, that making a car faster is costlier when the car is large ($c_{as} > 0$). To summarize succinctly: making a car better in any dimension is costlier when the car is already better in any dimension. Note that the elasticities of marginal cost with respect to size, acceleration, and fuel economy are $1 + \alpha_s$, $1 + \alpha_a$, 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:

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

where μ_s and μ_a capture the declining marginal benefits of size and acceleration while β_s and β_a are preference shifters.[10] Note that the marginal utility of size is $\beta_s s^{-(1+\mu_s)}$. Thus, the elasticity of marginal utility with respect to size is $-(1 + \mu_s)$, and likewise for acceleration. In our empirical application, we assume a uniform elasticity across consumers but allow for heterogeneity in the demand shifters (β_s, β_a).

Before moving to our necessary conditions, we pause briefly to explore the cost function in equation (3). Define the **marginal rate of technical substitution of attributes (MRTSA)** as the slope of the isocost curve in two-dimensional attribute space. The MRTSA measures how much one attribute must change to increase another attribute, with production costs held constant. The MRTSA is simply the (negative) marginal cost ratio for a given pair of attributes:

$$MRTSA_{gs} \equiv \left.\frac{\partial g}{\partial s}\right|_{\Delta c=0} = -\frac{c_s}{c_g} = \frac{\alpha_s}{\alpha_g}\frac{g}{s} > 0 \tag{5}$$

$$MRTSA_{ga} \equiv \left.\frac{\partial g}{\partial a}\right|_{\Delta c=0} = -\frac{c_a}{c_g} = \frac{\alpha_a}{\alpha_g}\frac{g}{a} > 0 \tag{6}$$

$$MRTSA_{as} \equiv \left.\frac{\partial a}{\partial s}\right|_{\Delta c=0} = -\frac{c_s}{c_a} = -\frac{\alpha_s}{\alpha_a}\frac{a}{s} < 0, \tag{7}$$

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

[10] Given that size and acceleration enter the Cobb–Douglas-like cost function with positive coefficients, this formulation ensures that level sets for preferences are well behaved with respect to level sets for costs, i.e., 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.

Meanwhile, the elasticity versions are given by:

$$\left.\frac{\partial g}{\partial s}\frac{s}{g}\right|_{\Delta c=0} = \frac{c_s}{c_g}\frac{s}{g} = \frac{\alpha_s}{\alpha_g} \tag{8}$$

$$\left.\frac{\partial g}{\partial a}\frac{a}{g}\right|_{\Delta c=0} = \frac{c_a}{c_g}\frac{a}{g} = \frac{\alpha_a}{\alpha_g} \tag{9}$$

$$\left.-\frac{\partial a}{\partial s}\frac{s}{a}\right|_{\Delta c=0} = \frac{c_s}{c_a}\frac{s}{a} = \frac{\alpha_s}{\alpha_a}, \tag{10}$$

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 size, with costs held 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). The signs flip, but nothing else changes.

3.1.3 Necessary conditions

Maximization of equation (4) with respect to each of the continuous car attributes implies the following necessary conditions:

$$s\colon\ \frac{\mu_s\beta_s s^{-\mu_s}}{s} = \frac{\alpha_s c(g,a,s)}{s} - \sigma_s\tau \tag{11}$$

$$a\colon\ \frac{\mu_a\beta_a a^{-\mu_a}}{a} = \frac{\alpha_a c(g,a,s)}{a} \tag{12}$$

$$g^{-1}\colon\ \frac{\beta_g pm(g^{-1})^{-1}}{g^{-1}} = \frac{\alpha_g c(g,a,s)}{g^{-1}} + \tau. \tag{13}$$

Here, note that we have maximized with respect to fuel economy (g^{-1}) rather than fuel consumption (g) in (13) to emphasize the symmetry among the three necessary conditions. The equivalent equation for g is

$$g\colon\ \beta_g pm = \frac{\alpha_g c(g,a,s)}{g} + \tau. \tag{14}$$

These conditions say that the marginal benefit from improving a car attribute (left side) should equal the marginal cost (right side). σ_s is the change in the standard for a small change in size. Second-order sufficient conditions are shown in the appendix.

Our empirical strategy below hinges on estimation of equation (13), which relates the demand for fuel economy (gas prices) to its marginal cost. We show that a regression of log fuel economy on log gas prices, log size, and log acceleration identifies the cost parameter on fuel economy (α_g), along with the other cost parameters (α_s, α_a) and attribute-neutral technical change (θ). Having estimated these parameters, we then show how to identify the preference parameters on size and acceleration (βs and μs) using the other two necessary conditions.

Now divide the left and right sides for each of the pairwise combinations of (11)–(13) to yield the following three equations involving ratios of marginal utility and marginal costs:

$$-\frac{\mu_s}{\mu_g}\frac{\beta_s}{\beta_g}\frac{(g^{-1})^{1+1}}{s^{1+\mu_s}} = -\frac{\alpha_s}{\alpha_g}\frac{g^{-1}}{s} \tag{15}$$

$$-\frac{\mu_a}{\mu_g}\frac{\beta_a}{\beta_g}\frac{(g^{-1})^{1+1}}{a^{1+\mu_a}} = -\frac{\alpha_a}{\alpha_g}\frac{g^{-1}}{a} \tag{16}$$

$$-\frac{\mu_a}{\mu_s}\frac{\beta_a}{\beta_s}\frac{s^{1+\mu_s}}{a^{1+\mu_a}} = -\frac{\alpha_a}{\alpha_s}\frac{s}{a}, \tag{17}$$

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 should equal the corresponding marginal rate of technical substitution of attributes (MRTSA) in production. Figure 2 illustrates this condition for equation (15) as points of tangency between level sets of utility (indifference curves) and level sets of production costs (isocost curves) in the fuel economy–size space.

3.1.4 Optimal attribute choices

Solving equations (11)–(13) for the three unknown car attributes and taking logs then yields equilibrium attribute choices as a function of the underlying preference and technology parameters:

$$\begin{aligned}\ln s = \frac{1}{\psi}\Big[&(\alpha_a + \mu_a + \alpha_g\mu_a)\ln\beta_s - \alpha_a\ln\beta_a - \alpha_g\mu_a\ln(\beta_g pm + \tau) + \mu_a\theta \\ &+ (\alpha_a + \alpha_g\mu_a + \mu_a)(\ln\mu_s - \ln\alpha_s) + \alpha_a(\ln\alpha_a - \ln\mu_a) + \alpha_g\mu_a\ln\alpha_g - \mu_a\ln k + \tau\frac{\sigma_s}{c_s}\Big]\end{aligned} \tag{18}$$

$$\begin{aligned}\ln a = \frac{1}{\psi}\Big[&(\alpha_s + \mu_s + \alpha_g\mu_s)\ln\beta_a - \alpha_s\ln\beta_s - \alpha_g\mu_s\ln(\beta_g pm + \tau) + \mu_s\theta \\ &+ (\alpha_s + \alpha_g\mu_s + \mu_s)(\ln\mu_a - \ln\alpha_a) + \alpha_s(\ln\alpha_s - \ln\mu_s) + \alpha_g\mu_s\ln\alpha_g - \mu_s\ln k\Big]\end{aligned} \tag{19}$$

$$\begin{aligned}\ln g^{-1} = \frac{1}{\psi}\Big[&(\alpha_s\mu_a + \alpha_a\mu_s + \mu_s\mu_a)\ln(\beta_g pm + \tau) - \alpha_s\mu_a\ln\beta_s - \alpha_a\mu_s\ln\beta_a + \mu_a\mu_s\theta \\ &- (\alpha_s\mu_a + \alpha_a\mu_s + \mu_a\mu_s)\ln\alpha_g + \alpha_s\mu_a(\ln\alpha_s - \ln\mu_s) + \alpha_a\mu_s(\ln\alpha_a - \ln\mu_a) - \mu_a\mu_s\ln k\Big],\end{aligned} \tag{20}$$

where $\psi \equiv \alpha_s\mu_a + \alpha_a\mu_s + \mu_s\mu_a + \alpha_g\mu_s\mu_a$ and we again model fuel economy (a good) rather than fuel consumption (a bad). Note that heterogeneity in the preference parameters $(\beta_s, \beta_a, \beta_g pm)$ 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 her preferences parameters 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 size ($\ln\beta_s$) and acceleration ($\ln\beta_a$) 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 ($\ln\beta_g pm$). Inspection of equations (18)–(20) reveals that consumers with stronger preferences to save fuel (larger $\ln\beta_g pm$) will choose smaller, slower, and more efficient cars. Further, because equations (18)–(20) are linear, any pair of logged attributes will be perfectly

Figure 2: *Optimal attribute choices for two consumers*

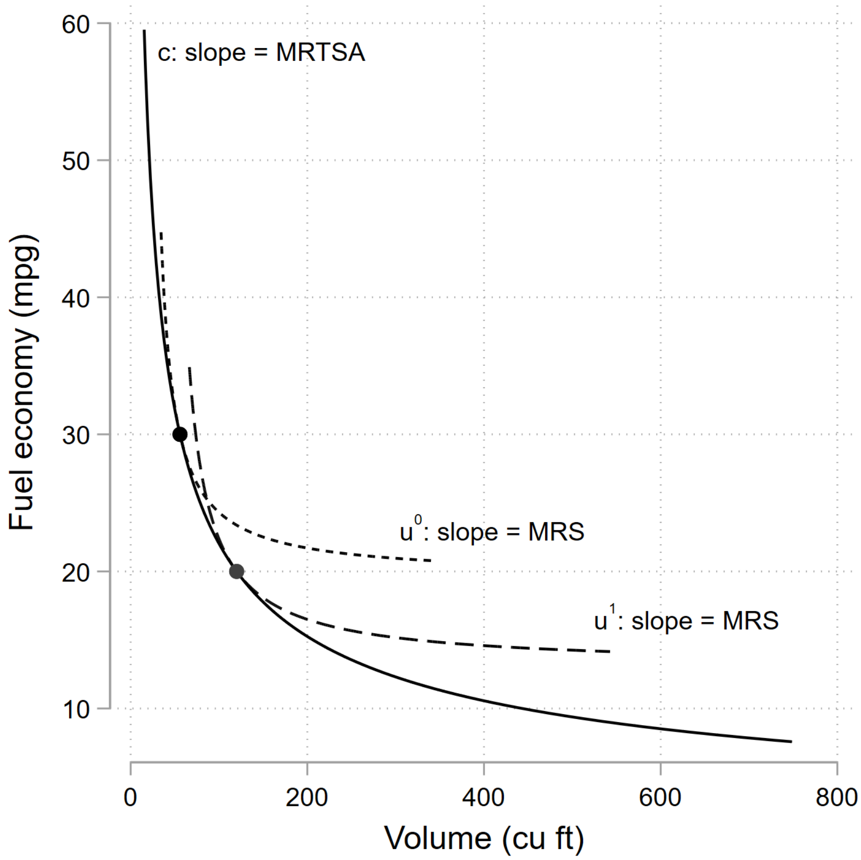

Note: This figure illustrates optimal attribute choices in two-dimensional (fuel economy and size) space as points of tangency between consumers' indifference curves (us) with slopes equal to the MRS and the isocost curve in production (c) with slope equal to the MRTSA.

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.

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 the 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 (13) using the specific functional form in (3), solve for gallons-per-mile (g), and then take logs to yield:

$$\ln g = \frac{1}{1+\alpha_g}\ln k - \frac{\theta}{1+\alpha_g} - \frac{1}{1+\alpha_g}\ln(\beta_g pm + \tau) + \frac{\alpha_s}{1+\alpha_g}\ln s + \frac{\alpha_a}{1+\alpha_g}\ln a \tag{21}$$

$$\approx \frac{1}{1+\alpha_g}\ln k - \frac{\theta}{1+\alpha_g} - \frac{1}{1+\alpha_g}\left[\ln p + \ln m - \ln(r+\rho+\delta) + \frac{\tau}{\beta_g pm} + \ln\beta_g\right] + \frac{\alpha_s}{1+\alpha_g}\ln s + \frac{\alpha_a}{1+\alpha_g}\ln a, \tag{22}$$

where the approximation in the second row follows from a first-order Taylor expansion about $tau = 0$. Note that the effect of τ is likely negligible. Back-of-the-envelope calculations suggest that pm is at least an order of magnitude larger than τ.[11] We return to this issue in our empirical analysis below, where we show that our results are robust to our controlling for variation in β_g due to τ. For estimation, we model lifetime miles (m) as the present-discounted value of lifetime miles, assuming a constant rate of time discounting (r), a constant rate of car scrappage (ρ), and a constant rate of decay in annual miles driven conditional on car survival (δ). The value of present-discounted lifetime miles is therefore given by $m = m(0)/(r+\rho+\delta)$, where $m(0)$ is miles driven in the car's first year. See appendix A.3 for details.

This equation suggests that we can recover the 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 size, log acceleration, 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, size, and acceleration 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 from the coefficients in this regression. Below, we show how to recover the cost parameters from household-level microdata by regressing log fuel economy choices on log gas prices, log miles, and log car attributes while controlling for household-level demographics. Using household-level microdata (as opposed to model-level attribute data) allows us to control for local gas prices, miles traveled, income, and other observables that might correlate with fuel economy preferences.

[11] Note that changes in $\ln\beta_g$ are given by $\Delta\ln(pm+\tau) \approx \frac{\Delta pm}{pm+\tau} + \frac{\Delta\tau}{\tau}\frac{\tau}{pm+\tau} = \left[\frac{\Delta pm}{pm} + \frac{\Delta\tau}{\tau}\frac{\tau}{pm}\right]\frac{pm}{pm+\tau}$ for small changes Δpm and $\Delta\tau$. Thus, the relative importance of percent changes in pm and τ in brackets is governed by the τ/pm term. What is the value of this term? Consider a fuel economy improvement from 20 mpg to 25 mpg, implying a $1/20 - 1/25 = 0.01$ reduction in gallons-per-mile. Assuming a fuel price of \$3 per gallon and 100,000 miles of lifetime travel (present value), this translates to $0.01 \cdot 3 \cdot 100,000 = \$3000$ of fuel savings. Meanwhile, the noncompliance penalty under U.S. fuel economy standards is \$55 per mile-per-gallon per car before 2007, which translates to $5 \cdot 55 = \$225$ for a 5 mile-per-gallon improvement. Note that the noncompliance penalty is an upper bound on the marginal compliance costs. Indeed, Anderson and Sallee (2011a) show that marginal costs are under \$30 in many years, although the shadow cost is estimated to be as high as \$65 after 2006, which translates to $5 \cdot 65 = \$325$ for a 5 mile-per-gallon improvement. These calculations imply that pm is an order of magnitude larger than τ. Thus, a percent change in pm has 10 times the effect on $\ln\beta_g$ as a percent change in τ.

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

How does our analysis differ from that in Knittel (2011)? Both analyses seek to estimate the slope of the isocost curve at a point in time and the shift in this curve over time due to technical change. However, 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:

$$\ln g = \frac{1}{\alpha_g} \ln k - \frac{\theta}{\alpha_g} + \frac{\alpha_s}{\alpha_g} \ln s + \frac{\alpha_a}{\alpha_g} \ln a - \frac{1}{\alpha_g} \ln c, \tag{23}$$

which is the form of Knittel's (2011) baseline regressions: log fuel economy on log attributes and year dummies to capture technical progress (θ).[12] Thus, our functional form for costs provides explicit microfoundations for Knittel (2011) and related papers.

Note that the coefficient on log size (α_s/α_g) captures the slope of the isocost curve in two-dimensional product space: percent changes in g that hold costs fixed given a percent change in s. The same applies for the coefficient on log acceleration (α_a/α_g). Note, however, that technology costs ($\ln c$) are unobserved in Knittel (2011). Thus, causal identification via regression requires that technology costs be uncorrelated with size and acceleration. Our analysis shows that this assumption is unlikely to hold. Conditional on other car attributes, consumers with a strong desire to save fuel purchase more efficient cars (equation (22)), putting them on a higher isocost curve. Meanwhile, these same consumers tend to choose different attributes (equations (18) and (19)), which implies different marginal costs for fuel economy (equation (3)). Thus, costly fuel-saving technologies will not be added to cars at random. We show this point theoretically below in our analysis of discrete technology adoption, and we confirm it empirically through an examination of penetration rates for various engine technologies.

Knittel (2011) concedes (pg. 3372) that this regression may not yield unbiased estimates for the slopes of isocost curves at a point in time or literal shifts in isocost curves over time due to technical change. He argues, however, that his estimates remain valid for predicting how fuel economy would have evolved in the presence of technical change had car attributes remained fixed. In effect, he argues that he is estimating the reduced-form time trend ($\theta/(1 + \alpha_g)$) in our equation (22). Our model clarifies that, for this argument to hold, it is crucial to control for shifts in fuel economy driven by consumer preferences ($\ln \beta_g$). Knittel (2011) partially addresses this concern, estimating an auxiliary 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) that isocost elasticities are not the same as counterfactual equilibria but contends that the estimates represent what is technologically feasible given the estimated rates of technical progress. We show in the next section that the equilibrium response to fuel economy forcing

[12] Our 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 the signs on the coefficients appropriately flipped.

mechanisms (gas prices and standards) is to trade off much less acceleration and size in equations (30) and (29) by adding fuel-saving technology. The isocost elasticity is the same as the equilibrium response only when the elasticity is zero, an elasticity that implies no trade-off is possible. The bias arising from interpreting isocost curves as proxies for equilibrium response to standards is increasing in the estimated isocost elasticity.

Finally, the time trend in the model measures the equilibrium trend in fuel economy relative to other attributes. The trend could easily be negative and still consistent with technical change. In our model, this would occur if $\alpha_g < 1$. For example, if average fuel economy stayed at a constant level throughout the time series but other attributes improved, the time trend would be negative, but there easily could have been technical change. We would not be able to say without accounting for consumer preferences.

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 equations (11) and (12) using the functional form in (3), solve for size (s) and acceleration (a), and then take logs to yield:

$$\ln s = \frac{1}{\alpha_s + \mu_s} \underbrace{\left(\theta - \alpha_a \ln a + \alpha_g \ln g - \tau \frac{\sigma_s}{c_s}\right)}_{\text{cost shifter for } s} + \frac{1}{\alpha_s + \mu_s} \left(\ln \mu_s \beta_s - \ln \alpha_s k\right) \tag{24}$$

$$\ln a = \frac{1}{\alpha_a + \mu_a} \underbrace{\left(\theta - \alpha_s \ln s + \alpha_g \ln g\right)}_{\text{cost shifter for } a} + \frac{1}{\alpha_a + \mu_a} \left(\ln \mu_a \beta_a - \ln \alpha_a k\right). \tag{25}$$

Here, note that we already know the relevant cost parameters (θ and αs) based on (22) above.

This equation suggests that 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. Intuitively, variation in acceleration and fuel economy shifts the marginal cost of size. In equilibrium, the change in size induced by this shift depends on the elasticity of demand for size (μ_s), along with the corresponding cost parameter (α_s). Thus, having estimated cost parameters in a first step, we can estimate the demand for size. Likewise, we can recover the demand parameter for acceleration (μ_s) by regressing log acceleration on a linear combination of log size and log fuel economy. Note that, before the 2011 footprint-based standards, $_s = 0$. Below, we show how to recover demand parameters in a second step after estimating costs using household-level microdata.

3.4 Gasoline prices and fuel economy standards

We now explore equilibrium responses to higher fuel prices. Recall that $\beta_g \equiv pm + \tau$ and assume that $\tau \approx 0$, such that percent changes in fuel prices correspond directly to percent changes in β_g. Differentiating equations (18)–(20) with respect to log preferences for fuel economy ($\ln \beta_g$) yields:

$$\frac{\partial \ln s}{\partial \ln p} = -\frac{\alpha_g \mu_a}{\psi} < 0 \tag{26}$$

$$\frac{\partial \ln a}{\partial \ln p} = -\frac{\alpha_g \mu_s}{\psi} < 0 \tag{27}$$

$$\frac{\partial \ln g^{-1}}{\partial \ln p} = \frac{\alpha_s \mu_a + \alpha_a \mu_s + \mu_s \mu_a}{\psi} > 0, \tag{28}$$

where recall that $\psi \equiv \alpha_s \mu_a + \alpha_a \mu_s + \mu_s \mu_a + \alpha_g \mu_s \mu_a$. Car size and acceleration decrease, while fuel economy increases. These changes are proportional to baseline attributes: consumers who strongly value fuel economy will increase it more, while consumers who value size and acceleration will decrease them more. Thus, the distribution of fuel economy expands, while the distributions of size and acceleration both contract.

Consider the percent change in size in (26). Note that larger αs and μs imply more inelastic supply and demand for car attributes. So size is more responsive to changes in gas prices 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 (27) follows a similar pattern. Now consider the percent change in fuel economy in (28). The numerator implies that fuel economy is more responsive to gas prices when the supply and demand for size and acceleration are more inelastic, again for a given ψ.

We now use equations (26)–(28) to calculate the percent increase in fuel economy induced by higher gas prices, relative to the percent decrease in size and acceleration:

$$-\frac{\partial \ln g^{-1}}{\partial \ln s}\bigg|_{\Delta \ln \beta_g} = \frac{\alpha_s}{\alpha_g} + \frac{\mu_s(\alpha_a + \mu_a)}{\alpha_g \mu_a} \tag{29}$$

$$-\frac{\partial \ln g^{-1}}{\partial \ln a}\bigg|_{\Delta \ln \beta_g} = \frac{\alpha_a}{\alpha_g} + \frac{\mu_a(\alpha_s + \mu_s)}{\alpha_g \mu_s} \tag{30}$$

Note that the equilibrium trade-off between fuel economy and size in (29) is larger than the iso-cost elasticity (α_s/α_g). Intuitively, consumers respond to higher gas prices not only by reducing car size holding cost fixed, but also by adding fuel-saving technology, or by reducing acceleration. Thus, the deviation from the iso-cost elasticity is large when the supply of fuel economy and demand for acceleration are highly elastic (small α_g and μ_a), especially when demand for size is inelastic (large μ_s). An analogous pattern holds for the equlibrium trade-off between fuel economy and acceleration in (30).

Note that increasing the stringency of the fuel economy standard (τ) has qualitatively similar effects. To see this, observe that $\partial \ln(\beta_g pm + \tau)/\partial \tau \approx 1/\beta_g pm$ when evaluated at $\tau \approx 0$. Thus, just re-scale each of the effects in equations (26)–(28) by $1/pm$ to yield the marginal percentage effects, where recall that m varies across consumers. This re-scaling relates to the finding in Jacobsen, Knittel, Sallee, and Van Benthem (2020) that fuel economy standards imperfectly target car emissions since they fail to account for heterogeneity in lifetime miles across car and consumer types.

Now consider the effect of higher gasoline prices on fuel economy, conditional on other attributes. Differentiating equation (22) with respect to log preferences for fuel economy ($\ln \beta_g$) yields:

$$\frac{\partial \ln g^{-1}}{\partial \ln p} = -\frac{1}{1+\alpha_g} > 0, \tag{31}$$

which shows that fuel economy increases conditional on size and acceleration. Thus, consumers optimally choose cars that are smaller and slower overall (equations 26 and 27). In addition, for cars of a given size and acceleration, consumers optimally choose more efficient and therefore more costly cars. Note that the partial derivative in equation (31) is only a function of the cost parameter (α_g). This result suggests that exogenous variation in fuel costs, controlling for size and acceleration, can be used to identify the cost coefficient on gallons-per-mile, which is essential to identifying the other cost parameters (α_s, α_a), as well as 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 parameter θ, which proportionally reduces costs, and which preserves the ratios of marginal costs in equations (5)–(7). Visually, increasing θ shifts the iso-cost curves outward in goods space such that, for any level of cost, all attributes may be improved.

First consider a marginal increase in parameter θ in equations (18)–(20). We can easily see that this increase leads to an improvement in all three attributes:

$$\frac{\partial \ln s}{\partial \theta} = \frac{\mu_a}{\psi} > 0 \tag{32}$$

$$\frac{\partial \ln a}{\partial \theta} = \frac{\mu_s}{\psi} > 0 \tag{33}$$

$$\frac{\partial \ln g^{-1}}{\partial \theta} = \frac{\mu_s \mu_a}{\psi} > 0, \tag{34}$$

where again recall $\psi \equiv \alpha_s\mu_a + \alpha_a\mu_s + \mu_s\mu_a + \alpha_g\mu_s\mu_a$. Thus, if $\mu_s = \mu_s = 1$, attribute-neutral technical change leads to an equal proportional improvement in size, acceleration, and fuel economy. In general, however, relative changes depend on the μ parameters:

$$\left.\frac{\partial \ln g^{-1}}{\partial \ln s}\right|_{\Delta\theta} = \mu_s, \tag{35}$$

$$\left.\frac{\partial \ln g^{-1}}{\partial \ln a}\right|_{\Delta\theta} = \mu_a, \tag{36}$$

$$\left.\frac{\partial \ln a}{\partial \ln s}\right|_{\Delta\theta} = \frac{\mu_s}{\mu_a}. \tag{37}$$

Thus, if $\mu_s > 1$ (demand for size sufficiently inelastic), then fuel economy improves more than size. Likewise for acceleration when $\mu_a > 1$. Meanwhile, if $\mu_s > \mu_a$ (demand for acceleration more elastic than demand for size), then acceleration improves more than size. These gains are all inversely proportional to the sum of cost and preference parameters in the denominator (ψ). The cost parameters (the αs) reflect the elasticity of marginal costs with respect to car attributes. Thus, a steeply increasing marginal cost for any car attribute is a drag on the attribute-enhancing role of technical change for all car attributes. Attribute changes scale proportionally with baseline attributes: consumers use technical change to buy more of the attributes they already value most. We can see this spreading-out effect clearly in figure (11) in the next section.

Now consider how a marginal increase in parameter θ affects fuel economy conditional on other attributes in equation (22). Differentiating with respect to θ yields:

$$\frac{\partial \ln g^{-1}}{\partial \theta} = \frac{1}{1+\alpha_g} > 0. \tag{38}$$

Thus, technical change manifests as a shift in equilibrium fuel economy conditional on other attributes. The size of this shift in percentage terms is the rate of technical change itself, scaled by the elasticity of marginal cost with respect to fuel economy $(1+\alpha_g)$. The larger this elasticity, the smaller the shift. Knittel (2011) interprets this conditional shift in fuel economy directly as the rate of technical change. But this is incorrect. In equilibrium, consumers allocate a portion of technical change to cost savings. Thus, the equilibrium shift understates the holding-cost-fixed rate of technical change. The correct rate is given by differentiating equation (23):

$$\frac{\partial \ln g(s,a,c;\theta)^{-1}}{\partial \theta} = \frac{1}{\alpha} = \frac{\partial \ln g(s,a;\theta)^{-1}}{\partial \theta}\frac{1+\alpha_g}{\alpha_g}, \tag{39}$$

which is the equilibrium shift scaled by $(1+\alpha_g)/\alpha_g$. Thus, estimating the cost parameter (α_g) is essential both to estimating attribute trade-offs and inferring the rate of technical change from market data. Alternatively, we might be interested in the change in all attributes that could be obtained via technical change, holding cost fixed. This rate is given by totally differentiating logged cost in equation (3), setting the total change to zero, and solving for the equivalent change in attributes:

$$\left.\frac{\partial \ln g(s,a,c;\theta)^{-1}}{\partial \theta}\right|_{\Delta \ln g^{-1}=\Delta \ln s=\Delta \ln a} = \frac{\partial \ln g(s,a;\theta)^{-1}}{\partial \theta}\frac{1+\alpha_g}{\alpha_a+\alpha_s+\alpha_g}, \tag{40}$$

which is the equilibrium shift scaled by $(1+\alpha_g)/(\alpha_a+\alpha_s+\alpha_g)$.

3.5.2 Attribute-biased technical change

Now consider attribute-biased technical change. To generate clear theoretical results, we model this form of technical change via a linear add-on to our original cost function:

$$\tilde{c}(g,a,s) = c(g,a,s) + \eta(g-g_0), \tag{41}$$

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) and the marginal cost of improving fuel economy is uniformly lower. Marginal cost is now:

$$\tilde{c}_g = -\frac{\alpha_g c}{g} + \eta = -\frac{\alpha_g c - \eta g}{g}, \tag{42}$$

which implies the following MRTSA between fuel economy and acceleration:

$$-\frac{\tilde{c}_a}{\tilde{c}_g} = -\frac{\alpha_a}{\alpha_g - \frac{\eta g}{c}} \frac{g}{a} < -\frac{\alpha_a}{\alpha_g} \frac{g}{a}, \tag{43}$$

where the last inequality assumes $\eta > 0$ with $\alpha_g > \eta g/c$. Thus, the iso-cost curves become steeper in fuel economy vs. acceleration space, implying a higher opportunity cost of improved acceleration.

How does this technical change affect attribute choices? Note that an increase in the technical change parameter (η) is mathematically equivalent to an increase in the fuel economy standard's credit price (τ), which we considered in the previous sub-section. Thus, biased technical change also leads to smaller, slower, and more fuel-efficient cars. Intuitively, biased technical change raises the opportunity cost of improving size and acceleration at the expense of fuel economy. So fuel economy goes up and these other attributes go down.

Figure 3 illustrates these points in fuel economy and acceleration space. As 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 steeper. The $MRTSA_{ga}$ has increased, implying that the opportunity cost of acceleration is now higher. Thus, the consumer now strictly prefers attribute bundles to the immediate northwest of the initial bundle: higher fuel economy and lower acceleration.

However, the combined effect of biased technical change and a change in gas prices is quite different. As noted above, biased technical change favoring fuel economy makes the slope of the iso-cost curve steeper. Equations (29) and (30) show that a steeper iso-cost elasticity increases the equilibrium response of fuel economy to gas prices, relative to size and acceleration. Thus, size and acceleration will be relatively less affected by gas taxes, fuel-economy standards, and energy price shocks in the presence of biased technical change.

Later we show evidence that attribute-biased technical change is important for understanding changes in attribute trade-offs for light-duty vehicles since the 1990s.

3.5.3 Discrete technology adoption

Carmakers periodically invent new technologies that can be added to cars to improve one or more attributes. For example, turbochargers boost the performance of an internal combustion by forcing additional compressed air into the combustion chamber. Adding a turbocharger to a car's drivetrain allows it

Figure 3: *Biased technical change*

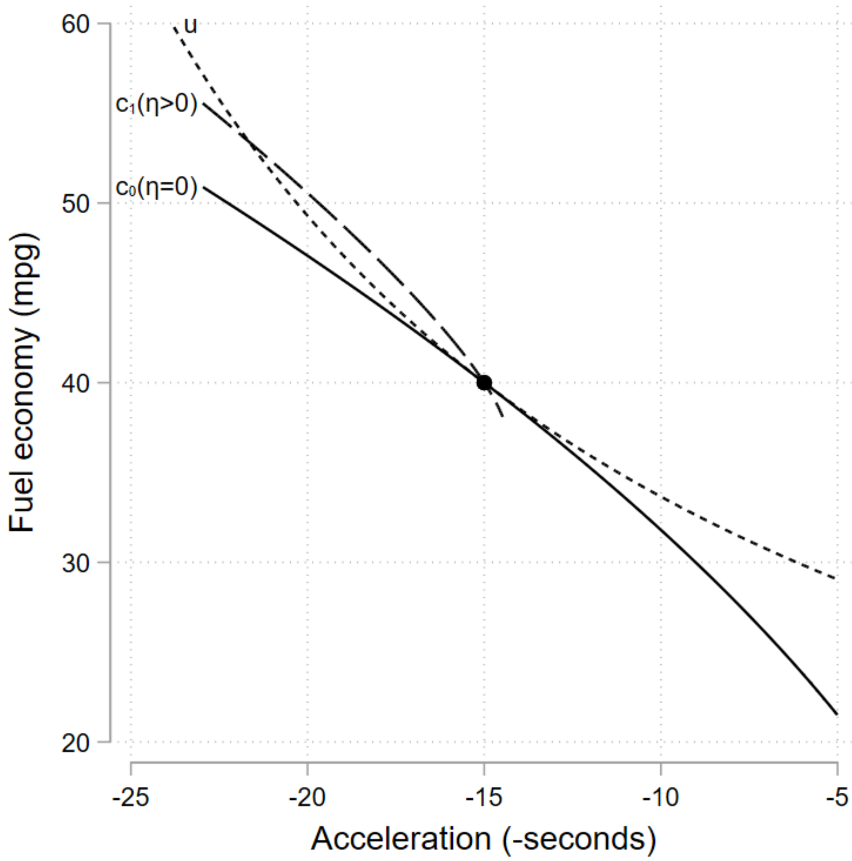

Note: This figure illustrates the effects of biased technical change in two-dimensional (fuel economy and acceleration) space. The figure shows an optimal attribute bundle as the point of tangency between a consumer's indifference curve (u) and an iso-cost curve prior to technical change ($c(\eta = 0)$).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

to accelerate faster. Carmakers often add turbochargers and simultaneously decrease engine displacement ("downsizing"), thereby improving fuel economy while holding acceleration constant (Shahed and Bauer 2009). Over time, the cost of adding a turbocharger has fallen—a form of technical change—leading to increased adoption. Other examples include continuously variable transmission, direct injection (replacing the carburetor), and electric or hybrid-electric drivetrains, all of which have experienced falling costs and increased adoption. How does this form of technical change manifest in the car market? Are new technologies adopted uniformly across the fleet? Or is their adoption biased toward specific consumer segments or cars with particular attributes? If so, what are the implications?

We consider a discrete technology targeting acceleration—to use the turbocharger example—that consumers may add to their car at a cost. This technology has two important features. First, we assume that the installation cost (κ) in a given year is the same for all cars regardless of baseline attributes. Second, we assume that the technology yields a fractional improvement (ω) in acceleration. Thus, cars with a higher baseline level of acceleration will see a larger absolute improvement.

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:

$$u(g, a, s; \omega) = y + \tau\sigma - \beta_s s^{-1} - \beta_a((1+\omega)a)^{-\mu_a} - \beta_g pmg - c(g, a, s), \tag{44}$$

where we have simply scaled perceived acceleration by $1+\omega$. By the envelope theorem, the utility gain from a marginal increase in ω at the optimum is given by: $\partial u(g, a, s; \omega)/\partial\omega = (1+\omega)^{-(1+\mu_a)}\mu_a\beta_a a^{-\mu_a}$. Thus, the net benefit from adopting the discrete technology is given by the following approximation:

$$\begin{aligned} u(g, a, s; \omega > 0) - u(g, a, s; \omega = 0) - \kappa &\approx \omega\mu_a\beta_a a^{-\mu_a} - \kappa \\ &= \omega\alpha_a c^* - \kappa, \end{aligned} \tag{45}$$

where c^* is the equilibrium 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 the equality in the second line follows from substituting the marginal benefit of acceleration with its marginal cost based on necessary condition (12).

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 craves speed. Or perhaps she craves size and fuel economy. Either way, she will opt to install the discrete technology if her car is sufficiently costly. We confirm this prediction below in section 4, showing that larger, faster, and more efficient cars are more likely to feature advanced engine technologies, such as turbocharchers and hybrid-electric drivetrains.

What happens to equilibrium attributes upon adoption? Consider equations (18)–(20) and note that adopting the acceleration-boosting technology is equivalent to replacing acceleration preferences ($\ln\beta_a$) with modified preferences ($\ln\beta - \mu_a\ln(1+\omega)$) in each of these equations. Differentiating each of these equations with respect to ω and evaluating at $\omega = 0$ then yields:

$$\left.\frac{\partial \ln s}{\partial \omega}\right|_{\omega=0} = \frac{\alpha_a \mu_a}{\psi} > 0 \tag{46}$$

$$\left.\frac{\partial \ln(1+\omega)a}{\partial \omega}\right|_{\omega=0} = \frac{\alpha_a \mu_s}{\psi} > 0 \tag{47}$$

$$\left.\frac{\partial \ln g^{-1}}{\partial \omega}\right|_{\omega=0} = \frac{\alpha_a \mu_s \mu_a}{\psi} > 0, \tag{48}$$

where the relevant attribute in the second row is log acceleration as perceived by the consumer.[13] Strikingly, consumers that adopt the discrete technology experience improvements in all attributes, i.e. even those not targeted by the technology. Even more strikingly, the improvements are identical to those in equations (32)–(34) for attribute-neutral technical change, just multiplied by the cost parameter on acceleration (α_a). Thus, a falling cost for discrete, attribute-boosting technologies is one way of motivating attribute-neutral technical change in section 3.5.1 above.

These same results all 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 as perceived by the consumer remains fixed but decreases by fraction ω in the cost function, driving down costs. Note that scaling the targeted attribute by $(1-\omega) < 1$ is equivalent to scaling the entire cost function by $(1-\omega)^{\alpha_a} < 1$, given our functional-form assumption. Thus, adoption will be most attractive to consumers choosing expensive cars, with the same effect in equilibrium as attribute-neutral technical change.

Importantly, there is nothing special about acceleration in our analysis. The same qualitative results hold for a technology targeting fuel economy. Thus, consumers with strong preferences in any dimension are more likely to adopt a technology targeting any attribute. Only as installation costs (κ) fall over time will others adopt. This result mirrors the stylized fact that new car technologies typically first appear in luxury segments. This result also has profound implications for policies to promote hybrid and electric cars. By lowering adoption costs (κ), a technology subsidy first spurs adoption among larger and faster cars. The initial boost to fuel economy is then reallocated, at least in part, to making these cars even larger and faster. Thus, the gasoline savings from a hybrid or electric vehicle technology subsidy in equilibrium may be much lower than what policymakers expect, if they naively assume that other car attributes remain fixed.

4 Data and descriptive statistics

We begin by describing our data sources. We then present descriptive evidence on car attributes, gasoline prices, and discrete technology adoption.

[13] Note that the second row follows from the fact that $\partial \ln(a(1+\omega)/\partial\omega = \partial \ln a/\partial\omega + 1$ when $\omega = 1$, and the fact that $\partial \ln a/\partial\omega = -(\alpha_s + \mu_s + \alpha_g \mu_s)\mu_a/\psi$ from equation (19).

4.1 Data

To estimate our model of car attribute production, we combine household microdata on car choices with model-level data on car attributes, including size, acceleration, and fuel economy, along with state-level data on gasoline prices. We supplement these data with estimates for the stringency of fuel economy standards, interest rates on new car loans, and information on car durability.

Our car choice data come from the National Household Transportation Survey (NHTS) public use microdata waves of 2001, 2009, and 2017. These data record the make, model, and model year of cars owned for individual households at the time of the survey, along with the number of miles the car was driven in the last 12 months. We calculate the age of each car based on the year of the survey and model year of the car. These data also record key household characteristics, including income, household size, and population density of the surrounding census tract. In some analyses, we wish to control for household characteristics that influence demand for fuel savings. We are unable to differentiate between cars purchased new vs. used. However, we argue above that interactions between car age and household characteristics help us control empirically for the mismatch in characteristics between the original car buyers and current owners, i.e., since buyers of new cars will implicitly channel the preferences of future used car buyers. See Appendix A.3 for details.

We match the car choice data to information on car attributes from several sources. Our size data come from Transport Canada, a Canadian federal government institution, and are compiled by the Canadian Association of Road Safety Professionals. These are the data offered by the National Highway Traffic Safety Administration's Vehicle Identification Number Decoder tool, so we are confident that these foreign data are applicable to US vehicles. With these data, we measure the volume of a box bounding the car's useful space: length times width times height. These data constrain the start of our study period to 1995, after NAFTA. 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 is the maximum distance from left to right. Height is the maximum distance between the ground and 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 EPA prior to matching to the NHTS data.

Our fuel economy and acceleration data come from the U.S. EPA. We measure fuel economy as the combined city–highway measure reported by the EPA. The EPA measures gallons-per-mile in a simulator designed to mimic both city and highway driving. The EPA calculates a weighted 55% city and 45% highway average of miles per gallon.[14] We proxy for acceleration as the horsepower-to-weight ratio since this measure is readily available for all cars and 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 fuel economy and acceleration data 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 our car choice data to state-level retail gasoline prices. These data are based on those of Davis and Kilian (2011), who start with pretax retail prices reported by the Energy Information Administration (EIA) and then painstakingly add percentage-based (ad valorem) and constant per-gallon (specific) gasoline taxes from myriad sources. Bates and Kim (2022) update these data to include additional years using data from GasBuddy.com and in some cases correcting errors in Davis and Kilian (2011).

We collect information on the shadow cost of fuel economy standards from several sources that have attempted to estimate this value. Before 2008, the shadow cost of standards varied by manufacturer. We take estimates from Anderson and Sallee (2011a) of the manufacturer-level shadow cost of compliance from 1996 to 2006. From 2007 to 2011, standards appeared to have not been binding (Yeh, Burtraw, Sterner, and Greene 2021). Starting in 2012, the NHTSA and EPA allowed compliance credits to be traded between firms. We take estimates of the fleet-wide shadow cost from Yeh et al. (2021) for 2012 to 2017. All estimates of shadow cost are an order of magnitude lower than the pecuniary benefits from fuel savings. Thus, our main specification does not use these data, but we test that our estimates are robust to their inclusion.

We measure the real annual interest rate on 48-month new car loans from the Federal Reserve Bank of Saint Louis (Reserve 2019). We combine these data with estimates of car scrap probabilities as reported by NHTSA (2006), along with our own estimates for the annual decay in miles for new vs. used cars based on the NHTS data. We use this information to calculate and control for the trend in interest rates plus durability and miles decay that would otherwise contaminate our estimate of technical change. Later, we use this information to calculate the present discounted value of lifetime miles for each car in our dataset for use in our counterfactual simulations. We describe these procedures in the appendix.

To test our theoretical predictions about the adoption of discrete fuel-saving technologies (e.g., turbochargers), we combine model-level data on car technology with data on car attributes, including size, weight,

[14] The EPA 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.

[15] Acceleration is often measured via the time it takes a car to accelerate from 0 mph to 60 mph in seconds. Our concept of acceleration as the "rate of velocity increase" is inversely related to this 0–60 measure. MacKenzie and Heywood (2012) show that the relationship between horsepower-to-weight and 0–60 time changed little from 1995 to 2010 (the end of their study period).

[16] For 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.

horsepower, and fuel economy. Our data on discrete fuel-saving technologies come from Wards Automotive and record, for a given car production model, the presence of a turbocharger, gasoline direct injection, 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 described above. Note that our unit of observation here is a single production model offered by a carmaker, rather than an individual household's choice of such model. These data are matched to EPA data to capture sales weights.

4.2 Correlations in car attributes

Figure 11 shows the distributions of fuel economy, size, and acceleration 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. size across car models in both time periods. Each dot corresponds to 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 distribution of car attributes across households. The bottom panel presents similar information for fuel economy vs. acceleration.

Echoing Knittel (2011), this figure reveals two striking facts. First, in any given time period, there is a strong negative correlation between fuel economy and size (inspecting the scatter diagrams). The same is true for fuel economy and acceleration. This pattern is consistent with strong technical trade-offs between fuel economy, size, and acceleration, with cost held fixed. However, our theory shows that the slope of these relationships also depends on the cost of adding fuel-saving technology. Second, cars have become larger, faster, and less efficient over time (as we observe by comparing the red and black histograms). However, cars of a given size and speed have become more efficient over time (as we see by comparing the red and black scatter diagrams). These shifts are consistent with growing consumer preferences for size and acceleration, along with improved car technology that has allowed even the largest and fastest cars to become more efficient—interpretations emphasized by Knittel (2011). Again, however, our theory shows that these shifts also depend on rising gasoline prices, which make fuel-saving technologies more valuable, along with falling interest rates and improved durability, which make all car attributes more valuable in present-value terms.

Identification of our model parameters relies largely on cross-sectional correlations among car attributes at a point in time, along with shifts in these relationships over time.

4.3 Gasoline prices and interest rates

Figure 5 shows trends in monthly real gasoline prices for three large and differentiated states: California, Michigan, and Texas. The top panel illustrates that state-level gasoline prices largely follow a common trend

Figure 4: *New car attributes in 1999–2001 and 2015–2017*

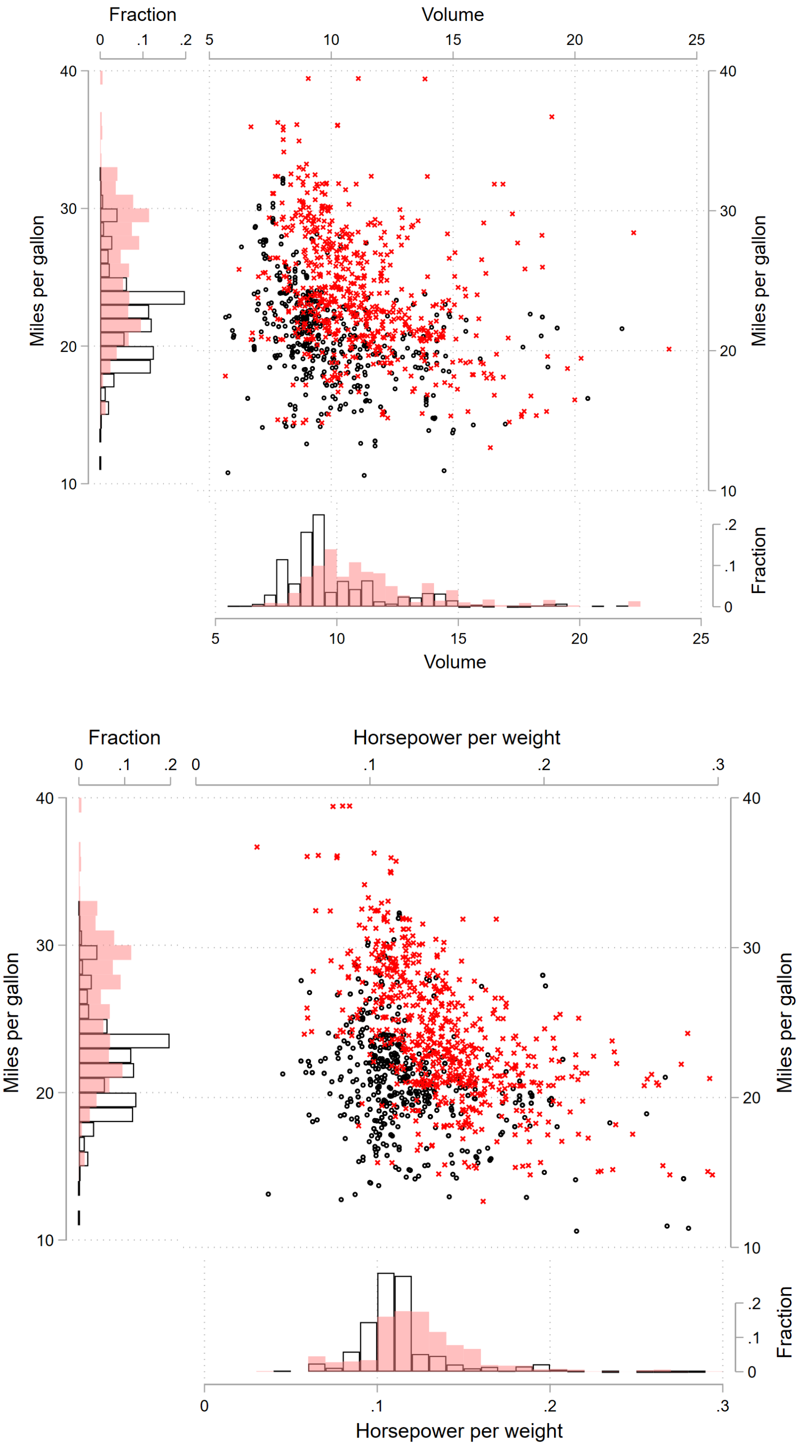

Note: This figure plots the distribution of new car fuel economy (miles per gallon), volume (cubic meters based on $l \times w \times h$ of car), and acceleration (as proxied by the horsepower per weight-in-pounds ratio). Attributes for model years 1999–2001 are shown in black, while attributes for 2015–2017 are shown in red. Scatter diagrams show the correlation between fuel economy and size or acceleration; 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 within the given range of model years.

Figure 5: *Real gasoline prices*

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

driven by global crude oil prices, with persistent gaps across states. However, the bottom panel shows that there is substantial variation in gasoline prices, even after we remove state and month fixed effects, with California's price rising relative to Michigan's and Texas's over time. This residual variation is driven by myriad supply-side factors, such as changes in state fuel taxes, differential environmental regulation of fuels and refineries, and shifts in refining and distribution costs, along with potential shifts in state-level demand running up against supply constraints. We revisit these issues below when discussing estimation.

Figure 10 shows how changes in interest rates might impact our estimates. The top panel shows the trend in nominal and real interest rates on 48-month new-car loans. Note that the interest rates fell by a factor of 2 over our sample period. Meanwhile, the bottom panel shows the trend in the (negative) logged interest rate plus scrappage plus miles decay ($\ln(r_t + \rho + \delta)$), which our theory shows to be a key driver of investment in fuel-saving technology via its impact on the present value of lifetime miles. The trend in this variable indicates that the present value of lifetime miles rose 24% over our sample period because of falling interest rates. Thus, it will be important to control for this trend to infer technical change in the up-front of drivetrain technology. Note that we do not attempt to estimate or control for time variation in durability as captured jointly by scrappage rates (ρ) or the decay in miles for cars that remain on the road (δ) since these could be construed broadly as reductions in drivetrain costs that should be included in our estimate of technical change. While the NHTSA estimates that scrap rates (ρ) have fallen over time, we find no evidence that the decay in miles (δ) fell across the three NHTS waves, nor do we find that first-year miles ($m(0)$) changed substantially. Thus, we control only for time-series variation driven by interest rates. We discuss the shadow cost of standards in 5.1.3 below.

5 Empirical estimation and results

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

5.1 Fuel economy conditional on size and acceleration

We start by estimating a version of equation (22) above:

$$\ln mpg_{ist} = f_\theta(t) + \gamma_p(\ln p_{ist} - \ln(r_t + \rho + \delta) + \tau/mp_{ist}) + \gamma_s \ln volume_{ist} + \gamma_a \ln hpwt_{ist} + \phi_s + \varepsilon_{ist}, \quad (49)$$

where $\ln mpg_{ist}$ is log fuel economy, $\ln volume_{ist}$ is log volume, and $\ln hpwt_{ist}$ is log horsepower-to-weight for car i in state s bought in year t; $\ln p_{ist} - \ln(r_t + \rho + \delta) + \tau/mp_{ist}$ is the log gas price in state s and year t minus the logged conversion rate from lifetime miles to first-year miles plus the shadow cost divided by discounted lifetime miles and gas price; $f_\theta(t)$ captures attribute-neutral technical change; ε_{ist} is an error term; and the γs are coefficients to be estimated. We include state dummies (ϕ_s) in all models. We control for trends in technology (f_θ) with a full set of year dummies. Alternatively, we drop the year dummies and capture the annual rate of technical change with a linear time trend ($\gamma_\theta t$). In some models, we control for continuous car age and log miles, along with their squares and interaction. In some models, we additionally control for categorical household demographics, including income, number of household members, number of adults, number of vehicles, home ownership, urban vs. rural residence, and population density.

We estimate this equation via ordinary least squares (OLS). Our identification assumption is that the error term in our model (ε_{ist}) is uncorrelated with state-level gasoline prices, along with car size and acceleration, conditional on our inclusion of state and year dummies. There are two main concerns. First, unobserved shifts in a state's demand for fuel economy over time might correlate positively with gas prices in that state, for example, via a surge in demand for miles, leading to biased estimation of γ_p. We address this concern by controlling for household-level demographics and miles traveled interacted with car age; but note that we observe miles only for survey years 2001, 2009, and 2017.[17] Second, our error term might reflect omitted car attributes that affect fuel economy and that are correlated with size and acceleration. For example, pickup trucks, which are designed for hauling, will typically have lower gear ratios, diminishing both fuel economy and acceleration in regular driving. Thus, towing capacity will be positively correlated with size and negatively correlated with fuel economy and acceleration, leading to biased estimates. We address this concern by including a pickup dummy. Estimating isocost trade-offs among car attributes from market data inevitably requires cross-sectional identification.

Note that the coefficient on log gasoline prices ($\gamma_p \equiv 1/(1+\alpha_g)$) identifies the cost parameter on gallons-per-mile (α_g). Thus, this coefficient can be used in combination with the coefficients on log size ($\gamma_s \equiv -\alpha_s/(1+\alpha_g)$) and log acceleration ($\gamma_a \equiv -\alpha_a/(1+\alpha_g)$) to identify the underlying cost parameters on size ($\alpha_s \equiv -\gamma_s/\gamma_p$) and acceleration ($\alpha_a \equiv -\gamma_a/\gamma_p$), along with the ratios capturing attribute trade-offs along isocost 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 ($\Delta\theta \equiv \Delta f_\theta(t)/\gamma_p$).[18] We do not observe the precise month in which a car is purchased. Thus, we include two years of lagged gasoline prices (coupled with the same log conversion factor from lifetime miles to first-year miles) since the prior year's gasoline prices will arguably be more relevant for car purchases toward the beginning of the year and to permit a more flexible model of beliefs about future gas prices. These lags also allow for gradual adjustment in car production and distribution following shocks to gas prices. Alternatively, we model the log of the three-year moving average of gas prices, finding nearly identical results.

Our theory implies that the coefficient on log gas prices $\times$ lifetime miles will identify the cost parameter on gallons-per-mile (α_g). Unfortunately, we do not observe the present discounted value of a car's lifetime miles as anticipated by the original buyers at the time of purchase. This is not a problem, however, given our log specification, which implies that the coefficient on log gasoline prices is sufficient to identify the relevant cost parameter. We do subtract the log discount rate from log gas prices prior to estimation. This has no

[17] In future work, we will instrument for gas prices using state gas taxes, following Davis and Kilian (2011).

[18] To interpret the estimated time trend directly as technical change, we must properly control for and net out changes in the interest rate variable. This variable contains only time-series variation in our dataset, which poses a challenge for identification, especially when we include year dummies. However, our theory implies that the coefficient on this variable (γ_r) equals the coefficient on log gas prices (γ_p). We impose this theoretical restriction by adding the interest rate and depreciation variable to log gas prices prior to estimation. This approach has zero effect on our other estimates given our inclusion of year dummies.

Table 2: *Fuel economy conditional on attributes*

	(1)	(2)	(3)	(4)	(5)	(6)
Year		0.017***				
		(0.001)				
Log 3-yr avg. gas price			0.111*	0.095*	0.093*	0.098*
			(0.042)	(0.042)	(0.043)	(0.045)
Log HP/weight (γ_a)	-0.509***	-0.517***	-0.509***	-0.509***	-0.508***	-0.507***
	(0.017)	(0.016)	(0.017)	(0.017)	(0.017)	(0.017)
Log volume (γ_s)	-0.469***	-0.479***	-0.469***	-0.469***	-0.467***	-0.467***
	(0.007)	(0.007)	(0.007)	(0.007)	(0.007)	(0.007)
Pickup	-0.068***	-0.064***	-0.068***	-0.067***	-0.065***	-0.065***
	(0.008)	(0.009)	(0.008)	(0.008)	(0.008)	(0.008)
Year FEs	X		X	X	X	X
State FEs	X	X	X	X	X	X
VMT x Age				X	X	X
Demographics					X	X
Shadow cost						X
Sum gas price (γ_g)			.	.	.	
			.	.	.	
$\frac{\partial \ln \text{mpg}}{\partial \ln \text{HP/Wt}}\|_{\Delta c=0} \left(\frac{\alpha_a}{\alpha_g} = \frac{\gamma_a}{1-\gamma_g}\right)$			-0.573	-0.563	-0.560	-0.562
			0.044	0.043	0.043	0.045
$\frac{\partial \ln \text{mpg}}{\partial \ln \text{volume}}\|_{\Delta c=0} \left(\frac{\alpha_s}{\alpha_g} = \frac{\gamma_s}{1-\gamma_g}\right)$			-0.528	-0.519	-0.515	-0.517
			0.031	0.030	0.030	0.032

Note: This table presents coefficient estimates from equation (49). Standard errors in parentheses are clustered by state. *, **, and *** indicate statistical significance at the 5%, 1%, and 0.1% levels, respectively.
Data source: NHTS, EPA, and Transport Canada.

effect on the coefficient estimate for log gas prices, given our inclusion of year fixed effects, but ensures that the year fixed effects, which we use to infer neutral technical change, are not contaminated by the trend in interest rates. In addition, we observe an odometer-based calculation or imputation of annual miles over the last 12 months for every car. Thus, we are able to control for log miles interacted with car age, along with their squares and interaction, plus our other demographic controls, to capture cross-sectional and temporal variation in demand for fuel economy.

5.1.1 Main results

Table 2 presents our results. Column (1) follows Knittel's (2011) approach. We regress log fuel economy on log attributes, along with a pickup truck dummy, and capture technical change with year dummies. We additionally control for state dummies, as in all of our regressions. Consistent with what we observe in figure 11, we find that cars with 10% higher acceleration and size in a given year have 5.1% and 4.7% lower fuel economy. Pickup trucks have 6.8% lower fuel economy conditional on these attributes. Figure 6 plots the coefficients on the year dummies and shows large increases in fuel economy conditional on attributes (gray xs with dotted line). Following Knittel, column (2) then replaces the year dummies with a linear time trend. Consistent with the dummy year variables, we estimate a 1.5% annual increase in fuel economy conditional on size and acceleration.

Column (3) presents the results of our basic model, derived from theory, which adds current and lagged gas prices to regression (1). The three gas price coefficients added together imply that a 1% increase in fuel prices leads to a statistically significant $\gamma_g = 0.096\%$ long-run increase in fuel economy conditional on other car attributes (see bottom of table). This coefficient in turn identifies the elasticity of drivetrain costs with respect to fuel consumption: $\alpha_g = (1-\gamma_g)/\gamma_g = 9.4$. The coefficients on acceleration and size do not change with the inclusion of gasoline prices, nor does the coefficient on the pickup truck dummy. Consistent with our theory, the cross-sectional correlations between log fuel economy, log size, and log acceleration remain stable in the presence of higher gasoline prices. However, the nonzero coefficient on gas prices implies that these correlations are not literally the slopes of the isocost curves. To recover these slopes, we must divide by $1-\gamma_g \approx 0.9$. Thus, the isocost curves are approximately 1.1 times steeper than the equilibrium relationships shown in figure 11, as fuel-saving technology is disproportionately added to larger and faster cars.

Column (4) adds our controls for car age and log miles, along with their squares and interactions. Column (5) further adds our full suite of demographic controls. These variables are intended to capture cross-sectional heterogeneity in lifetime miles and other factors that might influence demand for fuel economy conditional on attributes. The coefficients on gas prices and car attributes are virtually unchanged. Column (6) uses the log of the three-year moving average gas price. The coefficient is nearly identical to the sum of the three coefficients from column (5).

Figure 6 again plots the coefficients on the year dummies. Model (3) controls for gas prices. The year coefficients from this model (triangles with short-dashed 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.[19] Model (4) additionally controls for car age and log miles, along with their squares and interactions. The year coefficients from this model (diamonds with dashed line) are slightly higher than those from model (3). Model (5) additionally controls for demographics. The year coefficients from this model (squares with long-dashed line) are notably higher, crossing over the coefficients from model (1) but with a shallower slope. Finally, model (6) uses the three-year moving average gas price. The year coefficients from this model (black circles with solid line) are essentially identical to those from model (5).

[19] To 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 6.

Figure 6: *Estimated coefficients on year dummies*

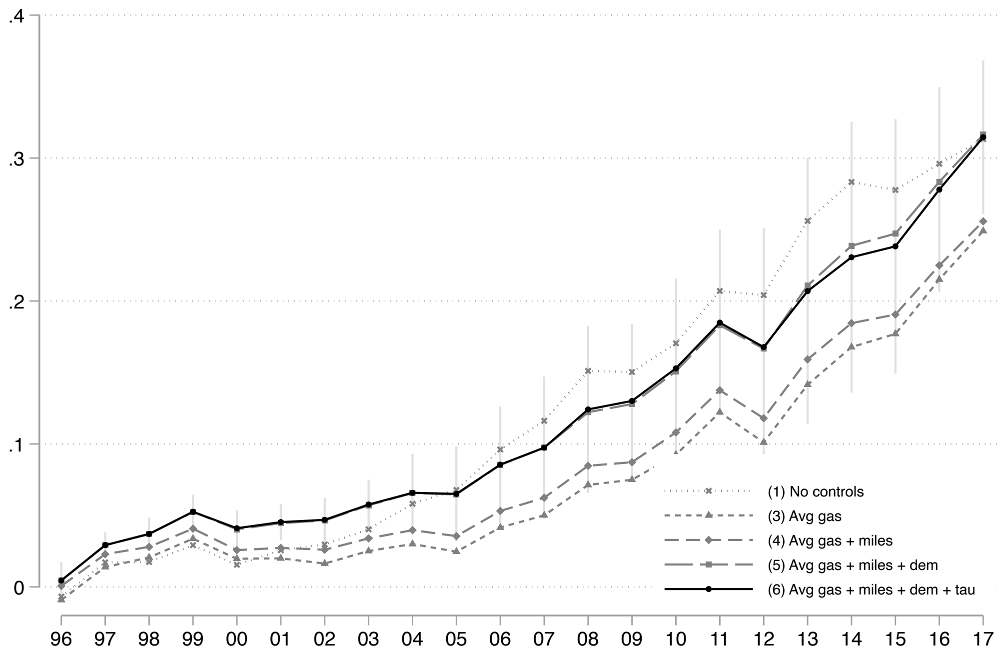

Note: This figure plots estimated coefficients on the year dummies from regressions (1), (3), (4), (5), and (6) from table 2. Vertical bars are the 95% confidence intervals from regression (5) based on standard errors clustered by state.

To recover our technical change parameter ($\Delta\theta$), we scale the coefficient on the year dummy for 2017 by $1/\hat{\gamma}_g$. Note that all year dummies are relative to 1995. The annualized rate of change in θ over time interval Δt is given by $r_\theta = \Delta\theta/\Delta t$. Based on our preferred specification, model (5), we estimate $r_\theta = 0.15$ over our 22-year sample period. This rate applies to theoretical drivetrain costs ($c(g, a, s)$), for which we lack a clear empirical touchstone. Thus, we divide by the sum of our cost parameters $\alpha_s + \alpha_a + \alpha_g = 21$ to yield the annual percent improvement in all car attributes that would hold drivetrain costs constant in the presence of attribute-neutral technical change. We recover the α parameters from our regression coefficients according to equation (71). Based on this approach, we estimate annual technical change in attributes of 0.7%.[20] Alternatively, we divide $r_\theta = 0.15$ by the cost parameter on fuel economy (α_g) to yield the annual percent improvement in fuel economy that would hold both drivetrain costs and other car attributes constant. Based on this approach, we estimate annual technical change in fuel economy of 1.5%.

[20] Equation (34) shows that consumers optimally choose attributes that increase by the same percentage in the presence of technical change. However, this equation shows that we divide $\Delta\theta$ by parameter $\phi = 1 + \alpha_s + \alpha_a + \alpha_g > \alpha_s + \alpha_a + \alpha_g$ to obtain these attribute changes. Thus, attributes increase by less than is implied by constant costs. This difference is small when $\alpha_s + \alpha_a + \alpha_g$ is large.

Table 3: *Biased technical change*

	(1) 1995-01	(2) 2002-09	(3) 2010-17
Log 3-yr ave. gas price	0.040	0.093**	0.115*
	(0.030)	(0.032)	(0.044)
Log HP/weight (γ_a)	-0.212***	-0.525***	-0.589***
	(0.016)	(0.016)	(0.013)
Log volume (γ_s)	-0.346***	-0.456***	-0.508***
	(0.017)	(0.006)	(0.010)
Pickup	-0.012	-0.077***	-0.098***
	(0.012)	(0.008)	(0.005)
$\frac{\partial \ln \text{mpg}}{\partial \ln \text{HP/Wt}}\vert_{\Delta c=0} \left(\frac{\alpha_a}{\alpha_g} = \frac{\gamma_a}{1-\gamma_g}\right)$	-0.221	-0.579	-0.666
	0.011	0.035	0.046
$\frac{\partial \ln \text{mpg}}{\partial \ln \text{volume}}\vert_{\Delta c=0} \left(\frac{\alpha_s}{\alpha_g} = \frac{\gamma_s}{1-\gamma_g}\right)$	-0.361	-0.503	-0.574
	0.010	0.024	0.039

Note: This table presents coefficient estimates from equation (49) in which we interact gas prices, acceleration, size, and the pickup dummy with era dummies (1995–2001, 2002–2009, and 2010–2017) to allow the coefficients to differ over time. All columns 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. *, **, and *** indicate statistical significance at the 5%, 1%, and 0.1% levels, respectively. Data source: NHTS, EPA, Transport Canada.

5.1.2 Biased technical change

Have the slopes of the isocost curves shifted over time? Table 3 shows OLS results from equation (49) in which we allow the coefficients on gas prices, acceleration, size, and the pickup dummy to differ across three different eras: 1995–2001, 2002–2009, and 2010–2017. This regression is based on model (6) from table 3, which includes controls for vehicle-miles traveled (VMT) and demographics. The results show increasingly steep isocost curves over time in both fuel economy vs. acceleration space and fuel economy vs. size space (see calculations at bottom of table). In particular, we estimate that a 10% reduction in acceleration led to a 2.21% gain in fuel economy in the late 1990s but a 6.66% gain in the 2010s, with costs held fixed. Likewise, a 10% reduction in car size led to a 3.61% fuel economy gain in the 1980s but a 5.74% gain in the 2010s. These results are consistent with technical change that is substantially biased toward fuel economy.

5.1.3 Robustness

We address three potential concerns with additional analysis, namely, 1) that the number of gasoline price lags used is arbitrary, 2) that the shape imposed by the functional form is overly constrictive, and 3) that the modeling assumptions about fuel economy and GHG standards may not hold.

Table 7 in the appendix repeats regression (5) in table 2 but varies the number of lags of the log gasoline price variable. The sums of the coefficients between regressions (2) and (4) are quite similar, ranging between 0.93 and 0.75. This gives us confidence that our estimates are not particularly sensitive to the number of lags beyond 1. The coefficient in the first regression, with no lags, is 0.65, and while it is not statistically different from our baseline estimate, is smaller than the others. This is not surprising as we do not know when in the year a new car is sold while we apply the t=0 gas price average of the entire calendar year.

Table 8 shows four regressions that interact log attributes with the tercile of that attribute. Regression (1) estimates a common attribute elasticity for all years, while regressions (2), (3), and (4) break up our sample period into three groups. Otherwise, these regressions use the same variables as regression (6) in table 2.[21] While the coefficients on attributes are statistically different from each other in many cases, the estimates lie in a remarkably small range—the lowest to the highest coefficients are all within 10% of each other. Further, regressions (2)–(4) show estimates of biased technical change similar to those in columns (1)–(3) in table 3. This gives us confidence that our functional form fits the data well and that our interpretation of biased technical change does not depend on this functional form.

Figure 13 shows estimates of the shadow cost of fuel economy and GHG standards from the literature. The panel on the left shows the shadow cost level, while the figure on the right shows the shadow cost as a percent of discounted lifetime fuel savings from a 1-mile-per-gallon (MPG) improvement under conservative assumptions. Note that there was no credit trading between manufacturers in corporate average fuel economy (CAFE) before 2007. This means that there may be different shadow costs of the standard for each manufacturer and for each vehicle class (car or light truck). For this period, Anderson and Sallee (2011b) use a loophole in CAFE—fuel economy credits from flex-fuel technology—to estimate the shadow costs for all manufacturers that were using this loophole. While the number of manufacturers and years is incomplete, the range is not large—between \$12 and \$23 or 3% to 4% of future discounted fuel savings. The penalty for being out of compliance was considerably higher—\$55 or between 8% and 12%.

Appealing to the law of one price,[22] we test three assumptions about shadow costs during this period: 1) that the shadow costs were equal to the \$55 penalty, 2) that the shadow costs were equal to the mean of the estimates from Anderson and Sallee (2011b), and 3) that the shadow costs were 0. All three assumptions include the estimated shadow cost from Yeh et al. (2021) after 2007. Using a first-order Taylor-series approximation about 0, we find that the effect that the shadow cost or tax (τ) has on $\log pm$ is to add $\frac{\tau}{pm}$ to the log gas price. m is lifetime discounted miles. Our method for computing m is described in section

[21] We use regression (6) for convenience, but the results are the same if we use lags of the log gas price as in regression (5) instead of the log of the three-year average gas price.

[22] This touches on issues of markups and market power that are beyond the scope of this paper. Estimates of individual manufacturer shadow costs are incomplete. Thus, estimating a model with manufacturer-specific shadow costs would require restrictions on our data. However, the assumption of a single shadow cost favors estimated changes to the year dummy variable and the coefficient on γ_g compared to and we estimate these effects to be small.

6.

Figure 14 shows the mean values of $\tau/p\hat{m}$ for high, medium, and low τ. Table 6 shows the coefficient estimates for the high, intermediate, and low shadow cost assumptions. Figure 15 shows the year dummy coefficient estimates. These estimates are remarkably consistent across assumptions and with our baseline estimates.

5.2 Consumer preferences

Having estimated the cost parameters, we now turn to estimating the preference parameters. We begin by estimating the preference parameters on size (μ_s) and acceleration (μ_a) by regressing log size and log acceleration on the variation in these attributes' marginal costs induced by technical change and variation in the other attributes. Table 4 reports the results from this regression.

Column (1) reports coefficients from a regression of log size on θ, $\alpha_s \ln s$, and $\alpha_a \ln g$, where we draw our cost parameters from regression (6) in table 2, which imposes time-constant αs.[23] Our theory implies that the coefficients on these variables should be identical; they are indeed strikingly similar. Column (2) imposes this restriction directly by including a single variable equal to the sum of these three variables, leading to coefficient estimate $\lambda_s = 0.093$. The bottom of the table reports the implied preference parameter ($\mu_s = 1/\lambda_s - \alpha_s$) and corresponding elasticity of demand for car size ($-1/(1 + \mu_s)$). The results imply that the elasticity of demand for size is -0.148. Note that column (3) further controls for year dummies, state dummies, and our full suite of demographic controls, including miles, income, and other variables. The estimated coefficients barely budge, alleviating concerns that unobserved shifts in demand for car size are correlated with acceleration and fuel economy, biasing our estimates. Columns (4)–(6) repeat this exercise for acceleration. The results imply that the elasticity of demand for acceleration is -0.185. Inelastic demand for size and acceleration limits the incentive to scale back these attributes in response to higher gas prices or biased technical change.

We finally turn to estimating the preference shifters for size and acceleration (β_s and β_a) relative to fuel economy (β_g) directly from equations (15)–(17). These equations relate the attribute trade-offs in utility (MRS) to the slopes of the isocost curves (MRTSA). We use these theoretical conditions, along with our estimated αs and μs (based on regression results with the most controls above), to infer each household's log marginal benefit for size in terms of forgone fuel economy ($\ln \beta_s/\beta_g$) and likewise for acceleration ($\ln \beta_a/\beta_g$).[24] Figure 11 in the appendix plots the joint distribution of these values for each car unweighted model represented in our dataset (scatter diagram), along with the marginal distributions across all survey

[23] Allowing for time-varying αs in the first step would require nonlinear estimation of μ_s in the second step, given that α_s appears in the denominator of $1/(\alpha_s + \mu_s)$ with μ_s.

[24] Given an estimate of the present-discounted value of lifetime fuel expenditures, we could estimate the willingness to pay (WTP) for size and acceleration in dollars.

Table 4: *Second-stage regressions to identify demand elasticities*

	(1) Vol	(2) Vol	(3) Vol	(4) HP/Wt	(5) HP/Wt	(6) HP/Wt
$\theta - \alpha_a \ln a + \alpha_g \ln g$ (λ_s)		0.093***	0.092***			
		(0.002)	(0.002)			
$\theta - \alpha_s \ln s + \alpha_g \ln g$ (λ_a)					0.102***	0.101***
					(0.001)	(0.001)
θ	0.122***			0.128***		
	(0.001)			(0.002)		
$-\alpha_a \ln a$	0.112***					
	(0.001)					
$-\alpha_s \ln s$				0.125***		
				(0.001)		
$\alpha_g \ln g$	0.086***			0.095***		
	(0.001)			(0.001)		
$\mu_s = 1/\lambda_s - \alpha_s$		5.739	5.885			
		0.232	0.199			
$\mu_a = 1/\lambda_a - \alpha_s$					4.394	4.457
					0.084	0.055
$\frac{\partial \ln s}{\partial mc_s} = -\frac{1}{1+\mu_s}$		-0.148	-0.145			
		0.005	0.004			
$\frac{\partial \ln a}{\partial mc_a} = -\frac{1}{1+\mu_s}$					-0.185	-0.183
					0.003	0.002

Note: This table presents second-stage coefficient estimates from regressions of log volume and log HP/Wt on the calculated marginal cost shifters as indicated in equations (24) and (25). Columns (1), (2), (4), and (5) contain no additional controls. Columns (3) and (6) control for state dummies, year dummies, and our full set of demographic controls. Standard errors in parentheses are clustered by state and do not currently account for first-stage estimation error (updated results coming soon). *, **, and *** indicate statistical significance at the 5%, 1%, and 0.1% levels, respectively.
Data source: NHTS, EPA, and Transport Canada.

respondents (histograms). The figure shows wide dispersion in preferences, where a one-log-point change implies a $\exp(1) \approx 2.72$ times higher marginal WTP (MWTP) for size and acceleration relative to fuel economy. The figure also shows a strong positive correlation between preferences for size and power. Figure 12 in the appendix shows that these preferences are strongly associated with income, population density, and family size. Large, high-income families living in suburban and rural areas drive larger and more powerful cars. Thus, viewed through the lens of our theoretical model, in which size and power come at the expense of lower fuel economy (and therefore higher fuel costs), these households have the highest willingness to pay.

Figure 7: *Discrete technology adoption*

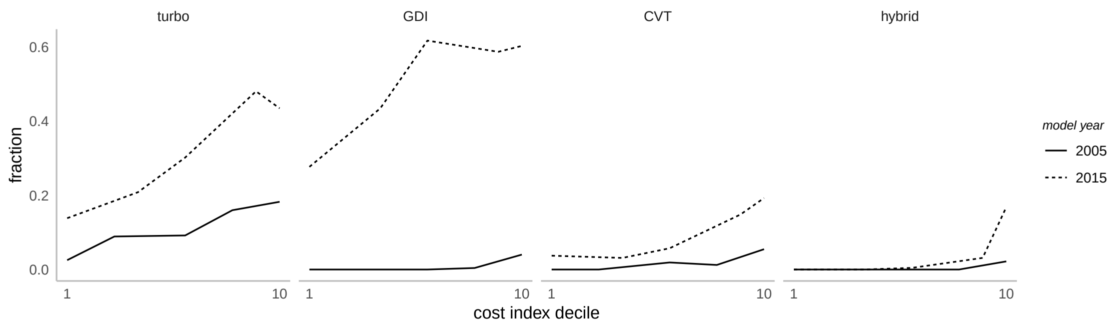

Note: Sales-weighted fractions of discrete technology adoption for cost index decile for 2005 model years and 2015 model years. Deciles are computed within each year, and the estimated cost index uses coefficients from our main specification.

5.3 Fuel-saving technologies are not evenly adopted in the fleet

Figure 7 shows the fraction of adoption for deciles of an estimated drivetrain cost index. The estimated drivetrain cost index is computed with our main specification and is a weighted sum of vehicles attributes:

$$cost\ ind. = \hat{\alpha}_s \ln vol + \hat{\alpha}_a \ln hpwt + \hat{\alpha}_g \ln mpg. \tag{50}$$

We compute the index at the trim level using WardsAuto data, which contain some detailed technology data. We then compute the within-year decile and plot the fraction of vehicles within each decile with each of four technologies: turbo- and superchargers, gasoline direct injection (GDI), continuous variable transmission (CVT), and hybrid gasoline–electric motors. We show these for two years, 2005 and 2015, to show both the increasing adoption and the pattern of adoption in an early and a later year. The figure shows that for both the earlier and the later year, costlier vehicles—those with higher attribute levels—are the most likely to have adopted any given technology.

We test this prediction more formally via regression of discrete technology adoption on car attributes that are relevant to consumers. We estimate equations of the following form:

$$technology_{it} = \phi_{bt} + \gamma_s \ln volume_{jt} + \gamma_a \ln hpwt_{jt} + \gamma_g \ln mpg_{jt} + \varepsilon_{jt}, \tag{51}$$

where $technology_{jt}$ is the technology dummy for a vehicle model trim j; ϕ_{bt} is the vehicle body style by year fixed effect and the other variables are defined in the same way as above. For this exercise, we opt not to match vehicle trims with the consumer microdata, as we would lose the trim-level detail, and we instead use trim-level sales weights from the EPA. Vehicle body styles here serve to capture larger consumer groups that have distinct preferences for acceleration and fuel economy.

The OLS results are given in table 5. While acceleration and fuel economy are physically affected by all of

Table 5: *Discrete technology adoption correlations*

	Turbo	GDI	CVT	Hybrid	Turbo	GDI	CVT	Hybrid
Est. cost index	0.007***	0.010***	0.035***	0.019***				
	(0.001)	(0.001)	(0.001)	(0.000)				
Log vol.					0.162***	−0.061***	0.014	0.116***
					(0.017)	(0.012)	(0.016)	(0.009)
Log hp/wt.					0.191***	0.108***	−0.381***	−0.336***
					(0.011)	(0.009)	(0.011)	(0.006)
Log mpg					0.161***	0.054***	0.331***	0.145***
					(0.015)	(0.011)	(0.014)	(0.007)
Num.Obs.	26 096	26 096	26 096	26 096	21 318	26 101	21 318	21 318
R2	0.031	0.148	0.143	0.069	0.138	0.154	0.272	0.277
R2 Adj.	0.031	0.148	0.143	0.069	0.138	0.154	0.272	0.277
RMSE	0.36	0.33	0.20	0.12	0.38	0.33	0.22	0.15
Year FE	X	X	X	X				
Veh. type × year FE					X	X	X	X

Note: Models (1)–(4) estimate the correlation between our estimated cost index from equation (50) and discrete technology adoption. Models (5)–(8) estimate the correlation between log attributes and discrete technology adoption within vehicle body types: pickup trucks, SUV/CUVs, sedans/wagons, hatchbacks, coupes/convertibles.
+ $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

these technologies, volume—which is highly correlated with all technologies within a vehicle body style—is not. These results further support our hypothesis that drivers who want large vehicles will adopt technologies at a faster rate, with other preferences held fixed. Additionally, the results support our concern that the correlation of attributes alone cannot yield estimates of isocost curves. Instead, technologies that contribute to drivetrain cost are highly correlated with attribute levels, which supports our approach.

6 What explains trends in car attributes?

In this section, we describe our counterfactual simulations, which explore the contribution of gas prices, technology, and preferences to trends in car attributes during 1995–2017. We begin by detailing our methods. We then describe our simulation results.

6.1 Model estimation and calibration

Estimating costs. We begin by re-estimating model (5) from 2, which controls for age, log miles, and our full set of demographics, but add two new variables to this model: log size and log acceleration interacted with a linear time trend. Thus, consistent with the regression results in table 3, we allow for biased technical change in addition to an overall improvement in technology. Figure 8 plots attribute-neutral technical change ($\theta/(1+\alpha_g)$) and trending cost parameters (αs) from this regression (upper-left panel).

Estimating demand. Given our estimated cost parameters, we estimate the elasticity of demand for size following the approach in section 5.2 by regressing log size on a linear combination of size and fuel economy. We use the same approach to estimate the elasticity of demand for acceleration. In constructing these linear combinations, we impose constant rather than time-varying cost parameters for size and acceleration (α_s and α_a).[25] We estimate demand elasticities similar to those reported in table 2 above.

Given our estimated cost parameters and demand elasticities (αs and μs), we calculate the heterogeneous preference shifters (βs) for each car in our sample according to the $MRTAS = MRS$ conditions in equations (15)–(16). Note that these conditions yield preference parameters for size and acceleration relative to fuel economy. Thus, we must calibrate fuel economy preferences to recover absolute preferences for size and acceleration.

We calibrate absolute preferences for fuel economy as state-level gas prices times lifetime miles for every car in our dataset. In log form:

$$\ln(\beta_g) = \ln p + \ln m(0) + \ln(r_t + \rho + \delta), \tag{52}$$

where $\ln p$ is the log price of gasoline, $m(0)$ is miles driven in the car's first year, r_t is the interest rate, ρ is the annual rate of scrappage, and δ is the annual rate of decay in miles conditional on survival. We calibrate $\ln p$ as the log of the three-year moving average of state-level gas prices, as in our econometric estimation. We calibrate $\rho = 0.05$ based on data in NHTSA (2006).[26] We estimate $m(0)$ and δ in our data by regressing log miles on car age, state dummies, and our full set of demographics, along with dummies for the NHTSA wave (2001, 2009, and 2017). The coefficient on age yields δ, while the fitted value when $age = 0$ yields $m(0)$. We find no evidence of a shift in first-year miles by survey year, i.e., the coefficients on NHTS wave are small and insignificant. Thus, we impose that these coefficients are zero when predicting $m(0)$. We also find no evidence that δ shifted over time. Thus, first-year miles $m(0)$ vary only in the cross-section based on observed demographics. Finally, we calibrate r_t based on the real interest rate on 48-month new-car loans from the St. Louis Fed. Figure 8 plots the resulting trends in preferences for fuel economy, size, and acceleration, all relative to their 1995 levels (upper-right panel).

Calibrating the cost residual. Our final step is to calibrate the cost shifter (k) based on equation (22), given our above calibrations for costs (αs and θ) and preferences (βs). Note that k is a free parameter in our model, which allows us to match baseline fuel economy for every individual car in our dataset. Thus, our simulation model exactly replicates the baseline trends in car attributes. One implication is that our careful calibration of first-year miles ($m(0)$) above makes zero difference in our simulations focused on mean car attributes. Whatever we choose for lifetime miles, we calibrate a cost residual (k) that rationalizes a car's

[25] Mechanically, we do this by regressing log fuel economy on log size and acceleration while imposing the cost parameter on fuel economy (α_g) that we estimate with the time-trending cost parameters; this yields a weighted average of the time-varying coefficients.

[26] Annual scrap rates conditional on survival change over time as a car ages: they are 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 may be 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 the weights given by discount factor $exp(-0.1age)$ to reflect both time discounting and the annual decay in miles. There is evidence that scrap rates have fallen in recent years, but mostly after 2017.

observed baseline fuel economy. Meanwhile, since our counterfactual simulations hold k and $m(0)$ fixed, they play zero further role. Figure 8 replicates the baseline trends in car attributes as generated by our simulation model (middle-left panel).

6.2 Structural decomposition via counterfactual simulation

We perform four counterfactual simulations that, respectively, hold gas prices ($\ln p$), technical change (θ & αs), attribute-biased technical change (ratios of αs), and preferences (β_a & β_s) fixed at their 1995 levels. We impose these 1995 parameter values one at a time and then simulate counterfactual choices of size, acceleration, and fuel economy for each car in our dataset according to equations (18)–(20), which give optimal attribute choices as a function of model primitives. Finally, we calculate mean attributes in every year. Figure 8 plots the resulting trends in fuel economy (middle right), size (bottom left), and acceleration (bottom right). We express counterfactual attributes as changes relative to the baseline level in each year, e.g., 0.1 in 2005 represents a 10% increase relative to that year's baseline value.

Overall, gas prices, technology, and preferences are all important for explaining trends in car attributes. When we hold gas prices fixed at their 1995 levels ("no gas"), fuel economy in 2017 is 7% lower than baseline, while size and acceleration are 4% higher. These effects are totally in line with expectation.

When we hold car technology constant ("no tech"), fuel economy is 28% lower in 2017, while size is unchanged and acceleration is 16% higher. These results illustrate the nuanced and multidimensional nature of technical change that we estimate in our model. Attribute-neutral technical change ($\Delta\theta$) lowers costs across the board and leads to improvements in all attributes. Meanwhile, biased technical change favoring fuel economy, as reflected in a rotating MRTSA (upward trends in α_a and α_s), leads to an increase in fuel economy and a decrease in size and acceleration. Intuitively, the opportunity cost of performance is increasing over time, causing consumers to choose more fuel economy and less performance. This effect is particularly pronounced for acceleration, whose opportunity cost increases by a factor of three over our sample period. Thus, biased technical change overpowers attribute-neutral technical change, leading to a decrease in acceleration. This effect is less pronounced for size, however, such that the attribute-neutral and biased technical change cancel each other out.

We next shut down the biased component of technical change ("no bias"). We do this by allowing the *sum* of the cost parameters (αs) to trend according to their baseline values while holding their ratios fixed at the 1995 levels. In this case, fuel economy is 9% lower in 2017, while size is 4% higher and acceleration is 21% higher. These results illustrate how biased technical change favoring fuel economy leads consumers to choose smaller, slower, more efficient cars.

Finally, when we hold consumer preferences for size and acceleration constant ("no prefs"), fuel economy in 2017 increases by 52%, while size decreases by 28% and acceleration decreases by 30%. These results

Figure 8: *Trends in technology, preferences, baseline car attributes, and counterfactual changes*

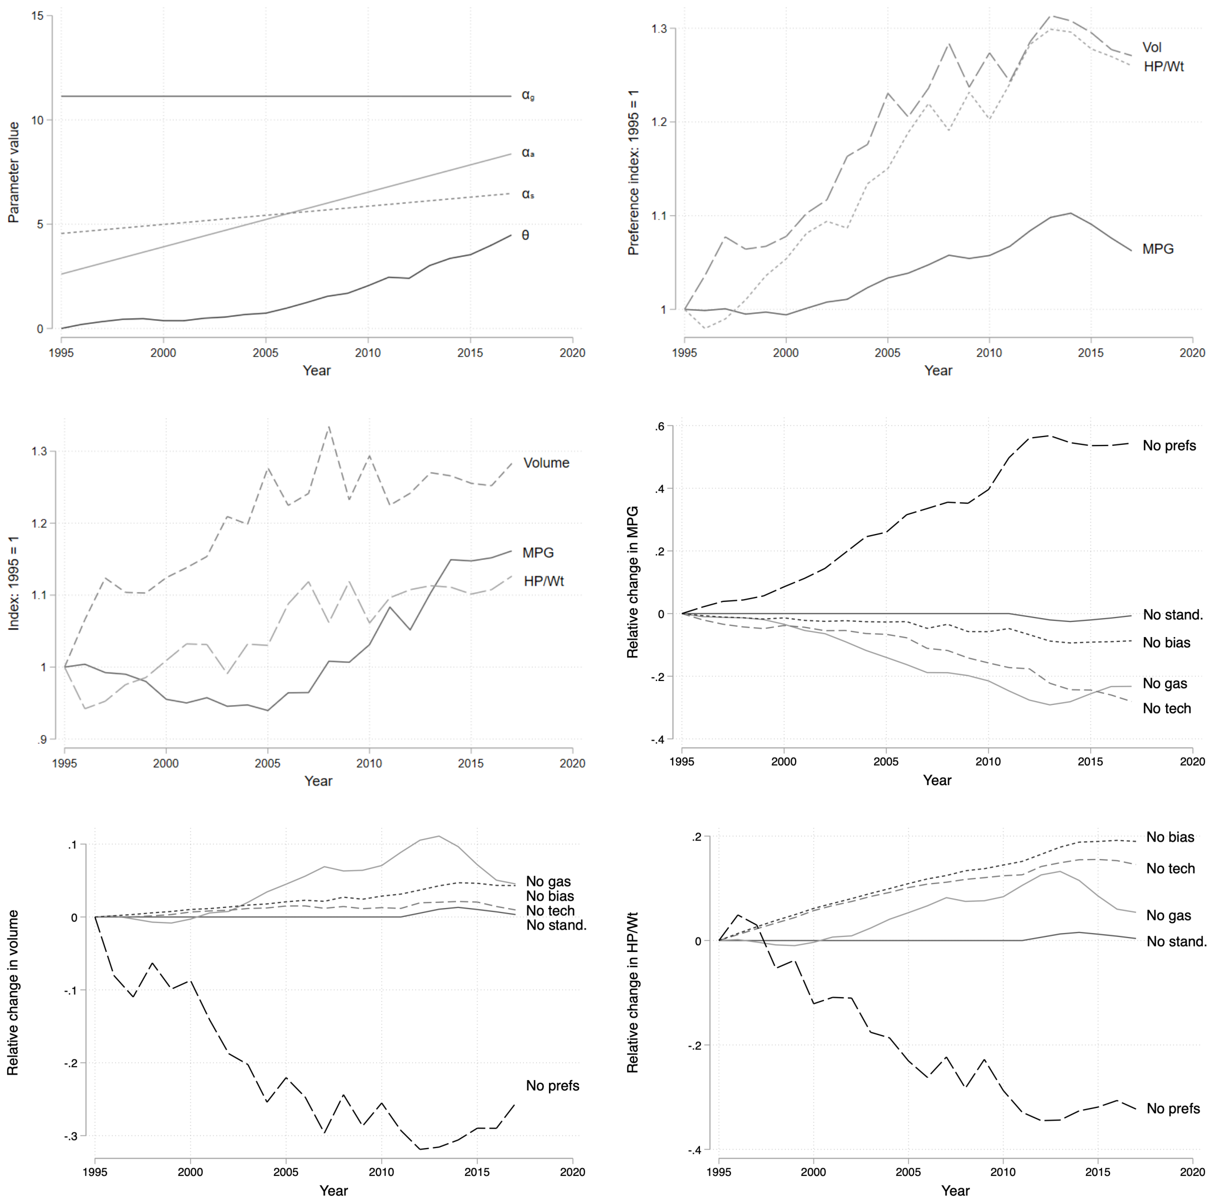

Note: This figure presents the parameter inputs to our counterfactual simulations, along with our results. The top-left panel shows the trend in attribute-neutral technical change (θ) and cost parameters (αs). The top-right panel shows the trends in average log consumer preferences ($\ln \beta_g$, $\ln \beta_a$, and $\ln \beta_s$) relative to their 1995 values. The middle-left panel shows trends in baseline attributes relative to their 1995 values as generated by our simulation model (the same as the actual trends in figure 1). The remaining panels show changes relative to baseline values in each year for fuel economy (middle right), volume (bottom left), and acceleration (bottom right) under four counterfactual scenarios. These scenarios set various model inputs constant at their 1995 levels: gas prices (labeled "no gas"), car technology ("no tech"), attribute-biased technical change ("no bias"), and consumer preferences ("no prefs"). See text for details.

illustrate the tremendous growth in consumer preferences for size and acceleration over time. They also illustrate an important interaction between preferences and biased technical change: increases in size and acceleration driven by consumer preferences entail larger and larger unrealized gains in fuel economy because of the increasing opportunity cost of performance.

7 Discussion and conclusion

Why have car size, acceleration, and fuel economy all increased in recent years, reversing the trend of the previous three decades? We find that preferences for size and acceleration have substantially increased while technical change has been biased toward fuel economy. This fuel-economy bias has increased the opportunity cost of size and acceleration, while demand for size and acceleration have also grown. The two trends have worked to counteract each other such that consumers are buying more of all attributes instead of trading off one for another. At the same time, rising gas prices and tightening fuel-economy standards have made fuel economy higher now than in the past, relative to reductions in size and acceleration. Biased technical change has made any binding standard approximately one-third cheaper than it would be under attribute-neutral technical change. Perhaps this is why the EPA has recently proposed the first standards in many years that are likely to be binding.

7.1 Implications for electric vehicles

Incorporating electric vehicles (EVs) into our model would require adding an essential car attribute: range. An EV's range depends on battery size and car weight. Larger batteries extend range but cost more and increase weight, reducing acceleration and energy efficiency. Conversely, charging speed is unlikely to interact with the marginal cost of size, acceleration, efficiency, or range in a meaningful manner.

We do not directly study EVs. Nonetheless, our model yields valuable insights about their likely impacts. Electric drivetrains simultaneously reduce cost-per-mile and improve acceleration. Thus, as EV costs decrease, more consumers will adopt EVs both to reduce fuel costs and to boost acceleration. Importantly, this dynamic implies that, in comparison to internal combustion vehicles, EVs will further reduce the trade-offs that drivers face when adhering to stringent standards. As standards become more demanding and EVs become more affordable, consumers who meet the standard with an EV will not be trading off between acceleration and efficiency but obtaining more of both.

Our analysis of discrete technology suggests that the adoption of EVs will first occur among drivers who prioritize reducing fuel costs, crave speed, and, to a lesser extent, like large cars. Indeed, EVs are already poised to replace highly efficient gas-powered cars, sports cars, and some SUVs. As standards become more stringent and promote greater EV adoption, EVs will replace larger, slower, and less efficient vehicles. As a consequence, increased incentives for EV adoption are likely to lead to increasing marginal social benefits from the forgoing of gas cars. The economic intuition is that EV technology has entered a market where consumers have already sorted by preferences for efficiency and so those who adopt first will replace the most efficient cars while those who adopt later will replace gas guzzlers. The effect may be magnified by EV range, which is currently an expensive attribute to produce, leaving those who drive the most to wait to adopt until battery costs come down. This deviates from the conventional model that assumes diminishing marginal benefits with escalating incentives. If this prediction holds, studies that measure the marginal benefits of EVs, such as Holland, Mansur, Muller, and Yates (2019), would need a very different interpretation.

7.2 Implications for fuel-economy standards

Our equilibrium approach contrasts with the EPA's regulatory impact analysis at one extreme and Knittel's (2011) empirical analysis at the other. The EPA's analysis holds car attributes fixed and calculates the cost of adding energy-saving technology to meet standards. Meanwhile, Knittel's (2011) analysis attempts to hold costs fixed and asks whether standards can be met via attribute-neutral technical change without substantial reductions in size and acceleration. We show that improving fuel economy conditional on attributes has become relatively more cost-effective over time, such that tighter standards reduce size and acceleration less than in the past. In short, the car market has become more aligned with the EPA's modeling approach over time. Among of the strengths of our model are its simplicity and transparency. These make it a good candidate for validating more complex regulatory analyses, such as the EPA's OMEGA model. The EPA's existing ALPHA model already generates an attribute-cost surface. The output from ALPHA could easily be paired with our model to facilitate simplified regulatory analysis.

7.3 Implications for directed technical change

We conclude with five lessons for policies aiming to speed up 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, welfare may increase in some dimensions (e.g., pollution) and decrease in others (e.g., fatal accidents). Second, subsidies for discrete technologies will boost adoption among consumers who choose high levels of all attributes. Such subsidies will tend to be regressive since wealthier drivers buy larger, faster cars. Further, the boost to fuel economy from a hybrid engine or other similar technology will be reallocated in part to larger, faster cars in equilibrium, eating into the fuel-economy gains. Third, directed technical change that reduces the marginal cost of efficiency relative to the cost of other attributes will pull consumers toward more efficient cars, without the need for a carbon tax or efficiency standard, and dampen the effect of such policies on size and acceleration. Thus, relative to pricing externalities, supporting innovation may have important political economy benefits. Finally, the equilibrium effects of carbon taxes, efficiency standards, and directed technical change are all mediated by consumer preferences, and shifts in consumer preferences for other attributes can undermine policy objectives aimed at reducing carbon emissions. This concern becomes particularly pertinent if these shifts arise from an arms race in an attribute with untaxed externalities, such as vehicle size (see Jacobsen 2013; Anderson and Auffhammer 2014; Bento, Gillingham, and Roth 2017).

References

Acemoglu, Daron. 2002a. "Directed Technical Change." *The Review of Economic Studies* 69 (4):781–809.

———. 2002b. "Technical Change, Inequality, and the Labor Market." *Journal of Economic Literature* 40 (1):7–72.

Acemoglu, Daron, Philippe Aghion, Leonardo Bursztyn, and David Hemous. 2012. "The Environment and Directed Technical Change." *American Economic Review* 102 (1):131–166.

Aghion, Philippe, Antoine Dechezleprêtre, David Hemous, Ralf Martin, and John Van Reenen. 2016. "Carbon Taxes, Path Dependency, and Directed Technical Change: Evidence from the Auto Industry." *Journal of Political Economy* 124 (1):1–51.

Anderson, Michael L and Maximilian Auffhammer. 2014. "Pounds That Kill: The External Costs of Vehicle Weight." *Review of Economic Studies* 81 (2):535–571.

Anderson, Soren T., Ryan Kellogg, and James M. Sallee. 2013. "What Do Consumers Believe About Future Gasoline Prices?" *Journal of Environmental Economics and Management* 66 (3):383–403.

Anderson, Soren T and James M Sallee. 2011a. "Using Loopholes to Reveal the Marginal Cost of Regulation: The Case of Fuel-economy Standards." *American Economic Review* 101 (4):1375–1409.

———. 2011b. "Using Loopholes to Reveal the Marginal Cost of Regulation: The Case of Fuel-Economy Standards." 101 (4):1375–1409. URL http://pubs.aeaweb.org/doi/10.1257/aer.101.4.1375.

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

Bates, Michael and Seolah Kim. 2022. "Estimating the Price Elasticity of Gasoline Demand in Correlated Random Coefficient Models with Endogeneity."

Bento, Antonio, Kenneth Gillingham, and Kevin Roth. 2017. "The Effect of Fuel Economy Standards on Vehicle Weight Dispersion and Accident Fatalities." URL http://www.nber.org/papers/w23340.pdf.

Calel, Raphael and Antoine Dechezleprêtre. 2016. "Environmental Policy and Directed Technological Change: Evidence from the European Carbon Market." *Review of Economics and Statistics* 98 (1):173–191.

Davis, Lucas W and Lutz Kilian. 2011. "Estimating the Effect of a Gasoline Tax on Carbon Emissions." *Journal of Applied Econometrics* 26 (7):1187–1214.

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

EPA. 2020. "Inventory of U.S. Greenhouse Gas Emissions and Sinks: 1990-2018." *United States Environmental Protection Agency* .

Holland, Stephen P, Erin T Mansur, Nicholas Z Muller, and Andrew J Yates. 2019. "Distributional Effects of Air Pollution from Electric Vehicle Adoption." *Journal of the Association of Environmental and Resource Economists* 6 (S1):S65–S94.

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

Jacobsen, Mark R. 2013. "Fuel Economy and Safety: The Influences of Vehicle Class and Driver Behavior." *American Economic Journal: Applied Economics* 5 (3):1–26.

Jacobsen, Mark R, Christopher R Knittel, James M Sallee, and Arthur A Van Benthem. 2020. "The use of Regression Statistics to Analyze Imperfect Pricing Policies." *Journal of Political Economy* 128 (5):1826–1876.

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

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

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

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

Leard, Benjamin, Joshua Linn, and Yichen Christy Zhou. 2023. "How Much Do Consumers Value Fuel Economy and Performance? Evidence from Technology Adoption." *Review of Economics and Statistics* 105 (1):158–174.

Lin, Yujie and Joshua Linn. 2023. "Environmental Regulation and Product Attributes: The Case of European Passenger Vehicle Greenhouse Gas Emissions Standards." *Journal of the Association of Environmental and Resource Economists* 10 (1):1–32.

MacKenzie, Don and John Heywood. 2012. "Acceleration Performance Trends and Evolving Relationship between Power, Weight, and Acceleration in U.S. Light-Duty Vehicles: Linear Regression Analysis." 2287 (1):122–131. URL http://journals.sagepub.com/doi/10.3141/2287-15.

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

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

NHTSA. 2006. “Vehicle Survivability and Travel Mileage Schedules.”

Reserve, Federal. 2019. “Consumer Credit-G. 19, January 2019.” *Retrieved August* 20:2021.

Shahed, SM and Karl-Heinz Bauer. 2009. “Parametric Studies of the Impact of Turbocharging on Gasoline Engine Downsizing.” *SAE International Journal of Engines* 2 (1):1347–1358.

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

Whitefoot, Kate S., Meredith L. Fowlie, and Steven J. Skerlos. 2017. “Compliance by Design: Influence of Acceleration Trade-offs on CO_2 Emissions and Costs of Fuel Economy and Greenhouse Gas Regulations.” 51 (18):10307–10315. URL http://pubs.acs.org/doi/10.1021/acs.est.7b03743.

Whitefoot, Kate S. and Steven J. Skerlos. 2012. “Design incentives to increase vehicle size created from the U.S. footprint-based fuel economy standards.” 41:402–411. URL https://linkinghub.elsevier.com/retrieve/pii/S0301421511008779.

Yeh, Sonia, Dallas Burtraw, Thomas Sterner, and David Greene. 2021. “Tradable Performance Standards in the Transportation Sector.” *Energy Economics* 102:105490.

A Mathematical appendix

A.1 Second-order sufficient conditions

The second derivatives of the utility function are

$$u_{ss} = -2\frac{\beta_s}{s^3} - \alpha_s(\alpha_s - 1)\frac{c}{s^2} \tag{53}$$
$$u_{aa} = -\alpha_a(\alpha_a + 1)\frac{c}{a^2} \tag{54}$$
$$u_{gg} = -\alpha_g(\alpha_g + 1)\frac{c}{g^2} \tag{55}$$
$$u_{sa} = \alpha_s\alpha_a\frac{c}{sa} \tag{56}$$
$$u_{sg} = \alpha_s\alpha_g\frac{c}{sg} \tag{57}$$
$$u_{ag} = -\alpha_a\alpha_g\frac{c}{ag}. \tag{58}$$

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

$$u_{ss} = -\alpha_s(\alpha_s + 1)\frac{c}{s^2} \tag{59}$$

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

$$D_2 = u_{ss}u_{gg} - u_{sg}^2 \tag{60}$$
$$= \alpha_s\alpha_g\frac{c^2}{s^2g^2}\Big((\alpha_s + 1)(\alpha_g + 1) - \alpha_s\alpha_g\Big) > 0. \tag{61}$$

The determinant of the third leading principle minor is

$$D_3 = u_{ss}u_{aa}u_{gg} - u_{ss}u_{ag}^2 + 2u_{sa}u_{sg}u_{ag} - u_{sa}^2u_{gg} \tag{62}$$
$$= -\alpha_s\alpha_a\alpha_g\frac{c^3}{s^2a^2g^2}\Big(\alpha_s + \alpha_a + \alpha_g + 2\alpha_s\alpha_a\alpha_g + 2\alpha_a\alpha_g\Big) < 0. \tag{63}$$

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

A.2 Margins and identification of parameters

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

$$\max_{g,a,s} u = y + v(g, a, s) - (1 + M) \times c(g, a, s). \tag{64}$$

The first order condition for g is then

$$[g] : \frac{\partial v}{\partial g} = (1+M)\frac{\partial c}{\partial g}. \tag{65}$$

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 manufacture and pickup dummy controls (X), the identifying assumption for unbiased $\hat{\alpha}_a$, $\hat{\alpha}_s$, and $\hat{\theta}$ is

$$\mathbf{E}[M|D,X,s,a,t] = \mathbf{E}[M|D,X]. \tag{66}$$

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 changes technical change. The assumption for identifying technical change is simply

$$\mathbf{E}[M|D,X,s,a,t] = \mathbf{E}[M|D,X,D,X,s,a]. \tag{67}$$

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.

A.3 Used cars, interest rates, and durability

Above, we implicitly assume that the costs and benefits of car ownership are realized simultaneously by a single consumer with fixed preferences at the time of purchase. In practice, cars are produced at a point in time, incurring some cost, and then generate flow benefits over time, both for the original owner and for any subsequent owners. We consider the implications in this section, highlighting the role improved durability as a mechanism of technical change and interest rates as a potential confounder. We also clarify how the preferences of used-car owners relate to up-front purchase decisions.

Assume that the marginal flow benefits from size, acceleration, and fuel economy are given by $b_s(t)$, $b_a(t)$, and $b_g(t)$. These values are indexed by car age (t) to capture physical depreciation and scrappage, maintenance costs, and declining miles over time as older cars are driven less intensively. Thus, our preference parameters from above (the βs) can be interpreted as the present-discounted values:

$$\beta_s = \int_{t=0}^{\infty} e^{-rt} b_s(t),\ \beta_a = \int_{t=0}^{\infty} e^{-rt} b_a(t),\ \text{and}\ \beta_g = \int_{t=0}^{\infty} e^{-rt} b_g(t), \tag{68}$$

where r is the rate of time discounting (interest rate). Note that the time-indexing accommodates both the possibility that marginal benefits will evolve over time for the original owner, as well as the possibility that the car will be sold in the used car market to an owner with different preferences. Note that a car buyer that only intends to own the car for a few years should still consider expected flow benefits to future owners, since these flow benefits determine resale value.

Recall our structural interpretation for the fuel economy preference parameter (β_g). A mile is a mile and a dollar is a dollar, regardless of who sits behind the wheel or fills the gas tank. Thus, from the perspective of the original car buyer, what matters is the owner's belief about the future price of gasoline and how much the car will be driven during its lifetime. Thus, we have $b_g(t) = pm(t) + r\tau$, where p is the gasoline price at the time of purchase assuming a no-change forecast (Anderson, Kellogg, and Sallee 2013), $m(t)$ is expected miles driven at some future date, and $r\tau$ is the fuel economy standard's credit price in annuity form. The present-discounted value is therefore given by:

$$\beta_g = p \int_{t=0}^{\infty} e^{-rt} m(t) dt + \tau, \tag{69}$$

which clarifies that the "miles" in our original formulation is the present-discounted sum of lifetime miles.

Suppose that annual miles decays exponentially: $m(t) = m(0)e^{-(\rho+\delta)t}$, where $m(0)$ is initial miles and $\rho + \delta > 0$ is the annual rate of exponential decay, reflecting both the rate of scrappage (ρ) and the decline in miles conditional on survival (δ). Then fuel economy preferences are given by:

$$\beta_g = \frac{pm(0)}{r + \rho + \delta}, \tag{70}$$

which shows that lower interest rates and improved durability (smaller r and $\rho+\delta$) both increase the up-front demand for fuel economy. Thus, equation (22) for the optimal choice of fuel economy conditional on other attributes becomes:

$$\ln g \approx \frac{k}{1+\alpha_g} - \frac{\theta}{1+\alpha_g} - \frac{1}{1+\alpha_g} \ln pm + \frac{1}{1+\alpha_g} \ln(r + \rho + \delta) + \frac{\alpha_s}{1+\alpha_g} \ln s + \frac{\alpha_a}{1+\alpha_g} \ln a, \tag{71}$$

where the approximation again follows from $\tau \approx 0$. This equation clarifies that a lower interest rate and improved durability both lead to higher fuel economy, conditional on other attributes. Durability is arguably an important mechanism of technical change, while the interest rates is an obvious confounder. Thus, we measure and control for both in our empirical application, to better identify technical change.

What about size and acceleration? In our empirical application, we analyze attribute choices for current car owners, who may have purchased the car long ago, or who may not even be the original buyers. How strongly should the flow benefits from car ownership correlate with expected up-front benefits at the time of purchase? To answer this question, we simply differentiate the equations in (68) with respect to the time-specific flow parameters ($\rho(t)$), which shows that flow benefits are all discounted by factor e^{-rt} in determining up-front benefits and therefore car attribute choices at the time of purchase. Intuitively, expectations for owner preferences in year 15 should matter little in determining up-front attribute choices, while preferences in years 1–2 should matter much more. To address this issue, in our empirical application we control for the demographics of current car owners to capture variation in flow benefits (ρ_s and ρ_a), and we interact these controls with car age to capture the effects of time discounting.

B Engineering simulation

Figure 9: *Simulated fuel economy vs. acceleration (via engine displacement)*

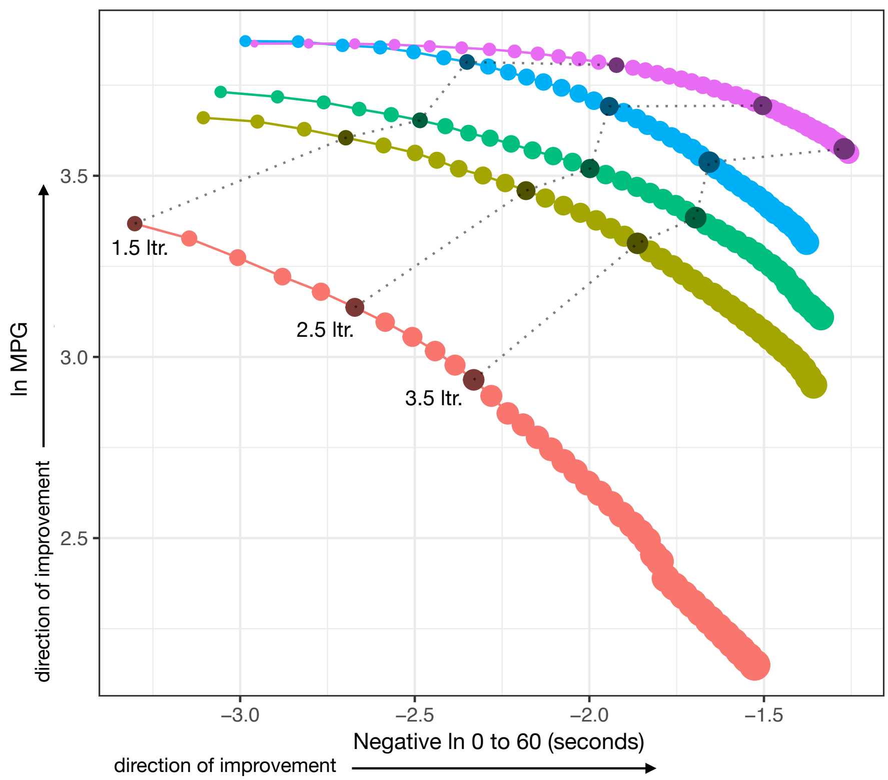

Note: This figure shows simulated log fuel economy vs. log acceleration for five distinct drive-trains, with different combinations of fuel economy and acceleration achieved by changing engine displacement. From bottom to top the five engines and transmissions are: 1980s carbureted (3-speed), 2007 Toyota 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 of displacement, power, and fuel economy, we show evidence from a physics-based full car simulation model—the Environmental Protection Agency's Advanced Light-Duty Powertrain and Hybrid Analysis (ALPHA) tool (Dekraker, Barba, Moskalik, and Butters (2018)), to investigate how the trade-off between fuel economy and power has shifted over time. Figure 9, which is reproduced from data in Moskalik, Bolon, Newman, and Cherry (2018), summarizes this change for five simulated cars. Each of the cars uses the same midsize sedan; the only difference is the drivetrain technology. This allows the simulation to maintain the same road-load across all simulated cars.[27] 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

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

to a six-speed transmission; (4) a 2017 Honda turbo-charged engine 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 are improving 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 has been greatest for smaller engines, while the gain in fuel economy has been largest for large engines (see the connecting dotted lines). Engines have clearly improved across-the-board. The correspondence between fuel economy and acceleration has flattened over time. A 1% decrease in acceleration is associated with a 0.75% gain in fuel economy in the 1980s, but only 0.33% gain in 2017. Thus, the opportunity cost of improving fuel economy via reductions in engine size has increased—improved fuel economy now comes at the expense of a much larger reduction in acceleration.

C Tables and figures

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

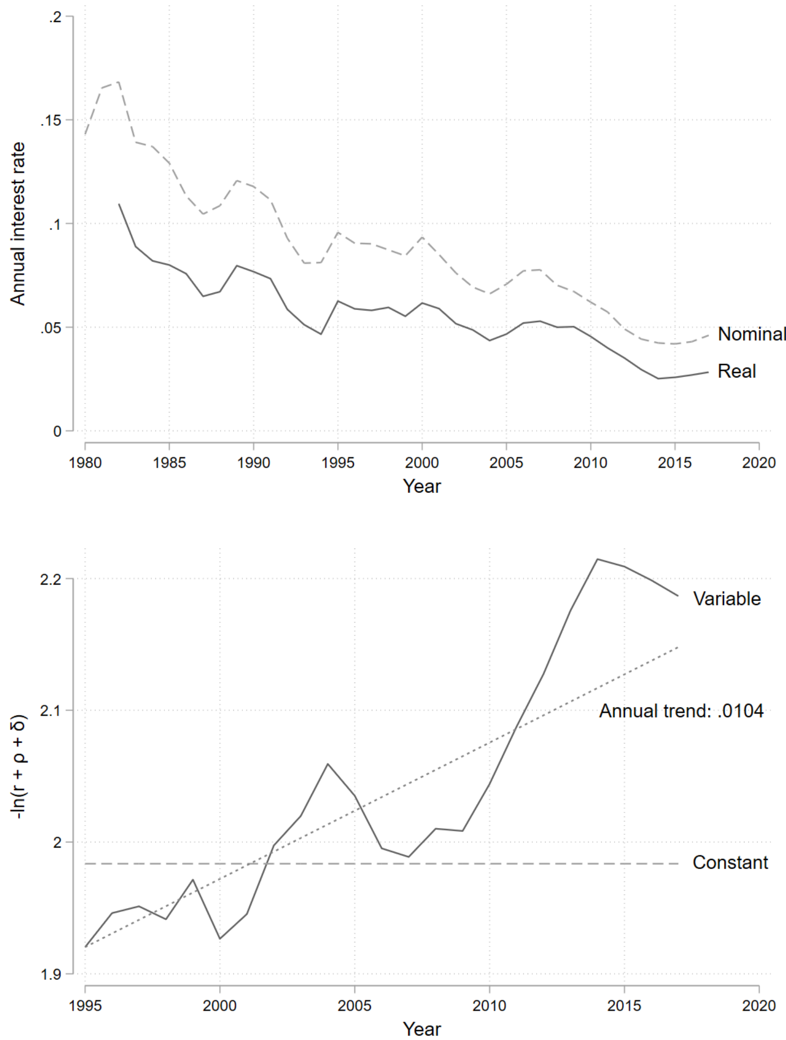

Note: The top panel plots the trend in nominal and real interest rates on 48-month new-auto loans. The bottom panel plots the trend of the present value multiplier, which is a function of the real interest rate.
Data source: Federal Reserve Bank of St. Louis.

Figure 11: *Consumer preferences*

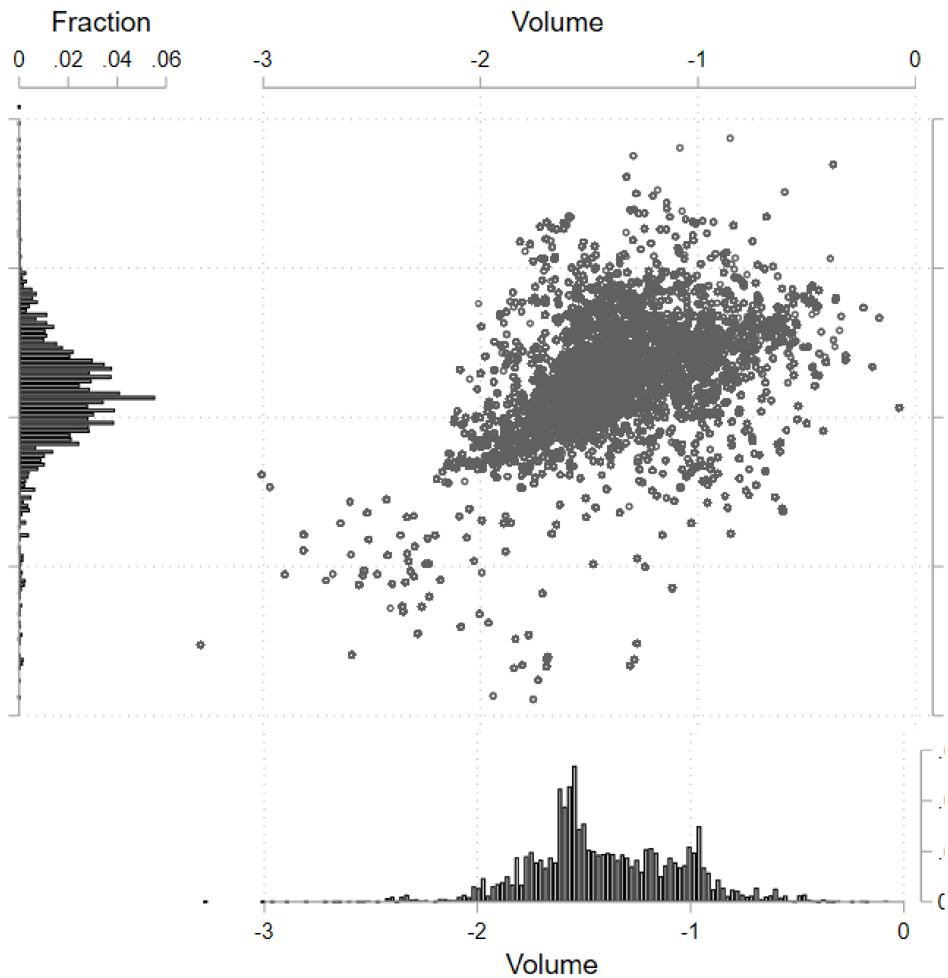

Note: This figure plots estimated preference parameters for acceleration ($\ln \beta_a/\beta_g$) and size ($\ln \beta_s/\beta_g$) for every car purchase in our final estimation sample, covering NHTS waves 2001, 2009, and 2017. The scatter diagram in the center plots preferences for size vs. acceleration, while the histograms along the axes show the full distributions across all years. We measure preferences for size and acceleration relative to preferences for fuel economy since these ratios are given directly by the $MRS = MRTSA$ conditions in equations (15)–(17) and do not depend on gas prices (which vary) or lifetime miles (which we do not observe). We calculate $MRS = MRTSA$ using the estimated cost parameters from model (5) in table 2 and observed choices of size and acceleration. We present these estimates in logs since the underlying distributions in levels are highly right skewed.

Figure 12: *Correlates of consumer preferences*

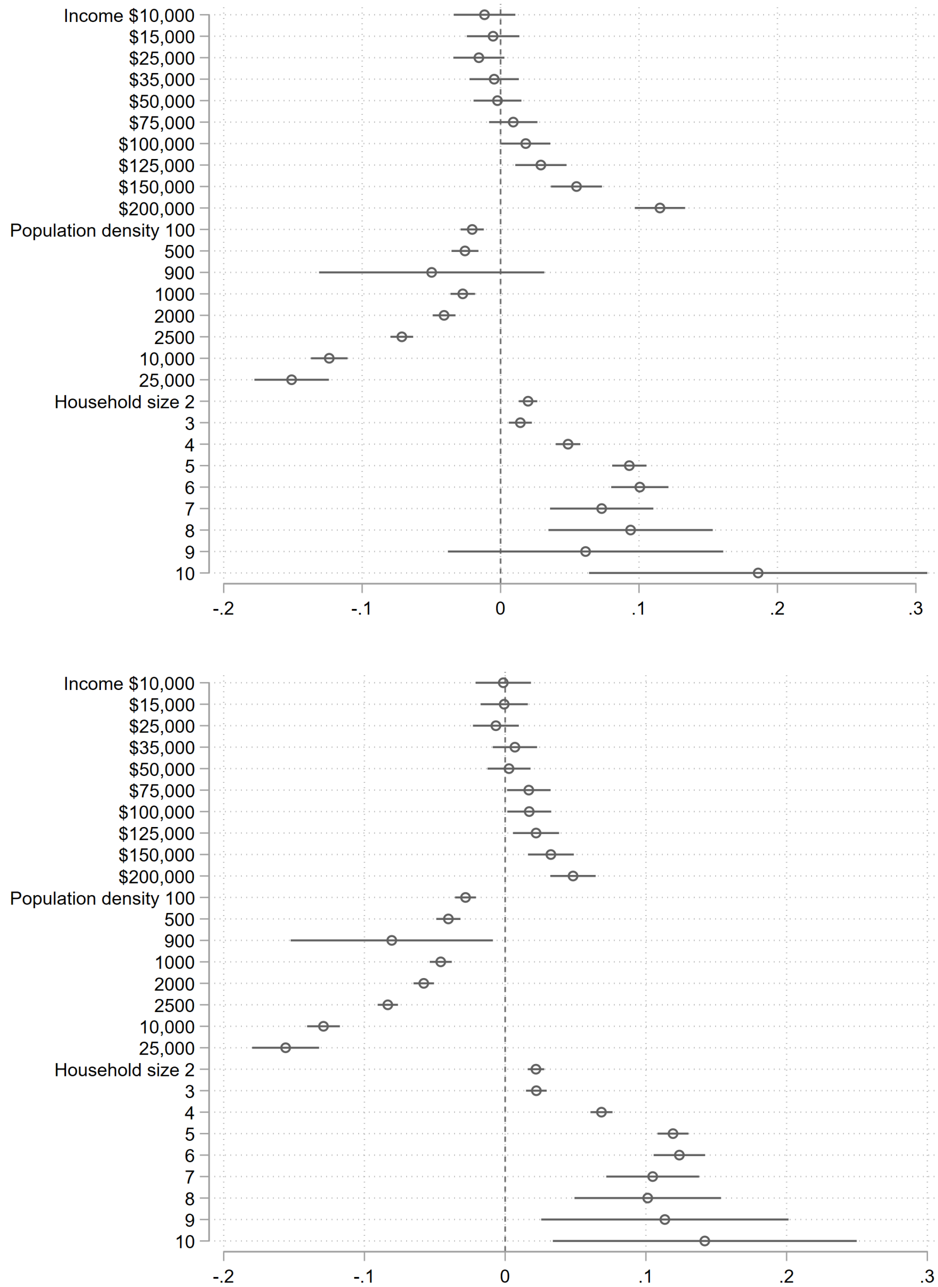

Note: This figure plots the estimated coefficients (hollow circles) and 95% confidence intervals (horizontal bars) from an OLS regression of estimated logged consumer preference ratios ($\ln \beta_a/\beta_g$ and $\ln \beta_s/\beta_g$) on dummy variables for household income, population density, and # household members. The top panel shows regression results for acceleration (dependent variable $\ln \beta_a/\beta_g$), while the bottom panel shows results for volume (dependent variable β_s/β_g). Coefficient estimates are relative to the excluded categories of income $< \$10,000$, density < 100 people per square mile, and one household member. Confidence intervals are based on standard errors clustered by state and do not account for first-stage uncertainty; i.e., they treat the dependent variable as known. Estimated preferences are based on model (5) from table 2.

Figure 13: *Estimates of shadow cost of standards from the literature*

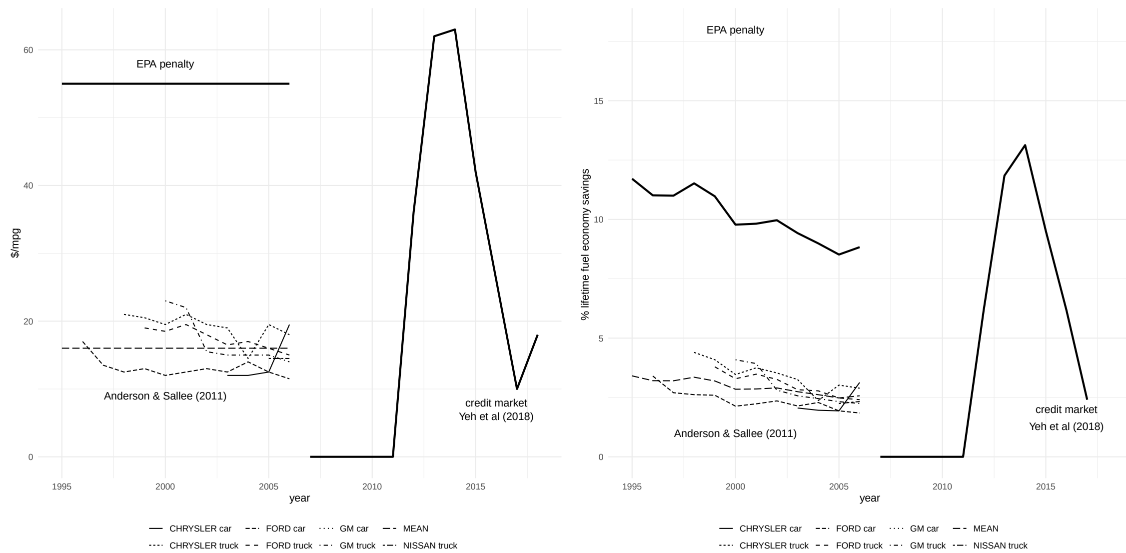

Note: The left figure shows the shadow cost estimates and the statutory penalty from fuel economy and greehouse gas (GHG) standards from the literature in dollars per mile-per-gallon (MPG). The two sources are Anderson and Sallee (2011b) and Yeh et al. (2021). The right figure shows the same as a percent of the savings from a 1 MPG improvement using 100,000 lifetime discounted miles ($m(0)/(r + \rho + \delta)$), the mean price of gasoline, and for the mean fuel economy of that year.

Figure 14: *Effect of standards on gas price variable assumptions*

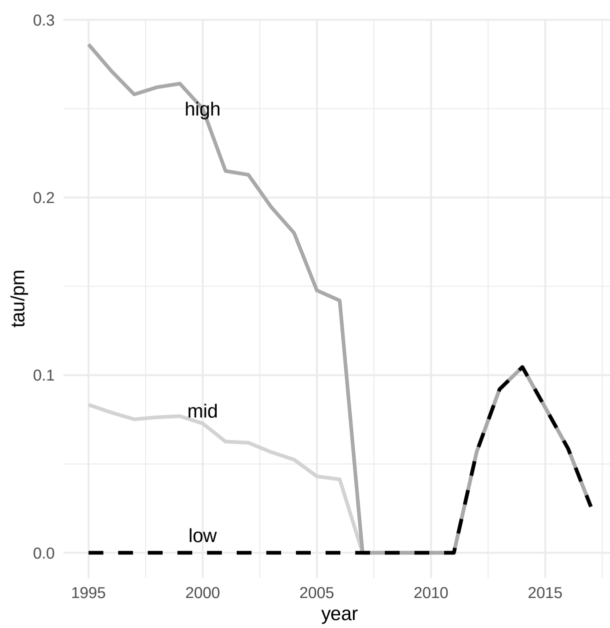

Note: The figure shows the mean of the estimated $\frac{\tau}{pm}$ for three assumptions about the period before 2007. "High" assumes the shadow price of the standards was set to the noncompliance penalty of \$55. "Mid" assumes the shadow price was the average of estimated shadow prices from Anderson and Sallee (2011b). "Low" assumes the shadow price was zero for this period.

Table 6: *Baseline regression with shadow price from standards*

	(1) Low	(2) Mid	(3) High
Log gas price + $\frac{\tau}{pm}$	0.091*	0.101*	0.115*
	(0.042)	(0.044)	(0.046)
ln hp/wt	−0.508***	−0.508***	−0.507***
	(0.017)	(0.017)	(0.017)
ln vol	−0.467***	−0.467***	−0.467***
	(0.007)	(0.007)	(0.007)
N	245 497	245 497	245 497
R^2	0.61	0.61	0.61

+ $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Note: The dependent variable is log MPG. All models replicate column (5) in table 2.

Figure 15: *Effect of shadow price on year dummies*

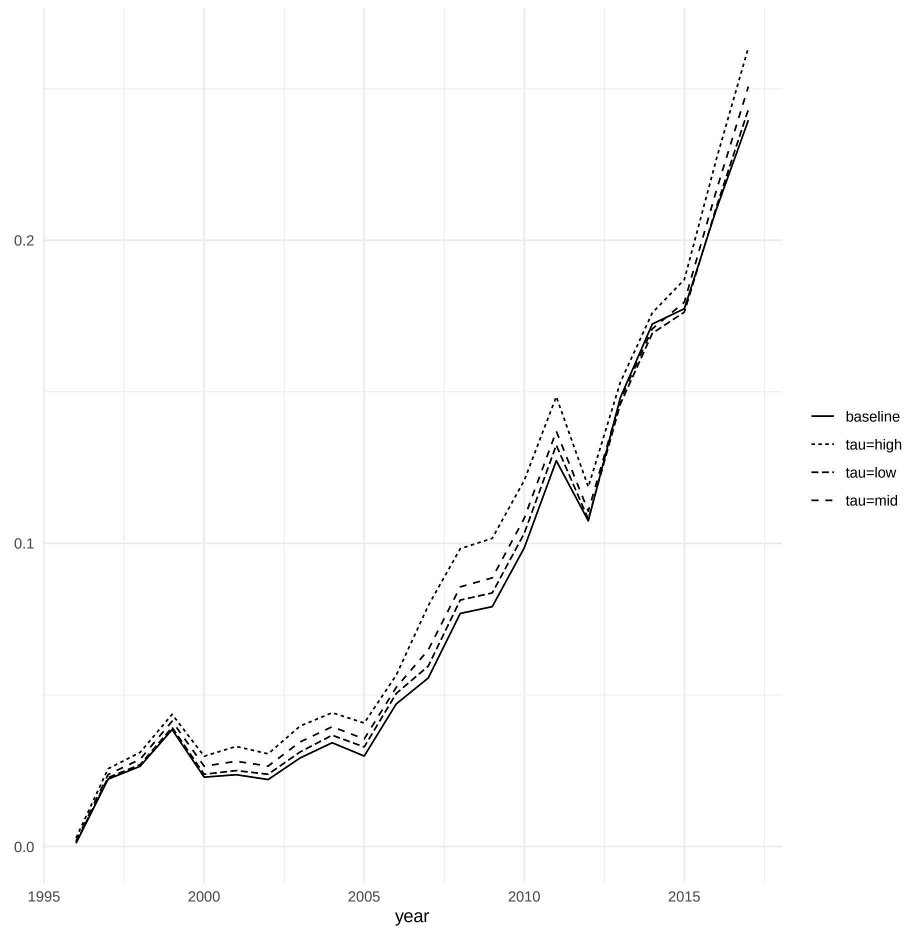

Note: This figure plots estimated coefficients on the year dummies. The baseline coefficients are from regression (6) in table 2. The remaining three lines come from regressions (1), (2), and (3) from in 6.

Table 7: *Gas price lag robustness checks*

	(1)	(2)	(3)	(4)	(5)
Log gas price$_t$	0.065*	0.034+	0.032+	0.032+	0.033+
	(0.025)	(0.019)	(0.018)	(0.017)	(0.018)
Log gas price$_{t-1}$		0.050**	0.039**	0.041**	0.041**
		(0.018)	(0.013)	(0.014)	(0.014)
Log gas price$_{t-2}$			0.022	0.035+	0.034+
			(0.026)	(0.019)	(0.017)
Log gas price$_{t-3}$				−0.033+	−0.037*
				(0.019)	(0.014)
Log gas price$_{t-4}$					0.008
					(0.027)
ln hp/wt	−0.507***	−0.507***	−0.507***	−0.507***	−0.507***
	(0.017)	(0.017)	(0.017)	(0.017)	(0.017)
ln vol	−0.467***	−0.467***	−0.467***	−0.467***	−0.467***
	(0.007)	(0.007)	(0.007)	(0.007)	(0.007)
Sum log gas price		0.084**	0.093**	0.075**	0.078**
		(0.032)	(0.043)	(0.047)	(0.055)
N	245 497	245 497	245 497	245 497	245 497
R^2	0.61	0.61	0.61	0.61	0.61

+ $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Note: The dependent variable is log MPG. Column (3) replicates column (5) in table 2 exactly. All other columns vary the number of log gas price lags but are otherwise identical.

Table 8: *Functional form robustness checks*

	(1) All years	(2) 1995–2001	(3) 2002–2009	(4) 2010–2017
ln gas price	0.096*	−0.001	−0.088	0.105***
	(0.040)	(0.084)	(0.061)	(0.017)
Log vol × tercile 1	−0.336***	−0.238***	−0.345***	−0.357***
	(0.009)	(0.019)	(0.019)	(0.010)
Log vol × tercile 2	−0.356***	−0.271***	−0.366***	−0.372***
	(0.008)	(0.019)	(0.018)	(0.009)
Log vol × tercile 3	−0.369***	−0.258***	−0.375***	−0.394***
	(0.008)	(0.017)	(0.015)	(0.010)
Log hp/wt × tercile 1	−0.552***	−0.164***	−0.615***	−0.592***
	(0.015)	(0.011)	(0.014)	(0.011)
Log hp/wt × tercile 2	−0.568***	−0.167***	−0.642***	−0.591***
	(0.014)	(0.012)	(0.014)	(0.010)
Log hp/wt × tercile 3	−0.578***	−0.153***	−0.664***	−0.597***
	(0.015)	(0.013)	(0.015)	(0.012)
Pickup	−0.065***	−0.003	−0.077***	−0.095***
	(0.007)	(0.012)	(0.007)	(0.002)
N	245 497	51 986	117 544	75 967
R^2	0.62	0.31	0.63	0.68

+ $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Note: The dependent variable is log MPG. All models replicate column (5) in table 2 while allowing the coefficient on log attributes to vary by attribute tercile. Column (2) includes a time trend interaction with attributes interacted with terciles.